

Mathematical programs with complementarity constraints and application to hyperparameter tuning for nonlinear support vector machines

Samuel Ward · Alain Zemkoho · Selin Ahipasaoglu

Received: March 2025 / unpublished preprint

Abstract We consider the Mathematical Program with Complementarity Constraints (MPCC). One of the main challenges in solving this problem is the systematic failure of standard Constraint Qualifications (CQs). Carefully accounting for the combinatorial nature of the complementarity constraints, tractable versions of the Mangasarian Fromovitz Constraint Qualification (MFCQ) have been designed and widely studied in the literature. This paper looks closely at two such MPCC-MFCQs and their influence on MPCC algorithms. As a key contribution, we prove the convergence of the sequential penalisation and Scholtes relaxation algorithms under a relaxed MPCC-MFCQ that is much weaker than the CQs currently used in the literature. We then form the problem of tuning hyperparameters of a nonlinear Support Vector Machine (SVM), a fundamental machine learning problem for classification, as a MPCC. For this application, we establish that the aforementioned relaxed MPCC-MFCQ holds under a very mild assumption. Moreover, we program robust implementations and comprehensive numerical experimentation on real-world data sets, where we show that the sequential penalisation method applied to the MPCC formulation for tuning SVM hyperparameters can outperform both the Scholtes relaxation technique and the state-of-the-art derivative-free methods from the machine learning literature.

Keywords mathematical program with complementarity constraints · hyperparameter optimisation · bilevel optimisation · penalisation · relaxation · support vector machines

Mathematics Subject Classification (2020) 49M37 · 65K05 · 65K15 · 90C30 · 90C33

The work of the first author is jointly funded by DAS Ltd and the School of Mathematical Sciences at the University of Southampton through a PhD studentship grant. The second author's research is partly funded by an EPSRC grant with reference EP/X040909/1.

Samuel Ward

ORCID: 0009-0001-3084-7099

E-mail: S.Ward@soton.ac.uk

School of Mathematical Sciences, University of Southampton, Southampton, UK

Alain Zemkoho

ORCID: 0000-0003-1265-4178

E-mail: A.B.Zemkoho@soton.ac.uk

School of Mathematical Sciences, University of Southampton, Southampton, UK

Selin Ahipasaoglu

ORCID: 0000-0003-1371-315X

E-mail: S.D.Ahipasaoglu@soton.ac.uk

School of Mathematical Sciences, University of Southampton, Southampton, UK

Contents

1	Introduction	2
2	Preliminaries on mathematical programs with complementary constraints	4
2.1	Constraint qualifications	5
2.2	Stationarity	9
2.3	Link to bilevel optimisation	10
3	Solution methods	11
3.1	Penalisation	11
3.2	Relaxation	16
4	Application to hyperparameter tuning for nonlinear support vector machines	18
4.1	Support vector machine	18
4.2	Modelling hyperparameter tuning	21
4.3	Theoretical properties of problem (\mathcal{M})	23
5	Numerical experiments on hyperparameter tuning for support vector machines	28
5.1	Implementation details	28
5.2	Starting point selection	30
5.3	Numerical behaviour of penalisation and relaxation methods	32
5.4	Performance evaluation in comparison to derivative-free optimisation	33
6	Conclusion	36
A	Table of notation	41
B	Further details of the dual SVM	41
C	Explicit formula of the Lagrangian of (\mathcal{M}) and derivatives	47
D	Semismooth Newton method	48
E	Summary of derivative-free methods	49
F	Further details of the classification datasets	50

1 Introduction

The Mathematical Program with Complementarity Constraints (MPCC) is now a well-established optimisation problem, with a wide range of applications including in economics [3], chemical engineering [4], energy distribution [32], transportation [52], and machine learning (see Section 4 for an important problem class, as well as some relevant references).

Most standard optimisation algorithms are ineffective at finding solutions of MPCCs, in part due to the systematic failure of classical Constraint Qualifications (CQs) and the combinatorial nature of the problem’s feasible set [63, Chapter 3]. To address the issue with standard CQs, many alternative tractable ones tailored to the complementarity constraints have been developed in the literature (see, e.g. [27, 75, 82, 84, 104]). Considering the importance of the Mangasarian Fromovitz Constraint Qualification (MFCQ) in optimisation, it has been particularly at the centre of attentions. Two categories of MFCQ-type CQs for the MPCC have emerged in the literature; one is based on the decomposition of the complementarity constraints, especially the constraints associated with the biactive index sets, to build suitable versions of MFCQ that can hold for MPCCs [27, 82, 83]. A second category, constructed from the lens of variational analysis, consists of reformulating the complementarity constraints either with non-smooth operator functions or sets that lead to normal cone operations that restore the fulfilment of such CQs; see, e.g., [18, 38, 104].

In the literature, the MPCC-MFCQ-T, where we use the label “T” to represent the fact that it is built from a tightened decomposition of the complementarity constraints, has emerged as one of the most used CQs for MPCCs; see, e.g., [27, 48, 82, 84]. However, looking at the MPCC-MFCQ-T very closely, one can observe that it boils down to the MPCC-LICQ, the well-known tailored linear independence constraint qualification for MPCCs, when the feasible set does not involve any inequality constraint apart from the ones in the complementarity constraints; cf. Remark 2 in Section 2. For this reason, and considering the fact that the MPCC-LICQ is very strong, the MPCC-MFCQ-T fails for many MPCC problems (see Example 1 and Proposition 4).

Alternatively, there is another version of MPCC-tailored MFCQ, studied in [79, 81], based on a relaxation of the complementarity constraints associated with the biactive index set. Hence, we denote it as MPCC-MFCQ-R with the label “R” representing the corresponding relaxation framework. We show in this paper that the MPCC-MFCQ-R is strictly weaker than the MPCC-MFCQ-T; see Theorem 1 and Example 1. Based on this observation, we revisit two classical algorithms for MPCCs that rely on MPCC-type CQs for their convergence results, namely, the

partial penalisation and Scholtes relaxation methods, which have been widely studied and used in the literature (see, e.g., [2, 57, 81] and [16, 44, 60, 83, 91], respectively).

The partial penalisation method consists of moving the product of the complementarity constraint functions from the feasible set to the objective. Two options have been pursued for this approach: (i) the *exact* penalisation, where the penalty parameter is fixed [81], and (ii) the *sequential* version, where the penalty parameter increases sequentially throughout the algorithm [57]. As it is well known in constrained optimisation, a big challenge with exact penalisation is identifying a suitable penalty parameter. Furthermore, the numerical performance of an exact penalisation method highly depends on this choice. Crucially, for sequential penalisation algorithms, the selection of the initial penalisation parameter is not too critical, as theoretically, under suitable assumptions, the algorithm gets to a point of stability where a solution is found.

Considering this practical advantage of the sequential approach, we focus on it in the analysis in this paper. The article [57] shows that the sequential partial penalisation algorithm, in conjunction with an interior-point method for the subproblems, converges to an S-stationary point, under the MPCC-LICQ, as a key assumption. In this paper, we establish the convergence of the sequential partial penalisation method under the MPCC-MFCQ-R, irrespective of the algorithm for solving the corresponding subproblems; see Theorem 3. Not only is MPCC-MFCQ-R far weaker than the MPCC-LICQ, but the subproblems can be solved with any method, and not necessarily with an interior-point method, as required in [57].

As for the Scholtes relaxation algorithm, its convergence has been established under the MPCC-MFCQ-T; see, e.g., [84, page 112, Theorem 7.3]. Here, we show that the MPCC-MFCQ-R is sufficient to establish this convergence; cf. Theorem 5. Moreover, to make the method more practically relevant, in our proof, we only require the computation of approximate Karush-Kuhn-Tucker (KKT) points for the relaxed problem at each iteration. This is a significant departure from the existing convergence results, where it is usually assumed that the KKT points of the relaxed problem are computed exactly, something that is typically not possible in a practical implementation.

Subsequently, we explore the application of our aforementioned results for general MPCCs to the problem of computing optimal hyperparameters for a Support Vector Machine (SVM). Note that the SVM represents one of the most potent tools for classification in machine learning. Since their introduction in 1995 by Cortes and Vapnik [15, 96], it has been prolifically used to solve multiple practical problems, including facial authentication [94], protein sequencing [80] and speech recognition [87]; it also sets a competitive benchmark in all classification problems.

One of the biggest challenges with SVMs is choosing their hyperparameters. Unlike model parameters, which are learnt from patterns in the training data, hyperparameters define the learning process itself and so must be considered separately. The two most common hyperparameters are the balance of empirical loss *regularisation parameter* and the *choice of kernel*. There is no one model that will be best suited to all scenarios [102]. Traditionally, in the machine learning literature, grid search is used to select the value that provides the best validation accuracy [66]. The method considers a sequence of candidate choices for each hyperparameter and then forms a grid of every combination. Each coordinate of the grid (representing a choice for each hyperparameter) requires a new model to be trained and evaluated. This scales very poorly with multiple hyperparameters and is only as precise as the spacing between grid points. Many other direct (general, non-convex, gradient-free) search methods have been employed to tune hyperparameters. For example, random search [37, 65], Bayesian optimisation [34, 49, 77, 88], Nelder-Mead simplex [71] and genetic algorithms [54]. A good survey of these approaches can be found in [98]. However, these methods all assume that training is a black box – that a model must be trained in its entirety before being evaluated. They are general methods that neglect the specific structure of the SVM problem.

More recently, a new approach to hyperparameter selection using bilevel optimisation has been proposed [31, 59, 74]. It enables the simultaneous training of the model while hyperparameter tuning is being conducted. The bilevel optimisation technique has been applied thoroughly and with great success to linear models. For example, the series of papers [67, 68, 69] proposes some non-smooth optimisation methods for selecting linear support vector regression and classification hyperparameters. Other authors have created a similar framework but propose different approaches to solving the bilevel program; e.g., a value function-based descent, global relaxation, and stochastic gradient descent methods have been proposed in [33], [58] and [47], respectively. The aforementioned works

are on linear models; thus, enabling some mathematical simplicity, but such a framework is less useful in practice, as real-world data sets are rarely linearly separable.

Of particular interest to us is the sequence of papers [6, 50, 51], as they suggest an extension of the bilevel optimisation framework to nonlinear kernel models, while proposing an MPCC reformulation for the corresponding problem. However, they do not provide much theoretical insight into the problem. The paper [14] begins addressing this gap. Our paper expands on the latter work by first rigorously constructing the MPCC reformulation for the bilevel optimisation problem of tuning the hyperparameters of nonlinear kernel SVMs. We then prove for the first time that the MPCC-MFCQ-R automatically holds at any feasible point of the MPCC reformulation of the bilevel hyperparameter optimisation problem, provided that the corresponding pair of hyperparameters is positive; cf. Proposition 3. Note that this positiveness is not a strong requirement, as for SVMs, those values cannot be zero in practice. As a consequence of this result, the sequential penalisation and Sholtes relaxation algorithms, discussed above, converge automatically, almost for free, for the MPCC reformulation of the problem. It is important to note that we also show that for the MPCC reformulation of our bilevel hyperparameter optimisation problem, the MPCC-MFCQ-T fails at any of its feasible points with at least two biactive indices; see Proposition 4.

As another key contribution of this paper, we conduct a robust implementation of the partial penalisation and Scholtes relaxation methods and compare them to each other, and also to classical hyperparameter tuning algorithms in the machine learning literature. First, we observe that the sequential penalisation method, on average, not only runs faster but also converges to a stationary point with a lower objective function value than the Scholtes relaxation algorithm. Moreover, we show that the sequential partial penalisation technique can outperform the traditional derivative-free methods such as grid, random, Bayesian and pattern search. Note that for the experiments, we use 18 well-known classification datasets from the literature (see Table 1 for references), where the median problem size is 1773 variables, 1015 constraints (excluding complementarity constraints) and 1010 pairs of complementarity constraints. To the best of our knowledge, these problems are by far larger than any other that have been used for experiments in the context of non-convex MPCCs. For illustration, observe that for the MacMPEC collection [55], which is classical for MPCC numerical experimentation, the median problem size is 49 variables, 20 constraints (excluding complementarity constraints) and 12 pairs of complementarity constraints.

For the remainder of the paper, note that Section 2 covers preliminaries with particular attention on the difference between the two MPCC-tailored CQs that are a key focus in this paper; i.e., the MPCC-MFCQ-T and MPCC-MFCQ-R. In Section 3, we present the sequential inexact penalisation and relaxation methods and prove new convergence results which allow these methods to be applicable to a broader range of problems. In Section 4, we construct the bilevel optimisation-based hyperparameter tuning model for SVMs and prove that its KKT/MPCC reformulation satisfies MPCC-MFCQ-R but not MPCC-MFCQ-T. Finally, Section 5 presents our numerical experiments and a demonstration that the MPCC model solved with sequential penalisation can outperform traditional derivative-free methods from the machine learning literature. Some conclusions are given in Section 6. Many technical details that could further facilitate the comprehension of some aspects of the paper, especially developments related to the SVM and the corresponding hyperparameter tuning problem, are provided in the appendices.

2 Preliminaries on mathematical programs with complementary constraints

For a number of concepts used throughout the paper, we refer to a general Nonlinear Program (NLP) defined, for given dimensions $n, p, q \in \mathbb{N}$, as

$$\begin{aligned} & \underset{z \in \mathbb{R}^n}{\text{minimise}} && f(z) \\ & \text{subject to} && g_i(z) \geq 0 && \text{for } i = 1, \dots, p, \\ & && h_i(z) = 0 && \text{for } i = 1, \dots, q, \end{aligned} \tag{NLP}$$

where the functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$, $g : \mathbb{R}^n \rightarrow \mathbb{R}^p$, $h : \mathbb{R}^n \rightarrow \mathbb{R}^q$ representing the objective, inequality, and equality constraint functions, respectively, are continuously differentiable. However, the main

focus of our analysis is the Mathematical Program with Complementary Constraints (MPCC), defined for given dimensions $n, p, q, r \in \mathbb{N}$, as

$$\begin{aligned} & \underset{z \in \mathbb{R}^n}{\text{minimise}} && f(z) \\ & \text{subject to} && g_i(z) \geq 0 && \text{for } i = 1, \dots, p, \\ & && h_i(z) = 0 && \text{for } i = 1, \dots, q, \\ & && G_i(z) \geq 0, H_i(z) \geq 0, G_i(z)H_i(z) = 0 && \text{for } i = 1, \dots, r, \end{aligned} \quad (\text{MPCC})$$

where the functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$, $g : \mathbb{R}^n \rightarrow \mathbb{R}^p$, $h : \mathbb{R}^n \rightarrow \mathbb{R}^q$ and $G, H : \mathbb{R}^n \rightarrow \mathbb{R}^r$ are continuously differentiable. G and H are referred to as complementary constraint functions and $G_i(z)H_i(z)$ as the product term. The last line of problem (MPCC) can be written compactly as $0 \leq G(z) \perp H(z) \leq 0$.

2.1 Constraint qualifications

In continuous constrained optimisation, constraint qualifications (CQs) are usually required to construct necessary optimality conditions and prove the convergence of numerical algorithms. To introduce a constraint qualification for problem (NLP), we define the *index sets*

$$I^g(\bar{z}) := \{i \in \{1, \dots, p\} : g_i(\bar{z}) = 0\} \quad \text{and} \quad I^h := \{1, \dots, q\}.$$

Based on this notation, the *Linear Independence Constraint Qualification* (LICQ) is satisfied at a feasible point $\bar{z} \in \mathbb{R}^n$ of the (NLP) iff the family of gradients

$$\{\nabla h_i(\bar{z}) : i \in I^h\} \cup \{\nabla g_i(\bar{z}) : i \in I^g(\bar{z})\} \quad \text{is linearly independent,} \quad (\text{LICQ})$$

while, the *Mangasarian-Fromovitz Constraint Qualification* (MFCQ) will be said to hold at a feasible point $\bar{z} \in \mathbb{R}^n$ of the (NLP) iff

$$\begin{aligned} & \{\nabla h_i(\bar{z}) : i \in I^h\} \quad \text{is linearly independent and} \\ & \exists d \in \mathbb{R}^n : \quad \begin{cases} \nabla h_i(\bar{z})^\top d = 0 & \forall i \in I^h, \\ \nabla g_i(\bar{z})^\top d > 0 & \forall i \in I^g(\bar{z}). \end{cases} \end{aligned} \quad (\text{MFCQ})$$

Obviously, the LICQ is stronger than the MFCQ. These CQs play a fundamental role in nonlinear constrained optimisation thanks to their versatility in deriving fundamental properties for problem (NLP); see the excellent tutorial in [89] for an overview of these properties.

Unfortunately, it is well known that LICQ and MFCQ are systematically violated at every feasible point of problem (MPCC) when viewed as a constrained problem with equality and inequality constraints in the format described in (NLP) [13]. To observe this, consider an index $i = 1, \dots, r$, a feasible point \bar{z} and the three constraints $G_i(\bar{z}) \geq 0$, $H_i(\bar{z}) \geq 0$ and $G_i(\bar{z})H_i(\bar{z}) = 0$. Assuming, without loss of generality, that $H_i(\bar{z})$ is zero, the gradient of the second constraint $\nabla H_i(\bar{z})$ and the gradient of the third constraint $G_i(\bar{z})\nabla H_i(\bar{z})$ (by the product rule) are not linearly independent.

The main challenge in addressing problem (MPCC) comes from the combinatorial nature of the complementarity constraints. For each index $i = 1, \dots, r$, either H_i or G_i may be positive, but not both. To further study this, we partition these cases by defining the following index sets:

$$I^G(\bar{z}) := \{i \in \{1, \dots, r\} : G_i(\bar{z}) = 0, H_i(\bar{z}) > 0\}, \quad (1a)$$

$$I^H(\bar{z}) := \{i \in \{1, \dots, r\} : G_i(\bar{z}) > 0, H_i(\bar{z}) = 0\}, \quad (1b)$$

$$I^{GH}(\bar{z}) := \{i \in \{1, \dots, r\} : G_i(\bar{z}) = 0, H_i(\bar{z}) = 0\}. \quad (1c)$$

Based on these index sets, at a given feasible point \bar{z} of problem (MPCC), we first introduce the corresponding *Tightened Nonlinear Program* (TNLP(\bar{z})):

$$\begin{aligned}
& \underset{z \in \mathbb{R}^n}{\text{minimise}} && f(z) \\
& \text{subject to} && g_i(z) \geq 0 && \text{for } i = 1, \dots, p, \\
& && h_i(z) = 0 && \text{for } i = 1, \dots, q, \\
& && G_i(z) = 0, \quad H_i(z) \geq 0 && \text{for } i \in I^G(\bar{z}), \\
& && G_i(z) \geq 0, \quad H_i(z) = 0 && \text{for } i \in I^H(\bar{z}), \\
& && G_i(z) = 0, \quad H_i(z) = 0 && \text{for } i \in I^{GH}(\bar{z}).
\end{aligned} \tag{TNLP(\bar{z})}$$

It is very important to remember that these index sets depend on the point \bar{z} , so the (TNLP(\bar{z})) may be a different program when constructed around a different point. The decomposition of the complementarity constraints in (MPCC) makes the fulfilment of standard CQs more likely. Hence, problem (TNLP(\bar{z})) is commonly used to define CQs tailored to problem (MPCC) as follows.

Definition 1 The MPCC-MFCQ-T holds at a feasible point \bar{z} of the (MPCC) if the corresponding tightened program (TNLP(\bar{z})) satisfies the MFCQ at the same point \bar{z} . Equivalently, the MPCC-MFCQ-T holds at a feasible point \bar{z} of the (MPCC) iff

$$\begin{aligned}
& \left\{ \nabla h_i(\bar{z}) : i \in I^h \right\} \cup \left\{ \nabla G_i(\bar{z}) : i \in I^G(\bar{z}) \cup I^{GH}(\bar{z}) \right\} \cup \left\{ \nabla H_i(\bar{z}) : i \in I^H(\bar{z}) \cup I^{GH}(\bar{z}) \right\} \\
& \text{is linearly independent and}
\end{aligned} \tag{2}$$

$$\exists d \in \mathbb{R}^n : \begin{cases} \nabla h_i(\bar{z})^\top d = 0 & \forall i \in I^h, \\ \nabla G_i(\bar{z})^\top d = 0 & \forall i \in I^G(\bar{z}) \cup I^{GH}(\bar{z}), \\ \nabla H_i(\bar{z})^\top d = 0 & \forall i \in I^H(\bar{z}) \cup I^{GH}(\bar{z}), \\ \nabla g_i(\bar{z})^\top d > 0 & \forall i \in I^g(\bar{z}). \end{cases} \tag{3}$$

Definition 1 has been widely used in the literature (see, e.g., [27, Definition 3.1], [48, Definition 4], [82, page 4] and [84, Definition 5.4]) and is usually referred to as the MPCC-MFCQ. However, we denote this as MPCC-MFCQ-T to emphasise the fact that it is based on the tightened program. In addition, this avoids conflict with the alternative definition that we will present next.

An alternative MPCC-tailored MFCQ relies on the Relaxed Nonlinear Program (RNLP(\bar{z})) defined, for a given feasible point \bar{z} of problem (MPCC), as

$$\begin{aligned}
& \underset{z \in \mathbb{R}^n}{\text{minimise}} && f(z) \\
& \text{subject to} && g_i(z) \geq 0 && \text{for } i = 1, \dots, p, \\
& && h_i(z) = 0 && \text{for } i = 1, \dots, q, \\
& && G_i(z) = 0, \quad H_i(z) \geq 0 && \text{for } i \in I^G(\bar{z}), \\
& && G_i(z) \geq 0, \quad H_i(z) = 0 && \text{for } i \in I^H(\bar{z}), \\
& && G_i(z) \geq 0, \quad H_i(z) \geq 0 && \text{for } i \in I^{GH}(\bar{z}).
\end{aligned} \tag{RNLP(\bar{z})}$$

This problem is similar to the tightened program but for the biactive indices $i \in I^{GH}(\bar{z})$, where inequality constraints rather than equality constraints are added for G_i and H_i . Analogously, we use the notation (RNLP(\bar{z})) to make clear that this program is defined at a feasible point \bar{z} . This leads to a second, different MPCC-tailored constraint qualification.

Definition 2 The MPCC-MFCQ-R holds at a feasible point \bar{z} of the (MPCC) iff the corresponding relaxed program (RNLP(\bar{z})) satisfies the MFCQ at the same point \bar{z} . Equivalently, the MPCC-

MFCQ-R holds at a feasible point \bar{z} of the (MPCC) iff

$$\left\{ \nabla h_i(\bar{z}) : i \in I^h \right\} \cup \left\{ \nabla G_i(\bar{z}) : i \in I^G(\bar{z}) \right\} \cup \left\{ \nabla H_i(\bar{z}) : i \in I^H(\bar{z}) \right\}$$

is linearly independent and

$$\exists d \in \mathbb{R}^n : \begin{cases} \nabla h_i(\bar{z})^\top d = 0 & \forall i \in I^h, \\ \nabla G_i(\bar{z})^\top d = 0 & \forall i \in I^G(\bar{z}), \\ \nabla H_i(\bar{z})^\top d = 0 & \forall i \in I^H(\bar{z}), \\ \nabla g_i(\bar{z})^\top d > 0 & \forall i \in I^g(\bar{z}), \\ \nabla G_i(\bar{z})^\top d > 0, \nabla H_i(\bar{z})^\top d > 0 & \forall i \in I^{GH}(\bar{z}). \end{cases} \quad (5)$$

Definition 2 is less common. It can be found in [79, Definition 2.5] and [81, Definition 2.5]. These papers call it MPCC-MFCQ as well, but to avoid confusion, we will call it MPCC-MFCQ-R, denoting that it is based on the relaxed program. We next provide the statement of MPCC-LICQ.

Remark 1 The MPCC-LICQ holds at point \bar{z} of problem (MPCC) iff any of the following three equivalent statements holds:

1. The corresponding tightened program (TNLP(\bar{z})) satisfies the LICQ at the same point \bar{z} .
2. The corresponding relaxed program (RNLP(\bar{z})) satisfies the LICQ at the same point \bar{z} .
3. The following family of vectors is linearly independent:

$$\left\{ \nabla h_i(\bar{z}) : i \in I^h \right\} \cup \left\{ \nabla g_i(\bar{z}) : i \in I^g(\bar{z}) \right\} \cup \left\{ \nabla G_i(\bar{z}) : i \in I^G(\bar{z}) \cup I^{GH}(\bar{z}) \right\} \cup \left\{ \nabla H_i(\bar{z}) : i \in I^H(\bar{z}) \cup I^{GH}(\bar{z}) \right\}. \quad (6)$$

This means that the concept of the MPCC-LICQ is the same regardless of whether it is defined from the tightened or relaxed program.

Observe that the MPCC-MFCQ-T and MPCC-MFCQ-R are defined above in primal space due to the involvement of feasible directions. Next, we introduce the dual forms of these CQs, which are usually more handy in practice. To proceed, consider two sets of vectors $A := \{a_i : i \in I_1\}$ and $B := \{b_i : i \in I_2\}$. We say that (A, B) is *positive-linearly dependent* iff there exist scalars $\{\lambda_i\}_{i \in I_1}$ and $\{\mu_i\}_{i \in I_2}$, not all of them zero, such that

$$\sum_{i \in I_1} \lambda_i a_i + \sum_{i \in I_2} \mu_i b_i = 0, \text{ where } \lambda_i \geq 0 \text{ for } i \in I_1 \text{ and } \mu_i \text{ free for } i \in I_2.$$

If no such scalars exist, we say (A, B) is *positive-linearly independent*. The concept of positive-linear independence, which might have been introduced in [78, page 965, Definition 2.1], has since been widely used, in particular in the MPCC literature; see, e.g., [44, page 260, Definition 2.1] and [84, page 63, Definition 4.5]. We now state a lemma that will be useful for many of the proofs below.

Lemma 1 *The following statements hold true:*

1. The MFCQ holds at a feasible point \bar{z} of problem (NLP) if and only if (A, B) is positive-linearly independent where:

$$A = \{\nabla g_i(\bar{z}) : i \in I^g(\bar{z})\}, \quad B = \{\nabla h_i(\bar{z}) : i \in I^h\}.$$

2. The MPCC-MFCQ-T is satisfied at a feasible point \bar{z} of problem (MPCC) if and only if (A, B) is positive-linearly independent where:

$$A = \{\nabla g_i(\bar{z}) : i \in I^g(\bar{z})\},$$

$$B = \{\nabla h_i(\bar{z}) : i \in I^h\} \cup \{\nabla G_i(\bar{z}) : i \in I^G(\bar{z}) \cup I^{GH}(\bar{z})\} \cup \{\nabla H_i(\bar{z}) : i \in I^H(\bar{z}) \cup I^{GH}(\bar{z})\}.$$

3. The MPCC-MFCQ-R is satisfied at a feasible point \bar{z} of problem (MPCC) if and only if (A, B) is positive-linearly independent where:

$$A = \{\nabla g_i(\bar{z}) : i \in I^g(\bar{z})\} \cup \{\nabla G_i(\bar{z}), \nabla H_i(\bar{z}) : i \in I^{GH}(\bar{z})\},$$

$$B = \{\nabla h_i(\bar{z}) : i \in I^h\} \cup \{\nabla G_i(\bar{z}) : i \in I^G(\bar{z})\} \cup \{\nabla H_i(\bar{z}) : i \in I^H(\bar{z})\}. \quad (7)$$

Proof Direct results from applying Motzkin's theorem of the alternative [64, page 27] to the definitions of MFCQ, MPCC-MFCQ-T and MPCC-MFCQ-R for the corresponding problems. \square

We have the following sequence of implications, where the first is well known [27, Corollary 3.2]. The second implication is presented for the first time. We include both proofs for completeness.

Theorem 1 *For any feasible point z of problem (MPCC), the following implications hold:*

$$\text{MPCC-LICQ} \implies \text{MPCC-MFCQ-T} \implies \text{MPCC-MFCQ-R}$$

Proof For the first implication, assume that the MPCC-LICQ holds at \bar{z} , but the MPCC-MFCQ-T does not hold at this point. By Lemma 1, we have the positive-linear dependence of the two sets of vectors given in (2). That is, there exist scalars $\hat{\lambda} \geq 0$ and μ, ν, ξ free, not all zero, such that

$$\sum_{i \in I^g(\bar{z})} \hat{\lambda}_i \nabla g_i(\bar{z}) + \sum_{i=1, \dots, q} \mu_i \nabla h_i(\bar{z}) + \sum_{i \in I^G(\bar{z}) \cup I^{GH}(\bar{z})} \nu_i \nabla G_i(\bar{z}) + \sum_{i \in I^H(\bar{z}) \cup I^{GH}(\bar{z})} \xi_i \nabla H_i(\bar{z}) = 0.$$

But this contradicts the linear independence of (6). Therefore, MPCC-MFCQ-T must hold.

For the second implication, assume that the MPCC-MFCQ-R does not hold at \bar{z} . By Lemma 1, we can conclude the positive-linear dependence of (7). That is there exist scalars $\hat{\lambda}, \hat{\nu}, \hat{\xi} \geq 0$ and μ, ν, ξ free, not all of them zero, such that

$$\begin{aligned} & \sum_{i \in I^g(\bar{z})} \hat{\lambda}_i \nabla g_i(\bar{z}) + \sum_{i \in I^{GH}(\bar{z})} \hat{\nu}_i \nabla G_i(\bar{z}) + \sum_{i \in I^{GH}(\bar{z})} \hat{\xi}_i \nabla G_i(\bar{z}) \\ & + \sum_{i=1, \dots, q} \mu_i \nabla h_i(\bar{z}) + \sum_{i \in I^G(\bar{z})} \nu_i \nabla G_i(\bar{z}) + \sum_{i \in I^H(\bar{z})} \xi_i \nabla H_i(\bar{z}) = 0. \end{aligned}$$

This contradicts the positive-linear independence of (2). Therefore, MPCC-MFCQ-T also fails. \square

MPCC-MFCQ-R is strictly weaker than MPCC-MFCQ-T as shown in the following example.

Example 1 Consider the following (MPCC):

$$\begin{aligned} & \underset{z \in \mathbb{R}^2}{\text{minimise}} && z_1 + z_2 \\ & \text{subject to} && z_1 + z_2 \geq 0, \\ & && 0 \leq z_1 \perp z_2 \geq 0 \end{aligned}$$

with $g(z) := z_1 + z_2$, $G(z) := z_1$ and $H(z) := z_2$. At the point $\bar{z} = [0, 0]^\top$, we are in the bi-active case where $G(\bar{z}) = H(\bar{z}) = 0$. Consider the derivatives of the constraints: $\nabla g(\bar{z}) = [1, 1]^\top$, $\nabla G(\bar{z}) = [1, 0]^\top$ and $\nabla H(\bar{z}) = [0, 1]^\top$. Choosing, for example, $d = [3, 4]^\top$, we get $d^\top \nabla g(\bar{z}) = 7 > 0$ and $d^\top \nabla G(\bar{z}) = 3 > 0$ and $d^\top \nabla H(\bar{z}) = 4 > 0$. Thus, the MPCC-MFCQ-R holds at \bar{z} . However, there does not exist any direction $d \in \mathbb{R}^2$ such that $d^\top \nabla g(\bar{z}) > 0$ and $d^\top \nabla G(\bar{z}) = 0$ and $d^\top \nabla H(\bar{z}) = 0$. This implies the failure of the MPCC-MFCQ-T at \bar{z} .

Another important example that satisfies MPCC-MFCQ-R but not MPCC-MFCQ-T is the main topic of Section 4. For further evidence of the very strong nature of the MPCC-MFCQ-T, we end on the following remark.

Remark 2 For any problem of the form (MPCC) with $p = 0$ (in other words with no inequality $g_i(z) \leq 0$ constraints but still possibly with equality $h_i(x) = 0$ and complementarity $0 \leq H_i(z) \perp G_i(z) \leq 0$ constraints), the MPCC-MFCQ-T is equivalent to the MPCC-LICQ.




Fig. 1: This diagram highlights the differences in the W, A, C, M and S stationary conditions. Areas coloured in red show possible values the multipliers ν_i and ξ_i may take for a stationary point \bar{z} where $G_i(\bar{z}) = H_i(\bar{z}) = 0$

2.2 Stationarity

This subsection states some well-known stationarity concepts for problem (MPCC), which have been extensively studied in the literature; see, e.g. [82], which seems to be one of the early papers on the subject. We begin by introducing a problem-specific Lagrangian function

$$\mathcal{L}(z, \lambda, \mu, \nu, \xi) := f(z) - \sum_{i=1}^p \lambda_i g_i(z) - \sum_{i=1}^q \mu_i h_i(z) - \sum_{i=1}^r \nu_i G_i(z) - \sum_{i=1}^r \xi_i H_i(z). \quad (8)$$

Observe that this function does not include the product terms $G_i(z)H_i(z)$ for $i = 1, \dots, r$.

Definition 3 The feasible point \bar{z} of problem (MPCC) is said to be Weakly (W)-stationary if there exist multipliers (λ, μ, ν, ξ) such that the following conditions hold:

$$\nabla_z \mathcal{L}(\bar{z}, \lambda, \mu, \nu, \xi) = 0 \quad (9a)$$

$$\lambda_i \geq 0 \quad \text{for } i \notin I^g(\bar{z}), \quad (9b)$$

$$\lambda_i = 0 \quad \text{for } i \in I^g(\bar{z}), \quad (9c)$$

$$\xi_i = 0 \quad \text{for } i \in I^G(\bar{z}), \quad (9d)$$

$$\nu_i = 0 \quad \text{for } i \in I^H(\bar{z}). \quad (9e)$$

The point \bar{z} will be said to be Alternative (A), Clarke (C), Mordukhovich (M) or Strongly (S)-stationary if additionally to (9a-e), we respectively have the conditions

$$\text{A-stationary:} \quad \nu_i \geq 0 \text{ or } \xi_i \geq 0 \quad \text{for } i \in I^{GH}(\bar{z}); \quad (9A)$$

$$\text{C-stationary:} \quad \nu_i \xi_i \geq 0 \quad \text{for } i \in I^{GH}(\bar{z}); \quad (9C)$$

$$\text{M-stationary:} \quad \nu_i, \xi_i \geq 0 \text{ or } \nu_i \xi_i = 0 \quad \text{for } i \in I^{GH}(\bar{z}); \quad (9M)$$

$$\text{S-stationary:} \quad \nu_i, \xi_i \geq 0 \quad \text{for } i \in I^{GH}(\bar{z}). \quad (9S)$$

It is important to note a few properties of these stationary conditions. They only differ by the nature of multipliers on the biactive index set I^{GH} . Figure 1 highlights these differences. So, if the problem is strictly complementary at \bar{z} , that is $G_i(z) > 0, H_i(z) = 0$ or $G_i(z) = 0, H_i(z) > 0$ for all $i = 1, \dots, r$, then all five concepts are equivalent. S-stationary is the strongest and is equivalent to the KKT conditions of the (MPCC) if the complementary constraints are viewed as ordinary inequality and equality constraints. The following relationships hold [27, page 60]:

$$\begin{array}{ccccccc} \text{S-stationary} & \implies & \text{M-stationary} & \begin{array}{l} \implies \\ \implies \end{array} & \begin{array}{l} \text{A-stationary} \\ \text{C-stationary} \end{array} & \begin{array}{l} \implies \\ \implies \end{array} & \text{W-stationary} \end{array}$$

2.3 Link to bilevel optimisation

As our support vector machine (SVM) application problem in this paper is primarily a bilevel optimisation problem, it is important to introduce the definition of a bilevel optimisation program and present its relationship to the MPCC. This is exactly what we do in this subsection.

First, recall that bilevel optimisation dates back to Stackelberg's economic leader-follower game theory [90], where the leader first gets to select variable, that we denote here by $z_1 \in \mathbb{R}^n$, to optimise their upper-level objective function $F(z_1, z_2)$ subject to some upper-level constraints $G(z_1) \geq 0$ and $H(z_1) = 0$. Whatever selection the leader makes for z_1 , the follower will then get to choose their own variable, represented here by $z_2 \in \mathbb{R}^m$, to optimise their lower-level objective function $f(z_1, z_2)$ subject to some lower-level constraints $g(z_1, z_2) \geq 0$ and $h(z_1, z_2) = 0$. However, as the leader's choices depend on the follower's variable z_2 , the upper-level player has to carefully choose their strategy. More precisely, given dimensions $n, m, p, q, r, s \in \mathbb{N}$ and differentiable functions F, f, G, H, g , and h from domain \mathbb{R}^{n+m} to range $\mathbb{R}, \mathbb{R}, \mathbb{R}^p, \mathbb{R}^q, \mathbb{R}^r$, and \mathbb{R}^s , respectively, the described *Bilevel Optimisation Program* can be written as

$$\begin{aligned} & \underset{z_1, z_2}{\text{minimise}} && F(z_1, z_2) \\ & \text{subject to} && G(z_1) \geq 0, \\ & && H(z_1) = 0, \\ & && z_2 \in S(z_1) := \underset{z_2}{\operatorname{argmin}} \{f(z_1, z_2') : g(z_1, z_2') \geq 0, h(z_1, z_2') = 0\}, \end{aligned} \tag{BP}$$

where the lower-level problem corresponds to the following program parametrised by z_1 :

$$\underset{z_2'}{\text{minimise}} f(z_1, z_2') \quad \text{subject to} \quad g(z_1, z_2') \geq 0, h(z_1, z_2') = 0. \tag{10}$$

Problem (BP) has a wide range of applications (see, e.g., [19] for more details) and our SVM hyperparameter tuning problem in Section 4 is one its practical applications.

General NLP assumptions and CQs can be applied to the lower-level program. For instance, we say that problem (BP) has a convex lower-level program if and only if for all $z_1 \in \mathbb{R}^n$, the functions $f(z_1, \cdot)$ and $g_i(z_1, \cdot)$, for $i = 1, \dots, q$, are convex and $h(z_1, \cdot)$ is affine linear. Additionally, we say that the *Lower-Level Linear Independence Constraint Qualification* (LLICQ) and respectively the *Lower-Level Mangasarian–Fromovitz Constraint Qualification* (LMFCQ) are satisfied at a point (z_1, z_2) if and only if (LICQ) and respectively (MFCQ) are satisfied at the point z_2 for the lower-level problem (10) parametrised by z_1 .

A common technique towards solving problem (BP) is to replace the lower-level problem with its Karush–Kuhn–Tucker (KKT) conditions, leading to the problem

$$\begin{aligned} & \underset{z_1, z_2, \lambda, \mu}{\text{minimise}} && F(z_1, z_2) \\ & \text{subject to} && G(z_1) \geq 0, \\ & && H(z_1) = 0, \\ & && \begin{cases} \nabla_{z_2} f(z_1, z_2) - \lambda^\top \nabla_{z_2} g(z_1, z_2) - \mu^\top \nabla_{z_2} h(z_1, z_2) = 0, \\ 0 \leq \lambda \perp g(z_1, z_2) \geq 0, \\ h(z_1, z_2) = 0, \end{cases} \end{aligned} \tag{KKTR}$$

where $\lambda \in \mathbb{R}^r$ and $\mu \in \mathbb{R}^s$ correspond to the Lagrange multiplier associated with the lower-level constrain functions $g(z_1, z_2)$ and $h(z_1, z_2)$ at the point (z_1, z_2) such that z_1 is the upper-level selection and z_2 is the corresponding lower-level optimal solution.

Clearly, problem (KKTR) is a special class of problem (MPCC). However, problem (KKTR) is not necessarily equivalent to the original problem (BP), as shown in [17]. Nevertheless, under suitable assumptions, the following global and local relationships are well-known.

Theorem 2 ([17]) *Let the lower-level problem in (BP) be convex.*

(i) *Let (z_1^*, z_2^*) be a global (resp. local) optimal solution of problem (BP) and assume that the LMFCQ holds at (z_1^*, z_2^*) . Then, for each choice of Lagrange multipliers*

$$(\lambda^*, \mu^*) \text{ such that } \begin{cases} \nabla_{z_2^*} f(z_1^*, z_2^*) - \lambda^{*\top} \nabla_{z_2^*} g(z_1^*, z_2^*) - \mu^{*\top} \nabla_{z_2^*} h(z_1^*, z_2^*) = 0, \\ 0 \leq \lambda^* \perp g(z_1^*, z_2^*) \geq 0, \\ h(z_1^*, z_2^*) = 0, \end{cases}$$

the point $(z_1^, z_2^*, \lambda^*, \mu^*)$ is a global (resp. local) optimal solution of problem (KKTR).*

(ii) *Let $(z_1^*, z_2^*, \lambda^*, \mu^*)$ be a global (resp. local) optimal solution of problem (KKTR) and assume the LLICQ holds at (z_1^*, z_2^*) (resp. LLICQ holds at (z_1, z_2) , for all $z_2 \in S(z_1)$, for all $z_1 : G(z_1) \geq 0, H(z_1) = 0$). Then, (z_1^*, z_2^*) is a global (resp. local) optimal solution of problem (BP).*

This theorem establishes global and local relationships between problem (BP) and its KKT/MPCC reformulation (KKTR). This will be useful in Section 4 when we construct a model for hyperparameter selection in support vector machines.

3 Solution methods

In this section, we introduce and study the partial penalisation and Scholtes relaxation methods. There are at least two motivations to study these algorithms. First, they are among the most commonly used in the literature on MPCCs (see, e.g., [48]). Secondly, there is a nice symmetry between the two approaches. In the penalisation method, the constraint $G(z)^\top H(z) = 0$ is moved from the feasible set to the objective function of the MPCC, while this same constraint is instead relaxed to form a larger feasible set in the context of the Scholtes relaxation approach. Another interesting feature in the dynamics of the two methods is that in the penalisation algorithm, the penalty parameter can be sequentially increased (this is the main focus of the analysis here), while the relaxation parameter instead decreases sequentially throughout the iterative process of the method.

As it will be clear in the next subsection, a common point between the two methods, which does not seem to have been discovered before, is that they converge under the same CQ; i.e., the MPCC-MFCQ-R, which is strictly weaker than any CQ that has been used so far to prove the convergence of these algorithms. However, a key departure point on the outcomes of the convergence of these two algorithms is that the penalisation method will be shown to converge to a S-stationary point, while we get C-stationarity from the Scholtes relaxation method.

3.1 Penalisation

We start by recalling that the penalisation method works by moving the problematic product term $G(z)^\top H(z)$ from the constraints of problem (MPCC) to its objective function. This leads to the penalised problem with π being the penalty parameter:

$$\begin{aligned} & \underset{z \in \mathbb{R}^n}{\text{minimise}} && f(z) + \pi G(z)^\top H(z) \\ & \text{subject to} && g_i(z) \geq 0 && \text{for } i = 1, \dots, p, \\ & && h_i(z) = 0 && \text{for } i = 1, \dots, q, \\ & && G_i(z) \geq 0, \quad H_i(z) \geq 0 && \text{for } i = 1, \dots, r. \end{aligned} \tag{P_\pi}$$

Interestingly, (P_π) is more likely to satisfy standard constraint qualifications. In fact, if (MPCC) satisfies MPCC-MFCQ-R at a S-stationary point z^* , then (P_π) will satisfy the MFCQ at z^* [81, Theorem 5.1]. This means that we can employ standard NLP solvers such as interior point method (IPM) or sequential quadratic programming (SQP) for (P_π) . However, when employing such solvers, we should not expect them to find a global minimum but instead terminate within a certain tolerance $\epsilon \geq 0$ of a KKT point. With this in mind, we define a point $z \in \mathbb{R}^n$ to be an




Fig. 2: This figure visualises applying the penalisation sequence to a hypothetical MPCC with a single complementary constraint. The vertical axis is the objective value, while the horizontal axis is the product of the two complementary constraint functions $G(x)H(x)$. The green line shows the objective value $f(x)$, and the red curve shows the penalised objective $f(x) + \pi G(x)H(x)$ for increasing penalty parameters π . The grey dots show infeasible stationary points that are slowly ironed out.

ϵ -KKT point for (\mathcal{P}_π) if there exists a Lagrange multiplier vector $(\lambda, \mu, \nu, \xi) \in \mathbb{R}^{(p+q+2r)}$ satisfying the following system:

$$\left\| \nabla f(z) + \pi H(z) [\nabla G(z)] + \pi G(z) [\nabla H(z)] - \sum_{i=1}^p \lambda_i \nabla g_i(z) - \sum_{i=1}^q \mu_i \nabla h_i(z) - \sum_{i=1}^r \nu_i \nabla G_i(z) - \sum_{i=1}^r \xi_i \nabla H_i(z) \right\| \leq \epsilon, \quad (11a)$$

$$\begin{aligned} g_i(z) &\geq 0, & \lambda_i &\geq 0, & g_i(z)\lambda_i &= 0, & \text{for } i = 1, \dots, p, \\ h_i(z) &= 0, & \mu_i &\text{free}, & & & \text{for } i = 1, \dots, q, \\ G_i(z) &\geq 0, & \nu_i &\geq 0, & G_i(z)\nu_i &= 0, & \text{for } i = 1, \dots, r, \\ H_i(z) &\geq 0, & \xi_i &\geq 0, & H_i(z)\xi_i &= 0, & \text{for } i = 1, \dots, r. \end{aligned} \quad (11b)$$

This concept is widely used in the literature; see, e.g., [25, Definition 3.1]. For $\epsilon = 0$ the system (11a-e) is the classical KKT-conditions for (\mathcal{P}_π) . The idea is then for a large enough penalty parameter π and small enough tolerance ϵ , finding a solution to (11a-e) that will correspond to a stationary point of problem (MPCC). If we know such values for π and ϵ , we only need to solve (\mathcal{P}_π) once. This is called *exact penalisation*. However, in practice, we may not have this information, or in the case where π is too large, (\mathcal{P}_π) may become ill-conditioned [57, Example 3]. Therefore, a more practical approach is to introduce an increasing sequence of non-negative penalty parameters $(\pi^t)_{t \in \mathbb{N}}$ and a decreasing sequence of tolerances $(\epsilon^t)_{t \in \mathbb{N}}$ and solve problem (\mathcal{P}_{π^t}) repeatedly using the previous solution as the starting point for the next iteration until our stopping criteria are met. We call this *sequential penalisation*. A precise description of the method is provided in Algorithm 1 and its theoretical convergence is established in next result below. A graphical illustration, for increasing penalty parameter values, is given in Figure 2.

Algorithm 1 Inexact sequential partial penalisation method for problem (MPCC)

Input: Access to the (MPCC)'s functions f, g, h, G, H and their first order derivatives; Starting point $z^0 \in \mathbb{R}^n$; Initial penalty parameter $\pi^1 > 0$; Initial KKT tolerance $\epsilon^1 > 0$; Exit tolerance $\epsilon_{\text{exit}} > 0$;

Procedure:

- 1: **repeat** for $t = 1, 2, 3, \dots$:
 - 2: Find an ϵ^t -KKT point z^t of (\mathcal{P}_{π^t}) with starting point z^{t-1} ;
 - 3: Increase the penalty parameter $\pi^{t+1} > \pi^t$;
 - 4: Decrease the KKT tolerance $\epsilon^{t+1} < \epsilon^t$;
 - 5: **until** Stopping condition: $G(z^t)^\top H(z^t) \leq \epsilon_{\text{exit}}$.
 - 6: Output z^t .
-

Theorem 3 Consider problem (MPCC) and a sequence of iterates $(z^t)_{t \in \mathbb{N}}$, generated by Algorithm 1, which is assumed to converge to the point z^* . Furthermore, suppose that MPCC-MFCQ-R holds at z^* and $G(z^*)^\top H(z^*) = 0$. Then z^* is a S-stationary point for problem (MPCC).

Proof For each iterate z^t , since it is an ϵ^t -KKT point of (\mathcal{P}_{π^t}) , let $(\lambda^t, \mu^t, \nu^t, \xi^t) \in \mathbb{R}^{(p+q+2r)}$ be a Lagrange multiplier vector satisfying (11a-e). We now define augmented multipliers:

$$\hat{\nu}_i^t := \nu_i^t - \pi^t H_i(z^t) \quad \text{and} \quad \hat{\xi}_i^t := \xi_i^t - \pi^t G_i(z^t). \quad (12)$$

In this way we can rewrite (11a) as

$$\left\| \nabla f(z^t) - \sum_{i=1}^p \lambda_i^t \nabla g_i(z^t) - \sum_{i=1}^q \mu_i^t \nabla h_i(z^t) - \sum_{i=1}^r \hat{\nu}_i^t \nabla G_i(z^t) - \sum_{i=1}^r \hat{\xi}_i^t \nabla H_i(z^t) \right\| \leq \epsilon^t. \quad (13)$$

We wish to prove that the sequence $a_t := (\lambda^t, \mu^t, \hat{\nu}^t, \hat{\xi}^t)_{t \in \mathbb{N}}$ has a convergent subsequence. Suppose towards a contradiction that it does not. It must be that a_t is unbounded and $\|a_t\| \rightarrow \infty$. Now, dividing equation (13) though by $\|a_t\|$ leads to

$$\left\| \frac{\nabla f(z^t)}{\|a_t\|} - \sum_{i=1}^p \frac{\lambda_i^t}{\|a_t\|} \nabla g_i(z^t) - \sum_{i=1}^q \frac{\mu_i^t}{\|a_t\|} \nabla h_i(z^t) - \sum_{i=1}^r \frac{\hat{\nu}_i^t}{\|a_t\|} \nabla G_i(z^t) - \sum_{i=1}^r \frac{\hat{\xi}_i^t}{\|a_t\|} \nabla H_i(z^t) \right\| \leq \frac{\epsilon^t}{\|a_t\|}. \quad (14)$$

We consider the normed sequence of the Lagrange multipliers $\frac{(\lambda^t, \mu^t, \hat{\nu}^t, \hat{\xi}^t)}{\|(\lambda^t, \mu^t, \hat{\nu}^t, \hat{\xi}^t)\|}$. Clearly, as it is bounded, by Bolzano–Weierstrass theorem, there exists some infinite subsequence that converges. Let $(\lambda', \mu', \hat{\nu}', \hat{\xi}')$ be the limit of that subsequence. We know $(\lambda', \mu', \hat{\nu}', \hat{\xi}')$ must be non-zero; otherwise, the original sequence a_t would also have a subsequence converging to zero. From the complementary slackness conditions in (11b) we conclude that: $\lambda'_i = 0$ for $i \in \{1, \dots, p\} \setminus I^g$ and $\hat{\nu}'_i = 0$ for $i \in I^H$ and $\hat{\xi}'_i = 0$ for $i \in I^G$ respectively. Using this indexing, we can say that it must be the $(\lambda'_{I^g}, \mu', \hat{\nu}'_{I^G \cup I^{GH}}, \hat{\xi}'_{I^H \cup I^{GH}})$ components of the sequence that are non-zero. Now, returning to (14), as $t \rightarrow \infty$, the terms $\frac{\nabla f(z^t)}{\|a_t\|}$ and $\frac{\epsilon^t}{\|a_t\|}$ go to zero,

$$\sum_{i \in I^g} \lambda'_i \nabla g_i(z^*) + \sum_{i \in I^H} \mu'_i \nabla h_i(z^*) + \sum_{i \in I^G \cup I^{GH}} \hat{\nu}'_i \nabla G_i(z^*) + \sum_{i \in I^H \cup I^{GH}} \hat{\xi}'_i \nabla H_i(z^*) = 0. \quad (15)$$

From the assumptions of the theorem, we have that z^* satisfies MPCC-MFCQ-R. By Lemma 1, the set of gradient vectors in (7) should be positively linearly independent. But that is contradicted by (15). Therefore, we conclude that a_t has a convergent subsequence. Let $(\lambda^*, \mu^*, \hat{\nu}^*, \hat{\xi}^*)$ be the limit of that convergent subsequence. By the continuous differentiability of f, g, h, G, H , and since the sequence $(z^t)_{t \in \mathbb{N}}$ converges in their domain, these functions and their derivatives must converge in range. Applying this to inequality (13), we get

$$\nabla f(z^*) - \sum_{i=1}^p \lambda_i^* \nabla g_i(z^*) - \sum_{i=1}^q \mu_i^* \nabla h_i(z^*) - \sum_{i=1}^r \hat{\nu}_i^* \nabla G_i(z^*) - \sum_{i=1}^r \hat{\xi}_i^* \nabla H_i(z^*) = 0. \quad (16)$$

Considering an index $i \in I^G$, by definition, we have $G_i(z^*) = 0$ and $H_i(z^*) > 0$. By the complementary slackness, we have $\xi_i^* = 0$. Thus, we can conclude from (12) that $\hat{\xi}_i^* = 0$ for all $i \in I^G$. In the same way, we can conclude that $\hat{\nu}_i^* = 0$ for all $i \in I^H$. Now considering a bi-active index $i \in I^{GH}$, $G_i(z^*) = 0$ and $H_i(z^*) = 0$. Thus $\hat{\nu}_i^* = \nu_i^* \geq 0$ and $\hat{\xi}_i^* = \xi_i^* \geq 0$. Hence,

$$\hat{\nu}_i^* = 0 \quad \forall i \in I^H, \quad \hat{\xi}_i^* = 0 \quad \forall i \in I^G, \quad \hat{\nu}_i^*, \hat{\xi}_i^* \geq 0 \quad \forall i \in I^{GH}. \quad (17)$$

Consider the feasible point z^* and vector of multipliers $(\lambda^*, \mu^*, \hat{\nu}^*, \hat{\xi}^*)$. From equations (16), (11b) and (17), we can conclude that conditions (9a), (9b-e) and (9S) are respectively satisfied. These are exactly the S-stationary conditions. \square

Recall that Ralph and Wright [81, Theorem 5.2] provide a nice proof for penalisation with a fixed parameter π . However, in practice, a user would not know a good value for the π ahead of time. If the chosen π is too small, the property $G(z^*)^\top H(z^*) = 0$ may not hold. On the other hand, if it is too large, the program may become ill-conditioned [57, Example 3] as is common with penalisation-based algorithms [70, 103]. This provides a motivation for the sequential approach. Also, observe that the paper by Leyffer et al. [57, Theorem 3.4] already presents a sequential penalisation theory in the case where the subproblem (\mathcal{P}_π) is specifically solved with the interior point algorithm. Theorem 3 provides two useful generalisations over their work. Firstly, it gives flexibility to the user in the selection of the algorithm for the penalised subproblem (\mathcal{P}_π) . Secondly, we only assume MPCC-MFCQ-R, which is a far weaker CQ compared to the MPCC-LICQ, required in [57]. It is also important to note all the aforementioned convergence results for penalisation, including ours, require the fulfilment of $G(z^*)^\top H(z^*) = 0$. However, it is unclear how to ensure that this is the case in practice. In the remainder of this section, we provide Theorem 4 and Example 2 to establish when this assumption may and may not be reasonable.

Theorem 4 *Let $(z^t)_{t \in \mathbb{N}}$ be a sequence of iterates generated by Algorithm 1 such that z^t is a global optimal solution of problem (\mathcal{P}_π) for $t \in \mathbb{N}$. If $\lim_{t \rightarrow \infty} z^t = z^*$, then $G(z^*)^\top H(z^*) = 0$.*

Proof We first want to show that $(G(z^t)^\top H(z^t))_{t \in \mathbb{N}}$ is a non-increasing sequence. We will apply the classical trick of taking a difference of terms at two indices. Let $\bar{t} > t$. Notice both $z^{\bar{t}}$ and z^t are feasible for both $(\mathcal{P}_{\pi^{\bar{t}}})$ and (\mathcal{P}_{π^t}) . Since z^t minimises $f(z) + \pi^t G(z)H(z)$ and $z^{\bar{t}}$ minimises $f(z) + \pi^{\bar{t}} G(z)H(z)$ we can conclude that

$$f(z^{\bar{t}}) + \pi^t G(z^{\bar{t}})H(z^{\bar{t}}) \geq f(z^t) + \pi^t G(z^t)H(z^t), \quad (18a)$$

$$f(z^t) + \pi^{\bar{t}} G(z^t)H(z^t) \geq f(z^{\bar{t}}) + \pi^{\bar{t}} G(z^{\bar{t}})H(z^{\bar{t}}). \quad (18b)$$

Adding the two inequalities (18a) and (18b) gives

$$(\pi^{\bar{t}} - \pi^t) (G(z^t)H(z^t) - G(z^{\bar{t}})H(z^{\bar{t}})) \geq 0.$$

Since $(\pi^{\bar{t}} - \pi^t) > 0$, it follows that $G(z^t)H(z^t) \geq G(z^{\bar{t}})H(z^{\bar{t}})$.

Next, we show that $\lim_{t \rightarrow \infty} (G(z^t)^\top H(z^t)) = 0$. Let \bar{z} be any feasible point of (MPCC). Let z^1 be the minimiser of the penalised problem for $\pi^1 = 1$. Towards a contradiction, suppose that there exists an $\epsilon > 0$ such that for all $T \in \mathbb{N}$ we have $G(z^t)^\top H(z^t) > \epsilon$ for some $t > T$. Under that premise, we can choose an arbitrarily large t such that $\pi^t \geq \frac{1}{\epsilon} |f(\bar{z}) - f(z^1)| + 2$ and $G(z^t)H(z^t) > \epsilon$. Consider the infimum over the objective values of problem (MPCC):

$$\inf \{f(z) : g(z) \geq 0, h(z) = 0, 0 \leq G(z)^\top H(z) \leq 0\} \quad (19a)$$

$$\geq \inf \{f(z) + \pi^t G(z)H(z) : g(z) \geq 0, h(z) = 0, G(z) \geq 0, H(z) \geq 0\} \quad (19b)$$

$$= f(z^t) + \pi^t G(z^t)H(z^t) \quad (19c)$$

$$\begin{aligned} &\geq f(z^1) + \pi^t G(z^t)H(z^t) \\ &\geq f(z^1) + \pi^t \epsilon \\ &\geq f(z^1) + |f(\bar{z}) - f(z^1)| + 2\epsilon \\ &> f(\bar{z}). \end{aligned} \quad (19d)$$

It is helpful to remember that t here is chosen according to ϵ . The first inequality (19b) comes from the fact that any global minimum of (MPCC) is also feasible for (\mathcal{P}_π) . The second inequality (19c) comes from the fact that z^t is a global minimum of $f(z) + \pi^t G(z)H(z)$.

The conclusion of (19d) implies that \bar{z} achieves a smaller objective value than the infimum over objectives of the MPCC feasible points in problem (19a). But this contradicts the MPCC feasibility of \bar{z} . Therefore we conclude that $\lim_{t \rightarrow \infty} (G(z^t)^\top H(z^t)) = 0$. By the continuity of the functions G and H , we have $G(z^*)^\top H(z^*) = 0$. \square




Fig. 3: A graphical representation of Example 2. On the left are the contours of the objective function $f(z)$. When the complementary constraints $0 \leq G(z)H(z) \leq 0$ are added, the feasible set becomes the positive vertical and horizontal axes. In the centre are the contours of the penalised objective function $f(z) + \pi G(z)H(z)$. The two dark green wells at $(0,3)$ and $(3,0)$ are the two global minima. The light green well at $(3,3)$ is a local minimum. On the right is the constraint $G(z)$ plotted against z_1 .

Next, we provide an example, which shows how important it is for the sequence of points generated by Algorithm 1 to be global optimal solutions for the penalised problem (\mathcal{P}_π) . In this example, the objective and constraint functions $f, G, H : \mathbb{R}^2 \rightarrow \mathbb{R}$ are all continuously differentiable, and the MPCC-LICQ is satisfied. Nevertheless, its corresponding penalisation formulation has a sequence of local minimums that converge to an infeasible point for the corresponding (MPCC).

Example 2 Consider the following example of MPCC:

$$\begin{aligned}
 & \underset{z_1, z_2}{\text{minimise}} && f(z) := (z_1 - 3)^2 + (z_2 - 3)^2 \\
 & \text{subject to} && G(z) := \left(\frac{1}{3}z_1^3 - \frac{9}{4}z_1^2 + \frac{9}{2}z_1 \right) \geq 0, \\
 & && H(z) := \left(\frac{1}{3}z_2^3 - \frac{9}{4}z_2^2 + \frac{9}{2}z_2 \right) \geq 0, \\
 & && G(z)H(z) = 0.
 \end{aligned} \tag{20}$$

It is helpful to note that G or H are only equal to zero for $z_1 = 0$ or $z_2 = 0$, respectively. The feasible set is, therefore, the union of the two non-negative axes (see Figure 3). We begin by writing the derivatives of all the functions involved in the above problem:

$$\nabla f(z) = \begin{bmatrix} 2z_1 - 6 \\ 2z_2 - 6 \end{bmatrix}, \quad \nabla G(z) = \begin{bmatrix} (z_1 - \frac{3}{2})(z_1 - 3) \\ 0 \end{bmatrix}, \quad \nabla H(z) = \begin{bmatrix} 0 \\ (z_2 - \frac{3}{2})(z_2 - 3) \end{bmatrix}.$$

Program (20) has two local minima at $[3, 0]^\top$ and $[0, 3]^\top$. Furthermore, it satisfies MPCC-LICQ at all feasible points. This can be verified by checking the two strictly complementary cases of $(G(z) = 0, H(z) > 0)$ and $(G(z) > 0, H(z) = 0)$ and the biactive case $(G(z) = 0, H(z) = 0)$. For each of these, in order, the set of gradients of the active constraints are $\left\{ [0, \frac{9}{2}]^\top \right\}$ and $\left\{ [\frac{9}{2}, 0]^\top \right\}$ and $\left\{ [0, \frac{9}{2}]^\top, [\frac{9}{2}, 0]^\top \right\}$ which are each linearly independent sets.

Now consider the penalisation formulation (\mathcal{P}_π) of this example whose feasible set is the whole non-negative quadrant. Note (20) has only complementary constraint functions G, H and no inequality g nor equality h constraint functions. We introduce multipliers ν and ξ corresponding to G and H respectively and write the KKT conditions of the corresponding problem (\mathcal{P}_π) :

$$\begin{aligned}
 0 &= \nabla [f(z) + \pi^t G(z)H(z)] - \nu \nabla G(z) - \xi \nabla H(z), \\
 G(z) &\geq 0, \quad H(z) \geq 0, \quad \nu \geq 0, \quad \xi \geq 0, \quad G(z)\nu = 0, \quad H(z)\xi = 0.
 \end{aligned} \tag{21}$$

Consider the point $(z_1^*, z_2^*) = (3, 3)$, feasible for the corresponding version of problem (\mathcal{P}_π) for our example but not for problem (20) itself. At this point $G(z^*) = H(z^*) = 2.25$ and all the derivatives

are zero. $\nabla f(z^*) = \nabla G(z^*) = \nabla H(z^*) = 0$. Substituting these values into the KKT conditions for the penalisation problem (21) we can find that $(z_1^*, z_2^*, \nu, \xi) = (3, 3, 0, 0)$ is a KKT point of the penalisation problem (\mathcal{P}_{π^t}) for all parameters $\pi^t \geq 0$. Algorithm 1 could reasonably find the sequence of iterates $z^t = (3, 3)$ for each $t \in \mathbb{N}$. However, this converges to a point that is infeasible for the original MPCC. This highlights a key weakness of the penalisation method.

3.2 Relaxation

Scholtes [83] proposed the first regularisation scheme where problem (MPCC) is approximated by a sequence of parametrised nonlinear programs. Since then, many alternative relaxation schemes have been proposed (e.g., [16, 91]), and a good survey on the subject is given by Hoheisel et al. [44]. Typically, in relaxation schemes for MPCCs, the constraint involving the problematic product term $G(z)^\top H(z)$ of the complementarity system is relaxed. This enlarges the feasible set and can make the resulting problem more tractable. In the specific context of the Scholtes relaxation, the following subproblem is considered:

$$\begin{aligned} & \underset{z}{\text{minimise}} && f(z) \\ & \text{subject to} && g_i(z) \geq 0 && \text{for } i = 1, \dots, p, \\ & && h_i(z) = 0 && \text{for } i = 1, \dots, q, \\ & && G_i(z) \geq 0, H_i(z) \geq 0, G_i(z)H_i(z) \leq \tau && \text{for } i = 1, \dots, r. \end{aligned} \quad (\mathcal{R}_\tau)$$

Note that for τ close to zero, the feasible set of subproblem (\mathcal{R}_τ) closely approximates that of (MPCC); see Figure 4. This forms the basis of the relaxation method. The program (\mathcal{R}_{τ^t}) is solved multiple times with a sequence of decreasing parameters τ^t . For this, any classical NLP solver may be used.

In a similar fashion to our construction for penalisation, we shall not expect the NLP solvers to find the exact global minimum but instead terminate within a certain tolerance $\epsilon \geq 0$ of a KKT point. We say a point $z \in \mathbb{R}^n$ is an ϵ -KKT point for (\mathcal{R}_τ) if there exists a Lagrange multiplier vector $(\lambda, \mu, \nu, \xi, \omega) \in \mathbb{R}^{(p+q+3r)}$ such that

$$\begin{aligned} & \left\| \nabla f(z) - \sum_{i=1}^p \lambda_i \nabla g_i(z) - \sum_{i=1}^q \mu_i \nabla h_i(z) - \sum_{i=1}^r \nu_i \nabla G_i(z) - \sum_{i=1}^r \xi_i \nabla H_i(z) \right. \\ & \quad \left. + \sum_{i=1}^r \omega_i [G_i(z) \nabla H_i(z) + H_i(z) \nabla G_i(z)] \right\| \leq \epsilon, \end{aligned} \quad (22a)$$

$$\begin{aligned} & g_i(z) \geq 0, \quad \lambda_i \geq 0, & g_i(z) \lambda_i = 0, & \text{for } i = 1, \dots, p, \\ & h_i(z) = 0, \quad \mu_i \text{ free}, & & \text{for } i = 1, \dots, q, \\ & G_i(z) \geq 0, \quad \nu_i \geq 0, & G_i(z) \nu_i = 0, & \text{for } i = 1, \dots, r, \\ & H_i(z) \geq 0, \quad \xi_i \geq 0, & H_i(z) \xi_i = 0, & \text{for } i = 1, \dots, r, \\ & [\tau - G_i(z)H_i(z)] \geq 0, \quad \omega_i \geq 0, & [\tau - G_i(z)H_i(z)] \omega_i = 0, & \text{for } i = 1, \dots, r. \end{aligned} \quad (22b)$$

We now show that a sequence of such ϵ -KKT points generated by Algorithm 2, under certain assumptions, converges to a C-stationary point of the original (MPCC) problem.

Algorithm 2 Inexact sequential Scholtes relaxation method for problem (MPCC)

Input: Access to the (MPCC)'s functions f, g, h, G, H and their first order derivatives; Starting point $z^0 \in \mathbb{R}^n$; Initial relaxation parameter $\tau^1 > 0$; Initial KKT tolerance $\epsilon^1 > 0$; Exit tolerance $\epsilon_{\text{exit}} > 0$;

Procedure:

- 1: **repeat** for $t = 1, 2, 3, \dots$:
 - 2: Find an ϵ^t -KKT point z^t of (\mathcal{R}_{τ^t}) with starting point z^{t-1} ;
 - 3: Decrease the relaxation parameter $\tau^{t+1} < \tau^t$;
 - 4: Decrease the KKT tolerance $\epsilon^{t+1} < \epsilon^t$;
 - 5: **until** Stopping condition: $G(z^t)^\top H(z^t) \leq \epsilon_{\text{exit}}$.
 - 6: Output z^t .
-




Fig. 4: The feasible set for the relaxation subproblems with (left to right) decreasing relaxation parameters.

Theorem 5 Consider problem (MPCC). Let $(z^t)_{t \in \mathbb{N}}$ be a sequence of ϵ^t -KKT points of (\mathcal{R}_{τ^t}) parametrised by a strictly decreasing sequence of positive relaxation parameters $(\tau^t)_{t \in \mathbb{N}}$ and tolerances $(\epsilon^t)_{t \in \mathbb{N}}$ that both tend to zero. Let $z^* := \lim_{t \rightarrow \infty} (z^t)$ exist. Let MPCC-MFCQ-R hold at z^* . Then, z^* is a C-stationary point for problem (MPCC).

Proof For each iterate z^t , since it is an ϵ -KKT point of (\mathcal{R}_{τ^t}) , there exists a vector of Lagrange multipliers $(\lambda^t, \mu^t, \nu^t, \xi^t, \omega^t) \in \mathbb{R}^{(p+q+3r)}$ satisfying (22a-e). Consider the following constructions:

$$\hat{\nu}_i^t := \begin{cases} \nu_i^t & \text{if } i \in I^G(z^t) \cup I^{GH}(z^t), \\ -\omega_i^t H_i(z^t) & \text{if } i \in \{i : G_i(z^t)H_i(z^t) = \tau^t\}, \\ 0 & \text{otherwise,} \end{cases}$$

$$\hat{\xi}_i^t := \begin{cases} \xi_i^t & \text{if } i \in I^H(z^t) \cup I^{GH}(z^t), \\ -\omega_i^t G_i(z^t) & \text{if } i \in \{i : G_i(z^t)H_i(z^t) = \tau^t\}, \\ 0 & \text{otherwise.} \end{cases}$$

This allows us to rewrite equation (22a) as

$$\left\| \nabla f(z^t) - \sum_{i=1}^p \lambda_i^t \nabla g_i(z^t) - \sum_{i=1}^q \mu_i^t \nabla h_i(z^t) - \sum_{i=1}^r \hat{\nu}_i^t \nabla G_i(z^t) - \sum_{i=1}^r \hat{\xi}_i^t \nabla H_i(z^t) \right\| \leq \epsilon^t. \quad (23)$$

We also note that for each $i \in \{i : G_i(z^t)H_i(z^t) = \tau^t\}$, both $\hat{\nu}_i^t$ and $\hat{\xi}_i^t$ are non-positive, while for each $i \in I^G(z^t) \cup I^{GH}(z^t)$, both $\hat{\nu}_i^t$ and $\hat{\xi}_i^t$ are non-negative. For all other indices, they are zero. Therefore, we can conclude that

$$\hat{\nu}_i^t \hat{\xi}_i^t \geq 0 \quad \forall i \in \{1, \dots, r\}. \quad (24)$$

We now wish to prove that the sequence $a_t := (\lambda^t, \mu^t, \hat{\nu}^t, \hat{\xi}^t)_{t \in \mathbb{N}}$ has a convergent subsequence. Suppose towards a contradiction that it does not. It must be that a_t is unbounded and $\|a_t\| \rightarrow \infty$. Diving equation (23) through by $\|a_t\|$ leads to

$$\left\| \frac{\nabla f(z^t)}{\|a_t\|} - \sum_{i=1}^p \frac{\lambda_i^t}{\|a_t\|} \nabla g_i(z^t) - \sum_{i=1}^q \frac{\mu_i^t}{\|a_t\|} \nabla h_i(z^t) - \sum_{i=1}^r \frac{\hat{\nu}_i^t}{\|a_t\|} \nabla G_i(z^t) - \sum_{i=1}^r \frac{\hat{\xi}_i^t}{\|a_t\|} \nabla H_i(z^t) \right\| \leq \frac{\epsilon^t}{\|a_t\|}. \quad (25)$$

Now consider the normed sequence $\frac{(\lambda^t, \mu^t, \hat{\nu}^t, \hat{\xi}^t)}{\|(\lambda^t, \mu^t, \hat{\nu}^t, \hat{\xi}^t)\|}$. Clearly, it is bounded. It follows from the Bolzano–Weierstrass theorem that there exists some subsequence that converges. Let the point $(\lambda', \mu', \hat{\nu}', \hat{\xi}')$ be the limit of that subsequence. We know this must be non-zero; otherwise, the original sequence a_t would also have a subsequence converging to zero. From the complementary slackness conditions in (22b) we conclude that: λ'_i is non-zero only for $i \in I^g$ and $\hat{\nu}'_i$ is non-zero only for $i \in I^G \cup I^{GH}$ and $\hat{\xi}'_i$ is non-zero only for $i \in I^H \cup I^{GH}$. Using this indexing we can say that it must be the $(\lambda'_{I^g}, \mu', \nu'_{I^G \cup I^{GH}}, \xi'_{I^H \cup I^{GH}})$ components of the sequence that are non-zero.

Returning to (25), as $t \rightarrow \infty$, the terms $\frac{\nabla f(z^t)}{\|a_t\|}$ and $\frac{\epsilon^t}{\|a_t\|}$ go to zero, and we are left with

$$\sum_{i \in I^g} \lambda'_i \nabla g_i(z^*) + \sum_{i \in I^h} \mu'_i \nabla h_i(z^*) + \sum_{i \in I^G \cup I^{GH}} \hat{\nu}'_i \nabla G_i(z^*) + \sum_{i \in I^H \cup I^{GH}} \hat{\xi}'_i \nabla H_i(z^*) = 0. \quad (26)$$

From the assumptions of the theorem, we have that z^* satisfies MPCC-MFCQ-R. By Lemma 1, the set of gradient vectors in (7) should be positively linearly independent. But that is contradicted by (26). Therefore, we conclude a_t must have a convergent subsequence. Let $(\lambda^*, \mu^*, \hat{\nu}^*, \hat{\xi}^*)$ be the limit of that convergent subsequence. By the continuous differentiability of f, g, h, G, H , and since the sequence $(z^t)_{t \in \mathbb{N}}$ converges in their domain, these functions and their derivatives must converge in range. We can apply all this to (22a) to get the following:

$$\nabla f(z^t) - \sum_{i=1}^p \lambda_i^* \nabla g_i(z^*) - \sum_{i=1}^q \mu_i^* \nabla h_i(z^*) - \sum_{i=1}^r \hat{\nu}_i^* \nabla G_i(z^*) - \sum_{i=1}^r \hat{\xi}_i^* \nabla H_i(z^*) = 0. \quad (27)$$

From equations (27), (22b) and (24), we can conclude that (9a), (9b-e) and (9C) hold, respectively. These are exactly the C-stationary conditions. \square

A similar version of Theorem 5 was initially proven by Scholtes [83, page 922, Section 3] but required the stronger MPCC-LICQ. An updated proof is given by Schwartz [84, page 112, Theorem 7.3], which weakened the requirement to MPCC-MFCQ-T. We present a similar proof but weakening the CQ further to MPCC-MFCQ-R (cf. Theorem 1 and Example 1). Another unique feature of our result (Theorem 5) is that we only require an inexact KKT point to be computed at each iteration, unlike in the previous papers, where exact KKT points were used for the penalised problem. This allows the theorem to be applied to a much broader class of problems and for the numerical implementation to reflect the theory more closely.

4 Application to hyperparameter tuning for nonlinear support vector machines

In this section, we construct a model to optimally select the hyperparameters of a support vector machine (SVM) model. To proceed, we first introduce the problem of training a SVM (Subsection 4.1). Subsequently, we discuss the cross-fold validation method and then build on it to formulate a bilevel optimisation problem where the lower-level program trains the SVM model parameters while the corresponding upper-level program chooses the optimal hyperparameters. Based on the KKT reformulation introduced in Subsection 2.3, this bilevel program is transformed into a MPCC (Subsection 4.2), which is then carefully analysed in light of the results established in the previous two sections. In particular, we show that the MPCC reformulation of the bilevel hyperparameter tuning model for our nonlinear SVM automatically satisfies MPCC-MFCQ-R for any feasible point with positive hyperparameters. Additionally, we prove that the MPCC-LICQ and MPCC-MFCQ-T fail in general (Subsection 4.3).

4.1 Support vector machine

The concept of support vector machine (SVM) was introduced by Cortes and Vapnik [15, 96] and has since been a prolific research subject; see, e.g., [10]. Here, we provide an introduction to the topic as well as a setup of notation that will be useful for the discussion in the sequel.

The SVM is defined by a set of n training examples, denoted by (x_i, y_i) for $i = 1, \dots, n$, such that each example i is a vector of features collected by $x_i \in \mathbb{R}^d$ and a class label $y_i \in \{-1, +1\}$, where the aim is to find weights $w \in \mathbb{R}^d$ and a bias $b \in \mathbb{R}$ describing a hyperplane $\{x \in \mathbb{R}^d : w^\top x + b = 0\}$ and a margin around it, $\{x \in \mathbb{R}^d : -1 \leq w^\top x + b \leq 1\}$. The weights should be such that the positive label points lie above the margin ($w^\top x_i + b \geq 1 \quad \forall i : y_i = +1$) and the negative label points lie below the margin ($w^\top x_i + b \leq -1 \quad \forall i : y_i = -1$). This is equivalent to the constraint $1 - y_i(w^\top x_i + b) \leq 0$ for any i . Under these constraints, our goal is to maximise the width of the margin given by $\frac{2}{\|w\|}$. Intuitively, this means new unseen data points have the most leeway to

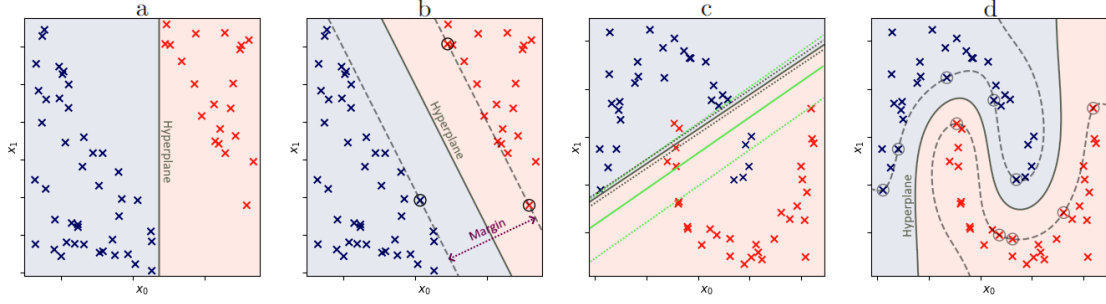


Fig. 5: Each of the four plots show blue and red data points in a two-dimensional space. Plots a and b (left) are the same linearly separable data set. Plots c and d (right) are sampled from the Moons data set. The first plot shows a naive separating hyperplane. Plot b shows a more intelligent hyperplane with a maximum margin found using the (Hard-Margin-SVM). Plot c shows two possible hyperplanes in black and green found by the (Soft-Margin-SVM) with different choices of hyperparameter C . Plot d shows a nonlinear boundary with a maximum margin. We found this using the (Dual-SVM) attached with a RBF kernel.

deviate beyond the range of the training data without crossing to the wrong side of the hyperplane; see Figure 5a-b. Once optimal weights w^* and bias b^* are found, we can predict the label of new data using the classifier $x \mapsto \text{sign}(w^{*\top}x + b^*)$ where the sign function returns -1 for negative inputs and +1 otherwise. This problem can be formulated as

$$\begin{aligned} & \underset{w \in \mathbb{R}^d, b \in \mathbb{R}}{\text{minimise}} && \frac{1}{2} w^\top w \\ & \text{subject to} && 1 - y_i (w^\top x_i + b) \leq 0 \quad \text{for } i = 1, \dots, n. \end{aligned} \quad (\text{Hard-Margin-SVM})$$

The weakness of the (Hard-Margin-SVM) becomes apparent when data is not linearly separable; the above problem is then infeasible. Thus, the (Hard-Margin-SVM) is rarely used in practice. A far more common approach is to use the hinge loss function

$$\text{HingeLoss}(x_i, y_i) := \max(0, 1 - y_i (w^\top x_i + b)),$$

which measures how much example (x_i, y_i) violates the margin constraint for a given weight and bias. It takes the value zero when the constraint is satisfied. By replacing each constraint with a hinge loss penalty in the objective function, we can form a model that discourages but allows some points to violate the constraint in favour of ensuring a large margin. A regularisation parameter C can be introduced to instruct the balance between these two objectives, with a large C ensuring that the margin constraint is less violated, whereas a smaller C is more tolerant to such violations where it can find a larger margin. The new objective is now to minimise the function

$$\frac{1}{2} w^\top w + C \frac{1}{n} \sum_{i=1}^n \max(0, 1 - y_i (w^\top x_i + b)).$$

The non-smooth max operator in this new objective function can be eliminated by introducing a variable ξ_i for each example (x_i, y_i) . This leads to the soft margin formulation of the SVM problem

$$\begin{aligned} & \underset{w \in \mathbb{R}^d, \xi \in \mathbb{R}^n, b \in \mathbb{R}}{\text{minimise}} && \frac{1}{2} w^\top w + C \frac{1}{n} \sum_{i=1}^n \xi_i \\ & \text{subject to} && \xi_i \geq 1 - y_i (w^\top x_i + b) \quad \text{for } i = 1, \dots, n, \\ & && \xi_i \geq 0 \quad \text{for } i = 1, \dots, n. \end{aligned} \quad (\text{Soft-Margin-SVM})$$

Now, we wish to extend this model so that it can find nonlinear boundaries to separate data. To achieve this, we consider a mapping $x_i \mapsto \phi(x_i)$ that lifts the features of each data point x_i , for $i = 1, \dots, n$, to a higher dimension. We can then find a hyperplane in this higher dimension, which corresponds to a nonlinear boundary in the original feature space; see Figure 6 for an illustration. However, it is often computationally hard to compute $\phi(x_i)$. The solution is the so-called kernel




Fig. 6: On the left, data in two dimensions that is not linearly separable is shown. On the right, the same data is lifted into a third dimension, where a separating hyperplane exists.

trick, which allows us to find a separating hyperplane in the higher dimension without ever explicitly evaluating $\phi(x_i)$. In our numerical experiments, we use the Radial Basis Function (RBF), which is defined by

$$K_\gamma(x_1, x_2) := \exp\left(-\gamma \|x_1 - x_2\|^2\right) \quad (28)$$

(with hyperparameter γ) and maps features to an infinite-dimensional space, making the kernel trick necessary. Other kernel functions commonly used in the literature [7, 10] include the linear, polynomial, and Sigmoid ones, respectively given as follows (for hyperparameters γ, δ):

$$K(x_1, x_2) := x_1^\top x_2, \quad K_{\gamma, \delta}(x_1, x_2) := (\gamma x_1^\top x_2 + \delta)^2, \quad \text{and} \quad K_{\gamma, \delta}(x_1, x_2) := \tanh(\gamma x_1^\top x_2 + \delta).$$

To apply the kernel trick, we need to work with the dual of the soft margin SVM optimisation problem. This was first discovered by Guyon, Boser and Vapnik in 1992 [7, 40] and is now standard. Nevertheless, for the paper to be self-contained, in Appendix B, we provide the details of the conversion from the primal problem (Soft-Margin-SVM) to the dual (Dual-SVM)

$$\begin{aligned} & \underset{\alpha \in \mathbb{R}^n}{\text{minimise}} && \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j y_i y_j K_\gamma(x_i, x_j) - \sum_{i=1}^n \alpha_i \\ & \text{subject to} && 0 \leq \alpha_i \leq C \quad \text{for } i = 1, \dots, n, \\ & && \sum_{i=1}^n \alpha_i y_i = 0, \end{aligned} \quad (\text{Dual-SVM})$$

which will be the main focus of our attention for the remainder of this section.

The (Dual-SVM) is highly effective at learning to separate real-world data. Many machine-learning libraries have become proficient at solving this; one of the most widely used is the open-source library LIBSVM [11]. At this point, it is worth discussing the link between the primal variable (w, b) and dual variable α . For each index i , if $\alpha_i = 0$, the solution represents a boundary that ensures data point i lies on the correct side of the margin. If $\alpha_i > 0$, then the data point i lies within the margin or on the wrong side. Details of how to reconstruct the primal variables and further explanations are provided in Appendix B.2.

We conclude this subsection by stating the Karush–Kuhn–Tucker (KKT) conditions for problem (Dual-SVM), which will be useful later. For the theory on KKT conditions, we refer again to Boyd and Vandenberghe [8, Section 5.5.3]. Since the program is convex and the constraints are linear, a point α is optimal if and only if there exist Lagrange multipliers $\underline{v}, \bar{v} \in \mathbb{R}^n, u \in \mathbb{R}$ such that

$$\begin{aligned} & \sum_{j=1}^n \alpha_j Q(\gamma)_{ij} - \underline{v}_i + \bar{v}_i + u y_i = 1 && \text{for } i = 1, \dots, n, \\ & 0 \leq \alpha_i \perp \underline{v}_i \geq 0 && \text{for } i = 1, \dots, n, \\ & 0 \leq (C - \alpha_i) \perp \bar{v}_i \geq 0 && \text{for } i = 1, \dots, n, \\ & \sum_{i=1}^n \alpha_i y_i = 0, \end{aligned} \quad (29)$$




Fig. 7: A visualization of 5-fold cross-validation with dedicated test data kept aside.

where $Q(\gamma)_{ij} := y_i y_j K_\gamma(x_i, x_j)$, $i = 1, \dots, n$, $j = 1, \dots, n$, denotes the so-called *kernel* matrix which we write as a function of γ as this will become a decision variable in the next section. For full details on this, see Appendix B.

4.2 Modelling hyperparameter tuning

A major aspect of machine learning is judging if a model is overfitted to its training data [37, page 28][43, Chapter 7]. For classifiers, this could mean unreasonably contorting the decision boundary to include outliers. By doing so, the classifier achieves near-zero empirical training error but misses the overarching pattern. SVMs' power comes from their ability to accept some level of empirical error in favour of a simpler boundary that generalises well to unseen data. These two objectives are balanced by the hyperparameter C . This leads to the question of how to evaluate the accuracy of a model or the choice of hyperparameters if training error does not provide the whole picture. Traditionally, data scientists have suggested withholding some (validation) data from the training process and using it separately to test the classifier's accuracy more fairly. However, this reduces the amount of data that the model has to learn from and can be sensitive to the choice of data withheld for validation. Our goal is to find the optimal hyperparameters without overfitting the model so we must use a more robust validation methodology.

The cross-validation procedure (see Figure 7 and literature [43, Section 7.10.1], [37, page 77], [92]) partitions the data indices into K folds. It then runs K experiment(s) such that at the k -th experiment ($k = 1, \dots, K$), new model weights w^k and bias b^k are learned using all the examples but withholding those from fold k . Then, the remaining data points from fold k can be used as unseen validation data to evaluate the model by calculating the average accuracy across the K experiments. This typically results in a less biased evaluation of the model.

Moving forward, K will represent the total number of folds, while k denotes the index of a given fold. For each experiment $k = 1, \dots, K$, the superscript k corresponds to variables related specifically to that experiment; i.e., precisely, we use the following notation for experiment k :

	Training	Validation
Number of examples:	n^k	\bar{n}^k
Examples:	$(x_i^k, y_i^k)_{i=1, \dots, n^k}$	$(\bar{x}_i^k, \bar{y}_i^k)_{i=1, \dots, \bar{n}^k}$
Kernel matrix:	$Q(\gamma)_{ij}^k := y_i^k y_j^k \exp(-\gamma \ x_i^k - x_j^k\)$	$\bar{Q}(\gamma)_{ij}^k := \bar{y}_i^k y_j^k \exp(-\gamma \ \bar{x}_i^k - x_j^k\)$

Note that in training, the kernel $Q(\gamma)_{ij}^k$ is calculated between two training examples (x_i^k, y_i^k) and (x_j^k, y_j^k) but, in validation, the kernel $\bar{Q}(\gamma)_{ij}^k$ is between one validation $(\bar{x}_i^k, \bar{y}_i^k)$ and one training example (x_j^k, y_j^k) . In a similar fashion, we define w^k , b^k , and α^k as the decision variables related to experiment k . This allows us to denote each experiment k independently of the other experiments.

With all the above notation in mind, we can now create our hyperparameter tuning model, which is a bilevel optimisation problem. The leader is to select hyperparameters C and ϕ that minimise the average K -fold cross-validation error measured by the hinge loss. The follower wishes to select optimal model parameters w^k and b^k according to (Soft-Margin-SVM), for each fold k ,

that minimises the training error:

$$\begin{aligned}
& \underset{C, \phi, w^1, \dots, w^K, b^1, \dots, b^K}{\text{minimise}} && \sum_{k=1}^K \frac{1}{\bar{n}^k} \sum_{i=1}^{\bar{n}^k} \max(0, 1 - \bar{y}_i^k (w^k{}^\top \phi(\bar{x}_i^k) + b^k)) \\
& \text{subject to} && C \geq 0, \\
& && \text{For each } k = 1, \dots, K : \\
& && w^k, b^k \text{ optimal for SVM problem parametrised by } (x^k, y^k)_{k=1, \dots, n^k}.
\end{aligned} \tag{31}$$

For the problem to be well-posed, as before, we will convert the upper-level variables into dual ones, by means of the kernel trick. This can be done by substituting $w^k = \sum_j \alpha_j^k y_j^k \phi(x_j^k)$ into the objective function and applying the kernel trick $\bar{y}_i^k y_j^k \phi(\bar{x}_i^k) \phi(x_j^k) = \bar{Q}(\gamma)_{ij}^k$. A few other technical arrangements with the bias b^k will lead to the dual objective:

$$\sum_{k=1}^K \frac{1}{\bar{n}^k} \sum_{i=1}^{\bar{n}^k} \max \left(0, 1 - \sum_{j=1}^{n^k} \alpha_j^k \bar{Q}(\gamma)_{ij}^k - \bar{y}_i^k b^k(\alpha^k, \gamma, C) \right). \tag{32}$$

More precise details on the process for this conversion can be found in Appendix B.

To obtain a smooth upper-level objective function, we remove the max operator by introducing slack variables ζ_i and pairs of constraints $\zeta_i \geq 0$ and $\zeta_i \geq 1 - \sum \alpha_j^k \bar{Q}(\gamma)_{ij}^k - \bar{y}_i^k b^k(\alpha^k, \gamma, C)$ as before. This must be done for each fold $k = 1, \dots, K$ and for each validation example in those folds $i = 1, \dots, \bar{n}^k$. To remove the embedded lower-level optimisation problem, we perform a KKT reformulation. This results in the following MPCC:

$$\begin{aligned}
& \underset{C, \gamma, \zeta, \alpha, \underline{v}, \bar{v}, u}{\text{minimise}} && \sum_{k=1}^K \frac{1}{\bar{n}^k} \sum_{i=1}^{\bar{n}^k} \zeta_i^k \\
& \text{subject to} && C \geq 0, \\
& && \gamma \geq 0, \\
& && \zeta_i^k \geq 0 \\
& && \left. \begin{aligned} \zeta_i^k &\geq 1 - \sum_{j=1}^{n^k} \alpha_j^k \bar{Q}(\gamma)_{ij}^k - \bar{y}_i^k b^k(\alpha^k, \gamma, C) \end{aligned} \right\} \begin{array}{l} \text{(Validation)} \\ \text{for } k = 1, \dots, K, \\ \text{for } i = 1, \dots, \bar{n}^k, \end{array} \\
& && \left. \begin{aligned} \sum_{j=1}^{n^k} \alpha_j^k Q(\gamma)_{ij}^k - 1 - \underline{v}_i^k + \bar{v}_i^k + u^k y_i &= 0 \\ 0 \leq \alpha_i^k \perp \underline{v}_i^k &\geq 0 \\ 0 \leq (C - \alpha_i^k) \perp \bar{v}_i^k &\geq 0 \\ \sum_j \alpha_j^k y_i^k &= 0 \end{aligned} \right\} \begin{array}{l} \text{(Training)} \\ \text{for } k = 1, \dots, K, \\ \text{for } i = 1, \dots, n^k. \end{array}
\end{aligned} \tag{M}$$

where the Lagrange multipliers $\underline{v}^k, \bar{v}^k, u^k$ for $k = 1, \dots, K$ which were introduced in the KKT conditions (29) are now decision variables. The constraints resulting from the upper-level are tagged as **(Validation)** given that they measure the error over the validation data. The **(Training)** tag corresponds to the KKT conditions (29) of the lower-level problem.

Problem (\mathcal{M}) is explicit but not concise. It uses a two-dimensional indexing scheme: First, it indices the fold $k = 1, \dots, K$, then the individual examples i within that fold. We now wish to flatten the indices into a one-dimensional sequence. Hence, we introduce the index sets

$$I := \{(k, i) : k = 1, \dots, K, i = 1, \dots, n^k\} \quad \text{and} \quad \bar{I} := \{(k, i) : k = 1, \dots, K, i = 1, \dots, \bar{n}^k\}. \tag{33}$$

This allows us to stack all the constraint functions into vectors via the notations

$$Z_i^k(z) := \zeta_i^k - 1 + \sum_{j=1}^{n^k} \alpha_j^k \bar{Q}(\gamma)_{ij}^k + \bar{y}_i^k b^k(\alpha^k, \gamma, C) \quad \text{for } (k, i) \in \bar{I}, \quad (34a)$$

$$\theta_i^k(z) := \sum_{j=1}^{n^k} \alpha_j^k Q(\gamma)_{ij}^k - 1 - \underline{v}_i^k + \bar{v}_i^k + u y_i \quad \text{for } (k, i) \in I, \quad (34b)$$

$$\varphi^k(z) := \sum_j \alpha_j^k y_j^k \quad \text{for } k = 1, \dots, K. \quad (34c)$$

where we have also introduced $z := [C, \gamma, \zeta, \alpha, \underline{v}, \bar{v}, u]^\top \in \mathbb{R}^N$ to collect all variables, which has the dimension $N := 3(K-1)n + n + K + 2$. This leads to the representation $f(z) := \sum_{k=1}^K \frac{1}{n^k} \sum_{i=1}^{n^k} \zeta_i$ for the upper-level objective function of (\mathcal{M}) , as well as its constraint functions

$$g(z) := \begin{bmatrix} C \\ \gamma \\ \zeta_1^1 \\ \vdots \\ \zeta_{n^K}^K \\ Z_1^1(z) \\ \vdots \\ Z_{n^K}^K(z) \end{bmatrix}, \quad h(z) := \begin{bmatrix} \theta_1^1(z) \\ \vdots \\ \theta_{n^K}^K(z) \\ \varphi^1(z) \\ \vdots \\ \varphi^K(z) \end{bmatrix}, \quad G(z) := \begin{bmatrix} \alpha_1^1 \\ \vdots \\ \alpha_{n^K}^K \\ (C - \alpha_1^1) \\ \vdots \\ (C - \alpha_{n^K}^K) \end{bmatrix}, \quad H(z) := \begin{bmatrix} \underline{v}_1^1 \\ \vdots \\ \underline{v}_{n^K}^K \\ \bar{v}_1^1 \\ \vdots \\ \bar{v}_{n^K}^K \end{bmatrix}, \quad (35)$$

where we obviously have $g(z) \in \mathbb{R}^{2+2n}$, $h(z) \in \mathbb{R}^{(K-1)n+K}$ and $G(z), H(z) \in \mathbb{R}^{2(K-1)n}$. Subsequently, using the above constructions for the functions f, g, h, G and H , we can re-write problem (\mathcal{M}) more compactly in the general form (MPCC).

4.3 Theoretical properties of problem (\mathcal{M})

We start here by showing some properties of the matrices $Q(\gamma)^k$ for $k = 1, \dots, K$.

Lemma 2 *Given distinct data points $x_1^k, \dots, x_{n^k}^k \in \mathbb{R}^d$ and labels $y_1^k, \dots, y_{n^k}^k \in \{-1, 1\}$ for each $k = 1, \dots, K$ and $\gamma > 0$, we have:*

1. *Each RBF kernel matrix $Q(\gamma)^k \in \mathbb{R}^{n^k \times n^k}$, which is defined by $Q(\gamma)_{ij}^k := y_i^k y_j^k \exp(-\gamma \|x_i^k - x_j^k\|^2)$ for $i = 1, \dots, n$ and $j = 1, \dots, n$, is positive definite.*
2. *The following block diagonal matrix $\mathcal{Q}(\gamma)$ formed from $Q(\gamma)^1, \dots, Q(\gamma)^K$ is also positive definite:*

$$\mathcal{Q}(\gamma) := \begin{bmatrix} Q(\gamma)^1 & & 0 \\ & \ddots & \\ 0 & & Q(\gamma)^K \end{bmatrix}. \quad (36)$$

Proof The Gaussian $x \mapsto \exp(-\gamma \|x\|^2)$ is a positive definite function on \mathbb{R}^d [100, Theorem 6.10]. Therefore, $P_{ij}^k := \exp(-\gamma \|x_i^k - x_j^k\|^2)$ is a positive definite matrix [100, Definition 6.1]. Hence, for any non-zero vector $v \in \mathbb{R}^n$, by the positive definiteness of P^k , it holds that

$$v^\top Q(\gamma)^k v = \sum_{i=1}^n \sum_{j=1}^n v_i v_j y_i^k y_j^k \exp(-\gamma \|x_i^k - x_j^k\|^2) = (vy)^\top P^k(vy) > 0,$$

where $vy \in \mathbb{R}^n$ is the vector $(vy)_i = v_i y_i^k$ for $i = 1, \dots, n$. Thus, $Q(\gamma)^k$ is positive definite and so has only positive eigenvalues. Applying either [35, Definition 1.59, Property a] or Schur determinant formula [105] to the characteristic polynomial of $Q(\gamma)$ we get:

$$\det(Q(\gamma) - \lambda \mathcal{I}_n) = \prod_{k=1}^K \det(Q(\gamma)^k - \lambda \mathcal{I}_{n^k}).$$

Hence, $Q(\gamma)$ has all positive eigenvalues and is therefore positive definite. \square

Next, we make an important assumption that for each fold, there exists at least one dual variable in the interior of its bounds. This is equivalent to assuming at least one data point lies on the margin.

$$\forall k = 1, \dots, K, \quad \exists i \in 1, \dots, n^k \quad \text{such that } 0 < \alpha_i^k < C. \quad (37)$$

We can now show that the lower-level problem (Dual-SVM) satisfies the assumptions of Theorem 2, and in the provided framework, problems (\mathcal{M}) and (31) will be locally equivalent.

Proposition 1 *Problem (\mathcal{M}) has a convex lower level program, and satisfies the LLICQ at every point $z := [C, \gamma, \zeta, \alpha, v, \bar{v}, u]^\top$ where $C > 0$, $\gamma > 0$ and (37) holds.*

Proof By Lemma 2, the objective function of (Dual-SVM) is a positive quadratic function. Each of the constraints is linear in α . Therefore, we have convexity. For each index i , only one of the constraints $0 \leq \alpha_i$ or $\alpha_i \leq C$ can be active. Therefore, each of our active inequality constraints is in a unique variable. Furthermore, by assumption (37), there exists some index i such that the variable α_i is involved in the equality constraint $\sum \alpha_i y_i = 0$ but none of the active inequality constraints. Thus, their derivatives are linearly independent, and (\mathcal{M}) fulfils LLICQ. \square

Next, we give the derivatives of the constraint functions of (\mathcal{M}) . This will be useful later in studying which constraint qualifications problem (\mathcal{M}) satisfies.

Proposition 2 *The transposed Jacobian matrix of the constraint functions of (\mathcal{M}) is*

		Constraint functions									
		$g(z)$				$h(z)$		$G(z)$		$H(z)$	
		C	γ	ζ	$Z(z)$	$\theta(z)$	$\varphi(z)$	α	$(C - \alpha)$	v	\bar{v}
Derivative w.r.t.	$\nabla_C \bullet$	1	0	0	$\nabla_C Z(z)$	0	0	0	1	0	0
	$\nabla_\gamma \bullet$	0	1	0	$\nabla_\gamma Z(z)$	$\nabla_\gamma \theta(z)$	0	0	0	0	0
	$\nabla_\zeta \bullet$	0	0	\mathcal{I}	\mathcal{I}	0	0	0	0	0	0
	$\nabla_\alpha \bullet$	0	0	0	$\nabla_\alpha Z(z)$	$Q(\gamma)$	\mathcal{Y}	\mathcal{I}	$-\mathcal{I}$	0	0
	$\nabla_v \bullet$	0	0	0	0	$-\mathcal{I}$	0	0	0	\mathcal{I}	0
	$\nabla_{\bar{v}} \bullet$	0	0	0	0	\mathcal{I}	0	0	0	0	\mathcal{I}
	$\nabla_u \bullet$	0	0	0	0	\mathcal{Y}^\top	0	0	0	0	0

(38)

with the column headers denoting which constraint function (refer to (35)) does the column corresponds to; the row headers $\nabla_v \bullet$ denote that their row corresponds to the derivative with respect to decision variable v ; \mathcal{I} is the identity matrix with inferred dimension; and $\mathcal{Y} \in \{-1, 0, 1\}^{|I| \times K}$ is the block diagonal matrix formed by

$$\mathcal{Y} := \begin{bmatrix} y^1 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & y^K \end{bmatrix} \quad \text{and} \quad y^k := \begin{bmatrix} y_1^k \\ \vdots \\ y_{n^k}^k \end{bmatrix} \quad \text{for } k = 1, \dots, K.$$

Using the standard MPCC index sets I^g, I^G, I^H, I^{GH} that were defined in (1) becomes labour-some and hides the structure of our problem. Recalling the indexing scheme of (33), we define the problem-specific index sets

$$\bar{I}^C := \{(k, i) \in \bar{I} : \zeta_i^k = 0\} \quad \text{and} \quad \bar{I}^Z := \{(k, i) \in \bar{I} : Z_i^k(z) = 0\}.$$




Fig. 8: This figure illustrates the geometry of our problem by showing example feasible points for each of the five index sets C_1, C_2, C_3, C_4 and C_5 . The horizontal axis represents values of α_i^k . The vertical axis represent non-negative values of \underline{v}_i^k and \bar{v}_i^k when $\alpha_i^k = 0$ and $\alpha_i^k = C$, respectively. The orange arrows show directions in the normal cone of the feasible set of the problem.

We then partition the index set I into five cases that better describe (\mathcal{M}) :

$$\begin{aligned} C_1 &:= \{(k, i) \in I : \alpha_i^k = 0, \underline{v}_i^k > 0, (C - \alpha_i^k) = C, \bar{v}_i^k = 0\}; \\ C_2 &:= \{(k, i) \in I : \alpha_i^k = 0, \underline{v}_i^k = 0, (C - \alpha_i^k) = C, \bar{v}_i^k = 0\}; \\ C_3 &:= \{(k, i) \in I : \alpha_i^k \in (0, C), \underline{v}_i^k = 0, (C - \alpha_i^k) \in (0, C), \bar{v}_i^k = 0\}; \\ C_4 &:= \{(k, i) \in I : \alpha_i^k = C, \underline{v}_i^k = 0, (C - \alpha_i^k) = 0, \bar{v}_i^k = 0\}; \\ C_5 &:= \{(k, i) \in I : \alpha_i^k = C, \underline{v}_i^k = 0, (C - \alpha_i^k) = 0, \bar{v}_i^k > 0\}. \end{aligned}$$

We denote the union of multiple cases in the subscript. For example, $C_{123} := C_1 \cup C_2 \cup C_3$. It is easy to see that these sets are disjoint and cover all possible cases. That is, $C_{12345} = I$. Figure 8 illustrates the geometry of these index sets. Next, we show that (\mathcal{M}) satisfies MPCC-MFCQ-R at every feasible point with positive hyperparameters.

Proposition 3 *Let $z := [C, \gamma, \zeta, \alpha, \underline{v}, \bar{v}, u]^\top$ be any feasible point of problem (\mathcal{M}) such that $C > 0$, $\gamma > 0$ and (37) holds. Then MPCC-MFCQ-R holds at this point.*

Proof Recall from Lemma 1 that (MPCC) satisfies MPCC-MFCQ-R if and only if the family of gradients (A, B) defined in (7) is positive-linearly independent. We write these gradients explicitly for (\mathcal{M}) using the problem specific variables, index sets \bar{I}^ζ, \bar{I}^Z , and partition C_l for $l = 1, \dots, 5$:

$$A := \left\{ \nabla_z \alpha_i^k, \nabla_z \underline{v}_i^k : (k, i) \in C_2 \right\} \cup \left\{ \nabla_z (C - \alpha_i^k), \nabla_z \bar{v}_i^k : (k, i) \in C_4 \right\} \\ \cup \left\{ \nabla_z \zeta_i^k : (k, i) \in \bar{I}^\zeta \right\} \cup \left\{ \nabla_z Z_i^k(z) : (k, i) \in \bar{I}^Z \right\}; \quad (39a)$$

$$B := \left\{ \nabla_z \alpha_i^k : (k, i) \in C_1 \right\} \cup \left\{ \nabla_z \underline{v}_i^k : (k, i) \in C_{345} \right\} \cup \left\{ \nabla_z (C - \alpha_i^k) : (k, i) \in C_5 \right\} \\ \cup \left\{ \nabla_z \bar{v}_i^k : (k, i) \in C_{123} \right\} \cup \left\{ \nabla_z \theta_i^k(z) : (k, i) \in I \right\} \cup \left\{ \nabla_z \varphi^k(z) : k = 1, \dots, K \right\}. \quad (39b)$$

Suppose there exist non-negative multipliers $\lambda^1, \lambda^2, \lambda^3, \lambda^4, \lambda^5, \lambda^6 \geq 0$ corresponding to the vectors in set A and free multipliers $\mu^1, \mu^2, \mu^3, \mu^4, \mu^5, \mu^6$ corresponding to the vectors in set B such that together they give rise to a linear combination whose sum is zero. More precisely,

$$\begin{aligned} & \sum_{(k, i) \in C_2} (\lambda_{ki}^1 \nabla_z \alpha_i^k + \lambda_{ki}^2 \nabla_z \underline{v}_i^k) + \sum_{(k, i) \in C_4} (\lambda_{ki}^3 \nabla_z (C - \alpha_i^k) + \lambda_{ki}^4 \nabla_z \bar{v}_i^k) \\ & + \sum_{(k, i) \in \bar{I}^\zeta} \lambda_{ki}^5 \nabla_z \zeta_i^k + \sum_{(k, i) \in \bar{I}^Z} \lambda_{ki}^6 \nabla_z Z_i^k(z) + \sum_{(k, i) \in C_1} \mu_{ki}^1 \nabla_z \alpha_i^k + \sum_{(k, i) \in C_{345}} \mu_{ki}^2 \nabla_z \underline{v}_i^k \\ & + \sum_{(k, i) \in C_5} \mu_{ki}^3 \nabla_z (C - \alpha_i^k) + \sum_{(k, i) \in C_{123}} \mu_{ki}^4 \nabla_z \bar{v}_i^k + \sum_{(k, i) \in I} \mu_{ki}^5 \nabla_z \theta_i^k(z) + \sum_{k=1}^K \mu_{ki}^6 \nabla_z \varphi^k(z) = \mathbf{0}. \end{aligned} \quad (40)$$

We have written ∇_{z^\bullet} to denote the derivative with respect to the full decision variable $z := [C, \gamma, \zeta, \alpha, v, \bar{v}, u]^\top$. Moving forward, we will ignore the derivatives with respect to C and γ and write the remaining components of the derivative as vectors of five blocks corresponding to the derivatives with respect to $[\zeta, \alpha, v, \bar{v}, u]$. For example $\nabla_{[\zeta, \alpha, v, \bar{v}, u]}(\alpha_i^k) = [0, e_i^k, 0, 0, 0]^\top$ where e_i^k is the standard unit vector of zeros with a 1 at index (k, i) . Thus we take (40) and rewrite explicitly

$$\begin{aligned}
& \sum_{(k,i) \in C_2} \left(\lambda_{ki}^1 \begin{bmatrix} 0 \\ e_i^k \\ 0 \\ 0 \\ 0 \end{bmatrix} + \lambda_{ki}^2 \begin{bmatrix} 0 \\ 0 \\ e_i^k \\ 0 \\ 0 \end{bmatrix} \right) + \sum_{(k,i) \in C_4} \left(\lambda_{ki}^3 \begin{bmatrix} 0 \\ -e_i^k \\ 0 \\ 0 \\ 0 \end{bmatrix} + \lambda_{ki}^4 \begin{bmatrix} 0 \\ 0 \\ 0 \\ e_i^k \\ 0 \end{bmatrix} \right) \\
& + \sum_{(k,i) \in \bar{I}^Z} \lambda_{ki}^5 \begin{bmatrix} e_i^k \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \sum_{(k,i) \in \bar{I}^\zeta} \lambda_{ki}^6 \begin{bmatrix} e_i^k \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \sum_{(k,i) \in C_1} \mu_{ki}^1 \begin{bmatrix} 0 \\ e_i^k \\ 0 \\ 0 \\ 0 \end{bmatrix} + \sum_{(k,i) \in C_{345}} \mu_{ki}^2 \begin{bmatrix} 0 \\ 0 \\ e_i^k \\ 0 \\ 0 \end{bmatrix} \\
& + \sum_{(k,i) \in C_5} \mu_{ki}^3 \begin{bmatrix} 0 \\ -e_i^k \\ 0 \\ 0 \\ 0 \end{bmatrix} + \sum_{(k,i) \in C_{123}} \mu_{ki}^4 \begin{bmatrix} 0 \\ 0 \\ 0 \\ e_i^k \\ 0 \end{bmatrix} + \sum_{k=1}^K \mu_{ki}^5 \begin{bmatrix} 0 \\ Y^k \\ 0 \\ 0 \\ 0 \end{bmatrix} + \sum_{(k,i) \in I} \mu_{ki}^6 \begin{bmatrix} 0 \\ Q(\gamma)_i^k \\ -e_i^k \\ e_i^k \\ Y \end{bmatrix} = \mathbf{0}. \tag{41}
\end{aligned}$$

Inspecting the first row of the vector derivatives, which represents the derivative with respect to ζ , we can have that $\lambda^5 + \lambda^6 = \mathbf{0}$. By their non-negativity we conclude $\lambda^5 = \mathbf{0}$ and $\lambda^6 = \mathbf{0}$.

We continue by inspecting subvectors of the multipliers based on C_{12345} . For the vector $\mu \in \mathbb{R}^{|I|}$ and matrix $Q \in \mathbb{R}^{|I| \times |I|}$, denote the subvector and submatrix according to index set $C_l \subseteq I$ as

$$\mu_{C_l} := \mu_{[(k,i) \in C_l]} \in \mathbb{R}^{|C_l|} \text{ and } Q_{C_l, C_{l'}} := Q(\gamma)_{[(k,i) \in C_l, (k',i') \in C_{l'}]} \in \mathbb{R}^{|C_l| \times |C_{l'}|},$$

respectively, for $l = 1, \dots, 5$. Note that Q_{C_l, C_l} is a principal submatrix of a positive definite matrix (Lemma 2); therefore, it is positive definite itself. It is still a function of gamma, but we drop the argument for brevity.

Consider the derivative with respect to v_i^k for $(k, i) \in C_2$; i.e., the third row in our block notation in (41). We have that $\lambda^2 - \mu_{C_2}^6 = \mathbf{0}$, which implies that $\mu_{C_2}^6$ is non-negative. Consider the derivative with respect to α_i^k for $(k, i) \in C_2$. i.e. the second row block in (41). We have $\lambda^1 + Q_{C_2, C_2} \mu_{C_2}^6 = \mathbf{0}$. Taking the inner product with $\mu_{C_2}^6$ gives

$$(\mu_{C_2}^6)^\top (\lambda^1) + (\mu_{C_2}^6)^\top (Q_{C_2, C_2}) (\mu_{C_2}^6) = 0.$$

By the positive definiteness of Q_{C_2, C_2} and since $\lambda^1, \lambda^2, \mu_{C_2}^6 \geq 0$ we must have $\lambda^1 = \lambda^2 = \mu_{C_2}^6 = 0$.

In a similar manner, we now consider the derivative with respect to \bar{v}_i^k for $(k, i) \in C_4$; i.e., the fourth row block in (41). We have that $\lambda^4 + \mu_{C_4}^6 = \mathbf{0}$ which implies that $\mu_{C_4}^6$ is non-positive. Considering the derivative with respect to α_i^k for $(k, i) \in C_4$, we have $-\lambda^3 + Q_{C_4, C_4} \mu_{C_4}^6 = \mathbf{0}$. Taking the inner product with $\mu_{C_4}^6$ leads to

$$-(\mu_{C_4}^6)^\top (\lambda^3) + (\mu_{C_4}^6)^\top (Q_{C_4, C_4}) (\mu_{C_4}^6) = 0.$$

By the positive definiteness of the principal submatrix Q_{C_4, C_4} of Q , and given that $\lambda^3, \lambda^4 \geq 0$ and $\mu_{C_4}^6 \leq 0$, we must have the equalities $\lambda^3 = \lambda^4 = \mu_{C_4}^6 = \mathbf{0}$.

At this point, we have shown that $\lambda^1, \lambda^2, \lambda^3, \lambda^4, \lambda^5, \lambda^6, \mu_{C_2}^6$ and $\mu_{C_4}^6$ are all zero. Now, consider the remaining derivatives in (41) whose multipliers we have not yet concluded are zero. We stack these vectors as columns of the matrix

	1 $(\theta_i^k(z))$ $(k,i) \in C_1$	2 $(\theta_i^k(z))$ $(k,i) \in C_3$	3 $(\theta_i^k(z))$ $(k,i) \in C_5$	4 $(\varphi^k(z))$ $k=1,\dots,K$	5 (α_i^k) $(k,i) \in C_1$	6 $(C - \alpha_i^k)$ $(k,i) \in C_5$	7 (\underline{v}_i^k) $(k,i) \in C_{345}$	8 (\bar{v}_i^k) $(k,i) \in C_{123}$
1 $(\nabla_{\alpha_i^k} \bullet)_{(k,i) \in C_1}$	\mathcal{Q}_{C_1, C_1}	\mathcal{Q}_{C_3, C_1}	\mathcal{Q}_{C_5, C_1}	\mathcal{Y}_{C_1}	\mathcal{I}_{C_1}	0	0	0
2 $(\nabla_{\alpha_i^k} \bullet)_{(k,i) \in C_3}$	\mathcal{Q}_{C_1, C_3}	\mathcal{Q}_{C_3, C_3}	\mathcal{Q}_{C_5, C_3}	\mathcal{Y}_{C_3}	0	0	0	0
3 $(\nabla_{\alpha_i^k} \bullet)_{(k,i) \in C_5}$	\mathcal{Q}_{C_1, C_5}	\mathcal{Q}_{C_3, C_5}	\mathcal{Q}_{C_5, C_5}	\mathcal{Y}_{C_5}	0	$-\mathcal{I}_{C_5}$	0	0
4 $(\nabla_{\underline{v}_i^k} \bullet)_{(k,i) \in C_1}$	\mathcal{I}_{C_1}	0	0	0	0	0	0	0
5 $(\nabla_{\underline{v}_i^k} \bullet)_{(k,i) \in C_{345}}$	0	$[-\mathcal{I}_{C_3}, 0, 0]^\top$	$[0, 0, -\mathcal{I}_{C_5}]^\top$	0	0	0	$\mathcal{I}_{C_{345}}$	0
6 $(\nabla_{\bar{v}_i^k} \bullet)_{(k,i) \in C_{123}}$	$[\mathcal{I}_{C_1}, 0, 0]^\top$	$[0, 0, \mathcal{I}_{C_3}]^\top$	0	0	0	0	0	$\mathcal{I}_{C_{123}}$
7 $(\nabla_{\bar{v}_i^k} \bullet)_{(k,i) \in C_5}$	0	0	\mathcal{I}_{C_5}	0	0	0	0	0
8 $(\nabla_{\varphi^k} \bullet)_{k=1,\dots,K}$	$\mathcal{Y}_{C_1}^\top$	$\mathcal{Y}_{C_3}^\top$	$\mathcal{Y}_{C_5}^\top$	0	0	0	0	0

that we denote by (\tilde{M}) , which is a submatrix of the transposed Jacobian in (38). The top-left element of (\tilde{M}) has column header $(\theta_i^k(z))_{(k,i) \in C_1}$ and row header $(\nabla_{\alpha_i^k} \bullet)_{(k,i) \in C_1}$. This helps the reader to observe that this element is the derivative of constraint function $(\theta_i^k(z))$ over indices $(k,i) \in C_1$ w.r.t. variables α_i^k over indices $(k,i) \in C_1$, which in this case is the submatrix \mathcal{Q}_{C_1, C_1} .

From here, we perform a sequence of elementary row operations to demonstrate that (\tilde{M}) is full rank. These are described in five steps. First, the row-blocks are enumerated 1, 2, 3, 4, 5, 6, 7, 8 are reordered to 1, 3, 5, 6, 2, 4, 7, 8. Secondly, the column-blocks are enumerated 1, 2, 3, 4, 5, 6, 7, 8 are reordered to 5, 6, 7, 8, 2, 1, 3, 4. Thirdly, we multiply rows in row-block 3 by -1 so that its leading entries are positive.

Observe that \mathcal{Q}_{C_3, C_3} has full rank, so there exist elementary row operations that map \mathcal{Q}_{C_3, C_3} into row echelon form. Denote with R the product of matrices that perform those elementary row operations so that $R\mathcal{Q}_{C_3, C_3}$ is in row echelon form. The fourth step is to premultiply row-block 2 by R .

Finally, take the rows of 8 and subtract scaled copies of rows of 2, 4 and 7 according to Gaussian elimination until the rows of 8 are zero in all columns of 2, 1, and 3. This leaves some non-zero matrix which we will denote S in row-block 8, column-block 4. Since the individual rows $k = 1, \dots, K$ of row-block 8 correspond to the same folds $k = 1, \dots, K$ as the columns $k = 1, \dots, K$ of column-block 4, the matrix S is diagonal. Furthermore, assumption (37) gives us that C_3 is non-empty for each fold $k = 1, \dots, K$; therefore, the diagonal elements of S are non-zero. After these five steps of row operations, the resulting matrix can be written in the following way:

	5 (α_i^k) $(k,i) \in C_1$	6 $(C - \alpha_i^k)$ $(k,i) \in C_5$	7 (\underline{v}_i^k) $(k,i) \in C_{345}$	8 (\bar{v}_i^k) $(k,i) \in C_{123}$	2 $(\theta_i^k(z))$ $(k,i) \in C_3$	1 $(\theta_i^k(z))$ $(k,i) \in C_1$	3 $(\theta_i^k(z))$ $(k,i) \in C_5$	4 $(\varphi^k(z))$ $k=1,\dots,K$
1 $(\nabla_{\alpha_i^k} \bullet)_{(k,i) \in C_1}$	\mathcal{I}_{C_1}	0	0	0	\mathcal{Q}_{C_3, C_1}	\mathcal{Q}_{C_1, C_1}	\mathcal{Q}_{C_5, C_1}	\mathcal{Y}_{C_1}
3 $-(\nabla_{\alpha_i^k} \bullet)_{(k,i) \in C_5}$	0	\mathcal{I}_{C_5}	0	0	$-\mathcal{Q}_{C_3, C_5}$	$-\mathcal{Q}_{C_1, C_5}$	$-\mathcal{Q}_{C_5, C_5}$	$-\mathcal{Y}_{C_5}$
5 $(\nabla_{\underline{v}_i^k} \bullet)_{(k,i) \in C_{345}}$	0	0	$\mathcal{I}_{C_{345}}$	0	$[-\mathcal{I}_{C_3}, 0, 0]^\top$	0	$[0, 0, -\mathcal{I}_{C_5}]^\top$	0
6 $(\nabla_{\bar{v}_i^k} \bullet)_{(k,i) \in C_{123}}$	0	0	0	$\mathcal{I}_{C_{123}}$	$[0, 0, \mathcal{I}_{C_3}]^\top$	$[\mathcal{I}_{C_1}, 0, 0]^\top$	0	0
2 $R(\nabla_{\alpha_i^k} \bullet)_{(k,i) \in C_3}$	0	0	0	0	$R\mathcal{Q}_{C_3, C_3}$	$R\mathcal{Q}_{C_1, C_3}$	$R\mathcal{Q}_{C_5, C_3}$	$R\mathcal{Y}_{C_3}$
4 $(\nabla_{\underline{v}_i^k} \bullet)_{(k,i) \in C_1}$	0	0	0	0	0	\mathcal{I}_{C_1}	0	0
7 $(\nabla_{\bar{v}_i^k} \bullet)_{(k,i) \in C_5}$	0	0	0	0	0	0	\mathcal{I}_{C_5}	0
8 \dagger	0	0	0	0	0	0	0	S

From this sequence of row reductions, we see that each column has a unique row index with its leading non-zero entry. Therefore, the matrix (\tilde{M}) has full rank. The equation (40) is satisfied for non-negative λ only when $\lambda = \mathbf{0}, \mu = \mathbf{0}$, giving us that (A, B) is positive-linearly independent and (\mathcal{M}) satisfies MPCC-MFCQ-R. \square

Remark 3 It is worth remarking that the above proof uses the dual form of MPCC-MFCQ-R. It would be equivalent to show for the sets defined in (39a,b) that B is linearly independent and there exists a direction d such that $Bd = \mathbf{0}$ and $Ad > \mathbf{0}$. This direction d is illustrated geometrically

by the orange arrows in Figure 8. Also, as a key consequence of Proposition 3, the assumptions of Theorem 3 and Theorem 5 hold accordingly.

We now show that problem (\mathcal{M}) in general fails the stronger constraint qualification MPCC-MFCQ-T, and therefore displaying that the second implication in Theorem 1 is strict for the application problem under consideration in this section.

Proposition 4 *The MPCC-MFCQ-T fails at all feasible points $z := (C, \gamma, \zeta, \alpha, \bar{v}, \underline{v}, u)$ of the program (\mathcal{M}) with at least two biactive indices; i.e., with $|I^{GH}(z)| > 2$.*

Proof MPCC-MFCQ-T requires that the following family of vectors be linearly independent:

$$\left\{ \nabla h_i(z) : i \in I^h \right\} \cup \left\{ \nabla G_i(z) : i \in I^G(z) \cup I^{GH}(z) \right\} \cup \left\{ \nabla H_i(z) : i \in I^H(z) \cup I^{GH}(z) \right\}. \quad (42)$$

Here, each vector is a derivative with respect to the decision variable z .

The constraint functions h , G and H are in the variable $(\alpha, \underline{v}, \bar{v}, u, C, \gamma)$, which has a dimension of $3(K-1)n + K + 2$, where n is the number of training examples and K is the number of folds in the cross-validation. Observe that all the vectors in (42) are constant with respect to ζ , so their derivatives are zero with respect to ζ . There are a total of $3(K-1) + K + |I^{GH}|$ vectors in this set (42). Therefore, if $|I^{GH}| > 2$, then there are more vectors than non-zero dimensions, and so they cannot be linearly independent. \square

By Theorem 1, we can conclude that MPCC-LICQ also fails at all these points. This program is a good practical example to show that the implication $\text{MPCC-MFCQ-T} \implies \text{MPCC-MFCQ-R}$ is strict. We have now shown that (\mathcal{M}) satisfies the relevant constraint qualifications to employ the solution methods from Section 3. Many of these propositions are conditional on the regularisation hyperparameter C being strictly positive. This is a very sensible restriction as any dataset where $C = 0$ or $\gamma = 0$ is optimal implies that the best classifier ignores the data entirely.

5 Numerical experiments on hyperparameter tuning for support vector machines

In this section, we present our numerical experiments on hyperparameter tuning for support vector machines where we solve large instances of the highly non-convex MPCC problem (\mathcal{M}) using the solution methods presented in Section 3 under various settings.

5.1 Implementation details

We begin by describing the hardware, software, and approach used for the implementation of the algorithms. Each program and system of equations we wish to solve, such as (Dual-SVM), (\mathcal{M}) , (\mathcal{R}_τ) , (\mathcal{P}_π) and (58), are written in the algebraic modelling language AMPL [30]. This allows us to robustly rearrange the equations and use automatic differentiation to calculate the derivatives of the objective and constraint functions. Furthermore, we exploit the fact that AMPL offers interfaces to all the solvers we wish to use. The algorithms from Section 3 are implemented in Python [95] (version 3.12) with use of the library NumPy [42] (version 1.26.4) for efficient linear algebra computation. All graphics are produced using Matplotlib Pyplot [46] (version 3.8.2) and seaborn [99] (version 0.13.2).

For solving the subproblems (\mathcal{R}_τ) and (\mathcal{P}_π) , we use two different solvers: Interior Point Optimizer (IPOPT) [97], an open source solver that implements a primal-dual interior point method with line search, and Artelys' Knitro [9], a commercial solver that implements an interior point trust region method. Both solve reliably for our model though our results are stated for Knitro as it seems to show faster runtimes. For solving a mixed system of equations with complementary terms, as studied by Michael Ferris and Steven Dirkse [22], we use the PATH solver [21, 23, 24]. All of these solvers interface with AMPL models and derivatives. Our experiments run on Southampton University's Iridis Compute Cluster [72]. Each experiment is scheduled on a single Intel Xeon

Dataset	n	d	Ref.	Description
Banknote	1372	4	[62]	Images of banknotes labelled as either genuine or forged.
BUPA	345	6	[29]	A collection of blood tests labelled by liver disorder diagnosis.
Circles	90	2	-	Coordinates of a positive class encircled by a negative class.
Cleveland	297	14	[20]	Various patient attributes used to predict heart disease.
Glass	146	9	[36]	Oxide content and refractive index of glass samples.
Hastie	100	10	[43]	This is a famous example from [43, equation (10.2)].
Ionosphere	351	34	[85]	Antennas radar signals used to classify structure in the ionosphere.
Iris	100	4	[1, 26]	Length and width of sepals and petals labelled by species.
Linear	100	10	-	Synthetically generated linearly separable dataset.
MNIST	600	784	[53]	Images of handwritten digits.
Mushroom	300	117	[61]	Observations of mushrooms labelled as either poisonous or edible.
Moons	54	2	-	Coordinates sampled from two interweaving crescent moons.
Palmer	342	5	[45]	Observations of penguins used to classify their species.
Pima	768	8	[86]	BMI, insulin and other features used in the diagnosis of diabetes.
Swiss roll	90	3	[66]	See [66, Figure 6.13] where one class is rolled inside another.
Sonar	208	60	[39]	Marine sonar signals used to detect sea mines.
Wisconsin	569	32	[93, 101]	Measurements of breast tumours labelled as malignant or benign.
XOR	100	2	-	Positive label assigned if and only if exactly one coordinate is positive.

Table 1: Summary of the eighteen datasets used throughout our numerical experiments.

Gold 6138 CPU. This CPU is compatible with most modern home PCs. It has 2 processors at 2.0 GHz/processor and twelve 16GB TruDDR4 2666 MHz RDIMM of installed RAM.

In our experiments, we focus on the RBF kernel (28) and the computation of the hyperparameters C and γ forming the upper-level variables of (\mathcal{M}) . It is common in the literature to set $K = 3$ (see [37, page 91]) and $K = 5$ (see [43, Section 7.10.1]) as the number of folds for K-fold cross-validation. We set $K = 3$ because it gives a more stable validation accuracy for small datasets. The stopping criterion for the subprograms is gradually decreased to a tolerance of 10^{-6} . The overarching (MPCC) loop in Algorithms 1 and 2 is terminated when one of each pair of complementary constraint functions is at most 10^{-6} . This condition can be written as

$$\max_i (\min(|G_i(z)|, |H_i(z)|)) \leq 10^{-6}. \quad (43)$$

For the penalisation method, we select the initial penalty parameter to be 100 and sequentially increase it by multiples of 10. For the relaxation method, we select the initial relaxation parameter to be 0.01 and sequentially decrease it by multiples of 10^{-1} .

Many publications have suggested using the so-called *exact penalisation* (see e.g. [81], resp. *exact relaxation* see e.g. [44]) approach. Here, the parameter π (resp. τ) is fixed at some chosen value, and only a single instance of the subproblem (\mathcal{P}_π) (resp. (\mathcal{R}_τ)) is solved. In our experiments here, we compare this exact approach to the sequential one presented in Algorithm 1 (resp. 2). To do so, we use values for the parameter π (resp. τ) resulting in an MPCC feasible solution to the corresponding subproblem. It is worth keeping in mind that, in practice, it is hard to find a good value for such a parameter.

We now present the data used for our experiments. We use eighteen datasets altogether; cf. Table 1. For each of them, we state a value for n , the number of examples, and the number of features d . Six of these (namely, *Circles*, *Hastie*, *Linear*, *Moons*, *Swiss roll* and *XOR*) are synthetically generated with code from scikit-learn [76]. This is useful as globally and locally, optimal classifiers can be known from inspection, but could still be hard for a SVM to find. This allows us to test the solution methods objectively. The remaining twelve datasets come from real-world observations. In these cases, the optimal values are unknown, and we must compare our solutions to those achieved by other authors in the literature. We perform standard preprocessing on these datasets before passing them to the solvers. All rows with missing values are dropped. For continuous features, we subtract the mean and divide by the standard deviation. The resulting scaled data has a mean of zero and a standard deviation of one. This helps avoid scaling issues that can cause our problem to become ill-conditioned.

Regarding categorical features, we either drop the column entirely or apply a one-hot encoding where 0 – 1 indicator variables are added when they are relevant to the classification. Finally,




Fig. 9: This heat map shows the hyperparameter space of the Moons dataset. Green regions correspond to (C, γ) combinations, which allow the SVM to achieve high validation accuracy. The eight coloured lines represent the path of iterates produced by the Relaxation method from eight different starting points (triangles) to various stationary points (stars).

we partition the multi-class labels into binary classes, which we replace with the labels -1 and +1. A better classifier can sometimes be achieved through principle component analysis or other feature manipulation, however, we stick to these very simple preprocessing steps to allow a direct comparison to the original papers where the data was published. Further details on the datasets and modifications that we made to them are provided in Appendix F.

When (\mathcal{M}) is parametrised with these data sets, it is large in comparison to other non-convex MPCCs that have been solved in the literature so far. The median problem size is 1773 variables, whereas for example, see Leyffer’s famous MacMPEC collection [55], which has a median problem size of 49 variables.

5.2 Starting point selection

In this subsection, we discuss multi-starts and initialisation strategies for the algorithms in Section 3, since each requires the user to provide a starting point $z^0 = [C^0, \gamma^0, \zeta^0, \alpha^0, \underline{v}^0, \bar{v}^0, u^0]^\top$ in the context of our problem (\mathcal{M}) . The methods are deterministic but perform very differently depending on the choice of z^0 . Recall also that the penalisation and relaxation methods converge to S and C-stationary points, respectively. Such points may not be global or even local minima. Therefore, initialising the methods from a range of different starting points can increase the chance of finding one. However, it is well-known that under a MPCC-specific second-order sufficient condition [81, Definition 2.8], S-stationary points correspond to strict local minima [81, Theorem 3.1]. We remind the reader here of the theoretical advantage that the penalisation method has by converging to S-stationary points and their potential to be local minima.

We propose a multi-start, where the upper-level variable (C^0, γ^0) shall be initiated according to a grid spread uniformly across the log space such as $\{10^n : n = -5, \dots, 4\} \times \{10^n : n = -3, \dots, 6\}$, for example. This encourages the solution paths to explore as much of the hyperparameter space as possible. Note that this grid is far coarser than the derivative-free grid search we compare against later. Figure 9 shows a two-by-four grid of starting points and the resulting path of iterates that the sequential penalisation algorithm produced for the Moons dataset. Of these, three converge to the global minimum, while the remaining five converge to other stationary points.

Considering that the variable ζ is only involved in two constraints, given an initialisation of α^0 , we can simply choose

$$\zeta_i^{0k} = \max \left(0, 1 - \sum_{j=1}^{n^k} \alpha_j^{0k} \bar{Q}(\gamma)_{ij}^k - \bar{y}_i^k b^k(\alpha^{0k}, \gamma^0, C^0) \right) \quad \text{for } k = 1 \dots K, i = 1, \dots, \bar{n}^k,$$

which ensures initial feasibility of the constraints $\zeta_i^{0k} \geq 0$ and $Z_i^k(z^0) \geq 0$; see (34a).

It remains to be discussed how to initialise the lower-level variables α^0 , \underline{v}^0 , \bar{v}^0 and u^0 . We propose three strategies. The first and simplest strategy shall be referred to as **LL-Zero**. This strategy initialises all of the lower-level variables to zero. Notice this satisfies the complementary constraints $0 \leq \alpha_i^k \perp \underline{v}_i^k \geq 0$ and $0 \leq (C - \alpha_i^k) \perp \bar{v}_i^k \geq 0$ but violates the stationarity constraint $\theta_i^k(z) = 0$ for each $(k, i) \in I$. The strategy has no computational cost and can perform well but must be treated with caution as it is infeasible for all the programs (\mathcal{M}) , (\mathcal{P}_π) and (\mathcal{R}_τ) .

We call the second strategy **LL-Centre**, which, for a given choice of C^0 and γ^0 , selects values for the lower-level variables as follows:

$$\alpha_i^{0k} = \frac{C^0}{2n^k} \sum_{j=1}^{n^k} (1 + y_i^k y_j^k), \quad \underline{v}_i^{0k} = \max \left(0, \sum_{j=1}^{n^k} \alpha_j^{0k} Q(\gamma)_{ij}^k - 1 + y_i \right), \quad (44)$$

$$u^{0k} = 1, \quad \bar{v}_i^{0k} = \max \left(0, - \sum_{j=1}^{n^k} \alpha_j^{0k} Q(\gamma)_{ij}^k + 1 - y_i \right) \quad (45)$$

for $k = 1, \dots, K$, $i = 1, \dots, n^k$. By construction, $\theta_i^k(z^0) = 0$ and $\varphi_i^k(z^0) = 0$ for $k = 1, \dots, K$, $i = 1, \dots, n^k$; see (34). For a balanced dataset, where there is a similar number of positive and negative labels y_i^k in each fold, the variables α_i^k are initialised close to $\frac{C}{2}$, the midpoint of its feasible interval $0 \leq \alpha_i^k \leq C$, for $k = 1, \dots, K$, $i = 1, \dots, n^k$. Since one of \bar{v} or \underline{v} is also positive, the complementary constraints are violated. This means that the initialisation is not MPCC feasible for problem (\mathcal{M}) but is feasible for the corresponding penalisation (\mathcal{P}_{π^0}) and relaxation (\mathcal{R}_{τ^0}) formulations for suitable choices of π^0 and τ^0 .

Finally, we call the third strategy **LL-Feasible**, which corresponds, for a given choice of (C^0, γ^0) , to find α^0 , \underline{v}^0 , \bar{v}^0 and u^0 that satisfy the KKT conditions (29), characterising the lower-level problem. This can be done by calling the PATH solver [21, 23, 24] K times, where K is the number of folds, to solve each of the K lower-level KKT systems. It is the most computationally expensive initialisation but ensures z^0 is feasible for all the programs (\mathcal{M}) , (\mathcal{P}_{π^0}) and (\mathcal{R}_{τ^0}) . The three strategies can be summarised in Table 2.

	LL-Zero	LL-Centre	LL-Feasible
Relative computational cost	low	low	high
Feasible z^0 for (\mathcal{M}) ?	no	no	yes
Feasible z^0 for $(\mathcal{P}_{\pi^0})/(\mathcal{R}_{\tau^0})$?	no	yes	yes

Table 2: Summary of the lower-level initialisation strategies.

When a feasible initialisation is made (i.e. **LL-Feasible**), at each index i , the complementarity choice (for example, choosing $G_i(z^0) = 0$ or $H_i(z^0) = 0$) is pre-decided. We observe that, for many indices, this is unlikely to flip throughout the iterations of the solution methods. However, when an infeasible initialisation (i.e. **LL-Centre**, **LL-Zero**) is chosen that violates the complementarity constraint (like having $G_i(z^0) = 1$ and $H_i(z^0) = 1$, for example), we observe that the same algorithms are more free to choose between $G_i(z^0) = 0$ and $H_i(z^0) = 0$ in the convergence process and arrive at better solutions. This phenomenon is demonstrated in the following experiment: For each dataset and solution method, we run 75 starting points, with 25 according to each of the strategies **LL-Zero**, **LL-Feasible** and **LL-Centre**. We then assess the quality of the solution by the distribution of objective values found. These results for dataset *Circles* are summarised as box plots in




Fig. 10: The columns correspond to fold $k = 1$ and indices $i = 1, 2, 3, 4$ (Left to right). The rows correspond to sequentially smaller relaxation parameters τ^t (top to bottom). The feasible regions for the relaxed program (\mathcal{R}_{τ^t}) are coloured in green. Each of the 16 small plots has a red point for the value of the complementary variables α_i (horizontal axis) and v_i (vertical axis) that were found to solve (\mathcal{R}_{τ^t}).

Figure 11, where the effect of the lower-level initialisation is particularly pronounced for the relaxation method. The algorithms are sensitive to the initial choice of (C, γ) . This is due to the highly non-convex nature of the hyperparameter space. The choice of lower-level variable initialisation only has a very small effect on most datasets. However, overall, it is clear that **LL-Centre** encourages the algorithms to converge to higher quality solutions and have a low computational cost. Therefore, we will use the **LL-Centre** initialisation strategy for the remainder of our experiments.

An interesting visualisation of the iteration of the sequential relaxation method solving problem (\mathcal{M}), parametrised by the Ionosphere dataset, is given in Figure 10. This shows the values of the complementary variables α_i and v_i for the first fold and first four indices $i = 1, 2, 3, 4$. Note that the variables are initially infeasible for the MPCC but feasible for the subproblem (\mathcal{R}_{τ}). As we proceed down the rows, the variables z converge to either the axis corresponding to $\alpha_i = 0$ or the axis corresponding to $v_i = 0$, and therefore generating an MPCC feasible final solution.

5.3 Numerical behaviour of penalisation and relaxation methods

Here, we discuss the behaviour of the methods from Section 3 when applied to problem (\mathcal{M}). First, for the sequential relaxation method, we observe that the relaxation parameter almost always decreases to 10^{-6} . That is, at each iteration t with $\tau^t = 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}$, the solutions found by Algorithm 2 for the subproblem (\mathcal{R}_{τ^t}) are not MPCC feasible (i.e. (43) does not hold) and the algorithm only terminates when the relaxation parameter reaches 10^{-6} and the complementary constraints are forced within the exit tolerance. On the other hand, the penalisation method terminates with a range of different penalty parameters between 10^3 and 10^6 . That is, for those values of π^t , the solvers find solutions to the subproblem (\mathcal{P}_{π^t}) that are MPCC feasible and so satisfy the exit condition. The distribution of these observations at termination is presented in Table 3. The value for MPCC feasibility, more precisely, the maximum violation over the constraints of (\mathcal{M}) is also given. Further research should investigate the structure of (\mathcal{M}) to quantify why the penalisation approach is able to find an MPCC feasible solution before the penalty parameter π^t potentially gets too large and causes (\mathcal{P}_{π^t}) to become ill-conditioned (e.g. [57, Example 3]), as it is usually observed for penalty-based algorithms [70, 103].

Note that the complementarity constraint functions of (\mathcal{M}) are all linear. Also, observe that for all indices $k = 1, \dots, K, i = 1, \dots, n^k$, the equality constraint (34b) could be rearranged and used to substitute out either \underline{v}_i^k or \bar{v}_i^k from the complementary constraints as is done by Coniglio et al. [14]. This would reduce the number of variables and constraints by $(K-1)n$. However, this would introduce a nonlinear term in the complementary constraints. In line with the advice of Fletcher

	Penalisation (Sequential)	Relaxation (Sequential)
Parameter π/τ Q_1	10^4	10^{-6}
Parameter π/τ Q_2	10^6	10^{-6}
Parameter π/τ Q_3	10^6	10^{-6}
MPCC Feasibility Q_1	$3.94e-07$	$1.00e-06$
MPCC Feasibility Q_2	$1.93e-08$	$9.98e-07$
MPCC Feasibility Q_3	$4.88e-09$	$2.01e-07$

Table 3: This table presents the observed distributions of final parameter π and τ as well as the MPCC feasibility at termination. Q_1, Q_2, Q_3 denote the first quartile, median and third quartile.




Fig. 11: These box plots summarise the distribution of objective values achieved when using each of the three different lower-level initialisation strategies for the Circles dataset.

and Leyffer [28, 56], we find that keeping the complementary constraint linear is more beneficial than reducing the number of variables via substitution.

Next, observe that the bilevel optimisation approach for hyperparameter optimisation, studied in this paper, sometimes overfits to the validation data. To monitor this, on top of the training and validation data that parameterises problem (\mathcal{M}), we require further data that is entirely unseen by the model; this will be called *test* data. For each of the synthetic datasets, this is generated according to the same distribution. For each of the real-world datasets, the test data is formed by putting aside 10% of examples and parametrising with the remaining 90%; refer back to Figure 7. The overfitting can then be observed when the classifiers, represented by the solutions of problem (\mathcal{M}), achieve higher validation accuracy than test accuracy. Evidence of this is presented in our later experiments, in particular, see the Ionosphere and Moons datasets in Table 4.

5.4 Performance evaluation in comparison to derivative-free optimisation

Across industry and academia, Derivative-Free Optimisation (DFO) remains by far the most used tool for hyperparameter tuning [66]. It is simple and effective. The defining characteristic of DFO is that it treats training as a black box, meaning it does not use any information from the training process other than the final evaluation of the model in optimising the hyperparameters. We compare the bilevel optimisation to four such methods: *grid search*, *random search*, *Bayesian optimisation*, and *pattern search*. For more details and references to the algorithms we used see Appendix E.

Our bilevel optimisation method has two clear advantages over DFO. The first advantage comes from the idea that two choices of hyperparameters with similar values often result in similar SVM states. For example, an optimal SVM with $(C, \gamma) = (15, 0.3)$ is likely to have many of the same support vectors as the optimal SVM with $(C, \gamma) = (12, 0.2)$. This means only a few decision variables α_i^k for $(k, i) \in I$ must be changed to move between local optima for the adjusted hyperparameters. When one of grid, random, Bayesian or pattern search moves between two nearby hyperparameter selections in this manner, it recalculates the optimal decision variables from scratch, whereas the bilevel optimisation method only adjusts the required variables. The second advantage comes from the fact that the bilevel optimisation approach makes use of the first and second-order derivatives of the objective function with respect to the hyperparameters to move in a descent direction through the hyperparameter space.

Note that (Dual-SVM) has n variables and $2n$ constraints. On the other hand, the MPCC model (\mathcal{M}) has $7n + 5$ variables, $2n$ inequality constraints, 3 equality constraints and $4n$ complementary constraints. Therefore, the comparison must be made as to whether better performance can be achieved by minimising the much smaller problem (Dual-SVM) many times according to DFO or by minimising the much larger (\mathcal{M}) with the solution methods from Section 3. This is answered empirically in the next subsection.

In the following, we evaluate both the exact and sequential type penalisation and relaxation methods from Section 3 and the derivative-free methods according to their runtime and quality of solution in solving problem (\mathcal{M}). Furthermore, the semi-smooth Newton method (see Appendix D) is added to the experiments. This is an entirely different approach to solving MPCCs and converges to an M-stationary point, making it an interesting comparison, given that the sequential




Fig. 12: These box plots show the distribution of objective values achieved by initialising from a range of different starting points. The four axes represent the datasets (left to right): Ionosphere, Palmer, Moons and Wisconsin. Within each axis, the box plots represent the solution method (left to right): 1. Penalisation (Exact), 2. Penalisation (Sequential), 3. Relaxation (Exact), 4. Relaxation (sequential), 5. Semi-Smooth Newton, 6. Grid Search, 7. Random Search, 8. Bayesian Optimisation, 9. Pattern Search.




Fig. 13: This performance profile graph plots the proportion of instances that were solved to the target test accuracy (vertical axis) by a given method using restarts (annotated 1 through 9) within the given time (horizontal axis).

penalisation and Scholtes relaxation algorithms that we primarily study in this paper have been shown to instead converge to S- and C-stationarity, respectively.

For each of the eighteen datasets listed in Table 1, 10% of the data is reserved for testing. Then, each of the solution methods are parametrised by the remaining 90% of data and run one hundred times from different initialisations. This whole process is repeated ten times with a different 10% sample reserved for testing. The runtime, objective function value and final test accuracy are measured. These datasets are selected because they have quite different structures. The optimal (C, γ) hyperparameter choice being for Ionosphere ($1.00e + 01, 3.16e - 02$), Palmer ($1.00e + 06, 1.00e - 05$), Moons ($1.78e + 01, 5.62e + 00$) and Wisconsin ($3.16e + 03, 1.78e - 04$). Furthermore, they show various different spreads of objective values and runtimes.

It is important to note that the penalisation method consistently converges faster than the other methods across all datasets. Additionally, the sequential approach, when compared to the exact approach, often achieves a smaller range of objective function values. To demonstrate this, we document the results of four datasets as quantiles in Table 4 and box-plots in Figure 12. It would be inappropriate to compare a single run of the penalisation, relaxation or semismooth Newton method to the derivative-free techniques. We see that the bilevel-based methods run much faster, but, on the other hand, converge to a large variety of stationary points, many of which have poor objective function values. To compare them on a fair footing, we attribute a target test accuracy to each of the datasets. We then multiply the mean runtime by $\frac{1}{p}$ where p is the proportion of initialisation from which the method achieves the target test accuracy. This adjusted runtime estimates the expected time to achieve the target test accuracy by repeatedly restarting the method from random initialisations. For example, if the average runtime is 10s but only one in five of the initialisations achieves the target test accuracy, then the adjusted runtime is recorded as 50s. These results are plotted as a performance profile in Figure 13.

		Penalisation (Exact)	Penalisation (Sequential)	Relaxation (Exact)	Relaxation (Sequential)	Semismooth Newton	Grid	Random	Bayesian	Pattern
Ionosphere	Runtime Q_1	4	9	16	55	386	64	69	400	83
	Runtime Q_2	11	12	29	90	589	98	106	880	132
	Runtime Q_3	17	17	40	149	1369	135	136	2152	133
	Objective Q_1	1.05	0.84	1.01	1.01	1.32	0.40	0.40	0.35	1.01
	Objective Q_2	0.62	0.47	0.42	0.40	0.87	0.37	0.38	0.33	0.47
	Objective Q_3	0.54	0.33	0.40	0.34	0.75	0.35	0.36	0.33	0.33
	Iterations	188	208	546	1569	891	49	50	50	8
	Validation acc.	84%	89%	86%	88%	79%	88%	87%	89%	89%
Palmer	Test acc.	83%	89%	86%	86%	79%	88%	87%	89%	89%
	Runtime Q_1	1	106	111	200	579	71	73	389	105
	Runtime Q_2	1	133	168	616	798	108	111	968	160
	Runtime Q_3	27	201	569	808	1639	143	143	2200	206
	Objective Q_1	0.89	0.69	0.91	0.42	1.26	0.20	0.04	0.04	0.90
	Objective Q_2	0.80	0.36	0.19	0.19	0.63	0.09	0.04	0.03	0.90
	Objective Q_3	0.44	0.03	0.19	0.19	0.46	0.04	0.04	0.03	0.04
	Iterations	54	235	290	849	516	49	50	50	8
Moons	Validation acc.	90%	99%	95%	95%	90%	99%	99%	99%	99%
	Test acc.	90%	99%	95%	95%	90%	99%	99%	99%	99%
	Runtime Q_1	1	2	2	6	3	49	49	447	72
	Runtime Q_2	1	2	2	8	4	82	104	678	103
	Runtime Q_3	1	3	3	16	112	104	111	979	217
	Objective Q_1	0.95	0.92	0.49	0.49	0.76	0.35	0.35	0.32	0.81
	Objective Q_2	0.60	0.38	0.48	0.32	0.54	0.34	0.34	0.31	0.47
	Objective Q_3	0.56	0.31	0.34	0.31	0.48	0.32	0.33	0.31	0.31
Wisconsin	Iterations	65	335	348	1125	665	49	50	50	8
	Validation acc.	83%	89%	88%	89%	74 %	89%	89%	89%	89%
	Test acc.	83%	87%	87%	89%	74 %	89%	89%	87%	87%
	Runtime Q_1	18	59	147	198	317	68	72	172	77
	Runtime Q_2	27	70	247	467	547	82	85	210	97
	Runtime Q_3	47	83	369	560	1431	94	95	591	104
	Objective Q_1	0.76	0.09	0.40	0.17	0.96	0.10	0.09	0.08	1.00
	Objective Q_2	0.52	0.08	0.18	0.13	0.51	0.09	0.08	0.08	0.12
Wisconsin	Objective Q_3	0.27	0.08	0.18	0.12	0.39	0.08	0.08	0.08	0.08
	Iterations	85	230	852	1565	979	49	50	50	8
	Validation acc.	93%	97%	94%	96%	90%	97%	97%	97%	96%
	Test acc.	93%	97%	94%	96%	90%	97%	97%	97%	96%

Table 4: A summary of the solution methods (columns) performance on each of the datasets (rows). The first quartile (Q_1) is the runtime/objective value achieved by 75% of initialisations; the second quartile (Q_2) is the median; and the third quartile (Q_3) is the runtime/objective value achieved by the top 25% of initialisations.

Overall, the obtained results make a convincing argument that our model (\mathcal{M}), being solved with the sequential (partial) penalisation method, can consistently outperform derivative-free methods in the tuning of SVM hyperparameters. The exact relaxation outperforms the derivative-free methods on a majority of instances. The semi-smooth Newton method performs well on small problems but fails to converge for larger problems - it seems to get stuck at the non-differentiable corners of the function F that arise from the use of the min operator. Between the derivative-free methods, grid and random search perform equally well on all but the very largest problems for which Bayesian optimisation is preferable. The penalisation methods clearly come out on top overall. Sequential penalisation is the fastest method to achieve the target test accuracy in 88% of instances. For the other 12%, exact penalisation is the fastest. However, it must be noted that, in practice, a user may prefer to always use sequential penalisation as it does not require the selection of a penalty parameter a priori.

6 Conclusion

We began by highlighting the distinction between the definitions of MPCC-MFCQ-T and MPCC-MFCQ-R, which is often overlooked in the literature. From there, we presented three solution methods for the (MPCC). For the relaxation and penalisation methods, we provided new convergence proofs that an inexact sequence of KKT points under the assumption of MPCC-MFCQ-R converge to S-stationary and C-stationary points, respectively. This provides a generalisation over the existing literature and allows the methods to be applied to a broader range of problems.

In the paper’s second half, we presented our model (\mathcal{M}) for the hyperparameter tuning of support vector machines. We then proved that this satisfies the MPCC-MFCQ-R, thus satisfying the assumptions of the theorems proved earlier in the paper. Finally, we ran an extensive set of experiments on real-world data sets. In this context, we showed the practical use of our methods in beating the derivative-free search that is most commonly used. We conclude that the sequential penalisation method both converges faster and is more likely to converge to a higher quality objective value than the other methods.

For future work, we will be looking to find more general sufficient conditions for the assumption $\lim_{t \rightarrow \infty} (G(z^t)^\top H(z^t)) = 0$ in Theorem 3 to hold. Additionally, we will develop first-order gradient descent methods to solve (\mathcal{M}) such as those being studied in the work of Bengio [5] and Chapelle et al. [12]. These have a distinct advantage in machine learning applications as they can be efficiently parallelised on modern graphical processing units (GPUs). Finally, in future works, we will develop a stochastic model where the gradients are approximated by subsampling a small number of data points. This will allow the model to scale to much larger datasets.

Acknowledgments

Thank you to Qingna Li from the Beijing Institute of Technology for her correction regarding assumption (37).

References

1. Anderson, E.: The species problem in Iris. *Annals of the Missouri Botanical Garden* **23**, 457–509 (1936). <https://www.biodiversitylibrary.org/part/4079>
2. Anitescu, M.: On solving mathematical programs with complementarity constraints as nonlinear programs. *Mathematics in Computer Science* (2002). <https://api.semanticscholar.org/CorpusID:16203968>
3. Arrow, K.J., Debreu, G.: Existence of an equilibrium for a competitive economy. *Econometrica* **22**(3), 265–290 (1954). <http://www.jstor.org/stable/1907353>
4. Baumrucker, B., Renfro, J.G., Biegler, L.T.: MPEC problem formulations and solution strategies with chemical engineering applications. *Computers & Chemical Engineering* **32**(12), 2903–2913 (2008). <https://doi.org/10.1016/j.compchemeng.2008.02.010>
5. Bengio, Y.: Gradient-based optimization of hyperparameters. *Neural Computation* **12**(8), 1889–1900 (2000). <https://doi.org/10.1162/089976600300015187>
6. Bennett, K.P., Hu, J., Ji, X., Kunapuli, G., Pang, J.S.: Model selection via bilevel optimization. In: *The 2006 IEEE International Joint Conference on Neural Network Proceedings*, pp. 1922–1929. IEEE (2006). <https://doi.org/10.1109/IJCNN.2006.246935>
7. Boser, B.E., Guyon, I.M., Vapnik, V.N.: A training algorithm for optimal margin classifiers. In: *Proceedings of the fifth annual workshop on Computational learning theory*, pp. 144–152 (1992). <https://doi.org/10.1145/130385.130401>
8. Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press (2004). <https://doi.org/10.1017/CB09780511804441>
9. Byrd, R.H., Nocedal, J., Waltz, R.A.: KNITRO: An integrated package for nonlinear optimization. In: G. Di Pillo, M. Roma (eds.) *Large-Scale Nonlinear Optimization, Nonconvex Optim. Appl.*, vol. 83, pp. 35–59. Springer US, Boston, MA (2006). https://doi.org/10.1007/0-387-30065-1_4
10. Cervantes, J., Garcia-Lamont, F., Rodríguez-Mazahua, L., Lopez, A.: A comprehensive survey on support vector machine classification: Applications, challenges and trends. *Neurocomputing* **408**, 189–215 (2020). <https://doi.org/10.1016/j.neucom.2019.10.118>
11. Chang, C.C., Lin, C.J.: LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)* **2**(3), 1–27 (2011)
12. Chapelle, O., Vapnik, V., Bousquet, O., Mukherjee, S.: Choosing multiple parameters for support vector machines. *Machine Learning* **46**, 131–159 (2002). <https://doi.org/10.1023/A:1012450327387>

13. Chen, Y., Florian, M.: The nonlinear bilevel programming problem: formulations, regularity and optimality conditions. *Optimization* **32**(3), 193–209 (1995). <https://doi.org/10.1080/02331939508844048>
14. Coniglio, S., Dunn, A., Li, Q., Zemkoho, A.: Bilevel hyperparameter optimization for nonlinear support vector machines (2022). Unpublished
15. Cortes, C., Vapnik, V.: Support-vector networks. *Machine Learning* **20**, 273–297 (1995). <https://doi.org/10.1007/BF00994018>
16. DeMiguel, V., Friedlander, M.P., Nogales, F.J., Scholtes, S.: A two-sided relaxation scheme for mathematical programs with equilibrium constraints. *SIAM Journal on Optimization* **16**(2), 587–609 (2005). <https://doi.org/10.1137/04060754x>
17. Dempe, S., Dutta, J.: Is bilevel programming a special case of a mathematical program with complementarity constraints? *Mathematical Programming* **131**, 37–48 (2012). <https://doi.org/10.1007/s10107-010-0342-1>
18. Dempe, S., Zemkoho, A.: On the Karush–Kuhn–Tucker reformulation of the bilevel optimization problem. *Nonlinear Analysis: Theory, Methods & Applications* **75**(3), 1202–1218 (2012). <https://doi.org/10.1016/j.na.2011.05.097>
19. Dempe, S., Zemkoho, A.: Bilevel optimization. In: Springer Optimization and its Applications, vol. 161. Springer (2020). <https://link.springer.com/book/10.1007/978-3-030-52119-6>
20. Detrano, R.C., János, A., Steinbrunn, W., Pfisterer, M.E., Schmid, J.J., Sandhu, S., Guppy, K., Lee, S., Froelicher, V.: International application of a new probability algorithm for the diagnosis of coronary artery disease. *American Journal of Cardiology* **64** 5, 304–10 (1989). <https://api.semanticscholar.org/CorpusID:23545303>
21. Dirkse, S., Ferris, M.C., Munson, T.: The PATH solver. <https://pages.cs.wisc.edu/~ferris/path.html>
22. Dirkse, S.P., Ferris, M.C.: MCPLIB: A collection of nonlinear mixed complementarity problems. *Optimization Methods and Software* **5**(4), 319–345 (1995). <https://doi.org/10.1080/10556789508805619>
23. Dirkse, S.P., Ferris, M.C.: The PATH solver: A non-monotone stabilization scheme for mixed complementarity problems. *Optimization Methods and Software* **5**(2), 123–156 (1995). <https://doi.org/10.1080/10556789508805606>
24. Dirkse, S.P., Ferris, M.C.: A path search damped Newton method for computing general equilibria. *Annals of Operations Research* **68**(2), 211–232 (1996). <https://doi.org/10.1007/BF02209613>
25. Dutta, J., Deb, K., Tulshyan, R., Arora, R.: Approximate KKT points and a proximity measure for termination. *Journal of Global Optimization* **56**(4), 1463–1499 (2013). <https://doi.org/10.1007/s10898-012-9920-5>
26. Fisher, R.A.: The use of multiple measurements in taxonomic problems. *Annals of Eugenics* **7**(2), 179–188 (1936). <https://doi.org/10.1111/j.1469-1809.1936.tb02137.x>
27. Flegel, M.L.: Constraint qualifications and stationarity concepts for mathematical programs with equilibrium constraints. Ph.D. thesis, Universität Würzburg (2005)
28. Fletcher, R., Leyffer, S.: Numerical experience with solving MPECs as NLPs. Numerical Analysis Report NA/210, Department of Mathematics, University of Dundee, Dundee, UK **8** (2002)
29. Forsyth, R.S.: Bupa liver disorders. BUPA Medical Research Ltd. (1990). <http://archive.ics.uci.edu/ml/machine-learning-databases/liver-disorders/bupa.data>
30. Fourer, R., Gay, D.M., Kernighan, B.W.: AMPL: A mathematical programming language. *Management Science* **36**(5), 519–554 (1990). https://www.ampl.com/_archive/first-website/REFS/amplmod.pdf
31. Franceschi, L., Frasconi, P., Salzo, S., Grazzi, R., Pontil, M.: Bilevel programming for hyperparameter optimization and meta-learning. In: J. Dy, A. Krause (eds.) *Proceedings of the 35th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol. 80, pp. 1568–1577. PMLR (2018). <https://proceedings.mlr.press/v80/franceschi18a.html>
32. Gabriel, S.A., Leuthold, F.U.: Solving discretely-constrained mpec problems with applications in electric power markets. *Energy Economics* **32**(1), 3–14 (2010). <https://doi.org/10.1016/j.eneco.2009.03.008>
33. Gao, L.L., Ye, J., Yin, H., Zeng, S., Zhang, J.: Value function based difference-of-convex algorithm for bilevel hyperparameter selection problems. In: K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, S. Sabato (eds.) *Proceedings of the 39th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol. 162, pp. 7164–7182. PMLR (2022). <https://proceedings.mlr.press/v162/gao22j.html>
34. Garnett, R.: Bayesian optimization. Cambridge University Press (2023)
35. George, R.K., Ajayakumar, A.: *A Course in Linear Algebra*. Springer (2024). <https://doi.org/10.1007/978-981-99-8680-4>
36. German, B.: Glass identification. UCI Machine Learning Repository (1987). <https://doi.org/10.24432/C5WW2P>
37. Géron, A.: *Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow*. O'Reilly Media, Inc. (2022). <https://www.oreilly.com/library/view/hands-on-machine-learning/9781492032632/>
38. Gfrerer, H., Ye, J.J.: New constraint qualifications for mathematical programs with equilibrium constraints via variational analysis. *SIAM Journal on Optimization* **27**(2), 842–865 (2017). <https://doi.org/10.1137/16M1088752>
39. Gorman, R.P., Sejnowski, T.J.: Analysis of hidden units in a layered network trained to classify sonar targets. *Neural Networks* **1**(1), 75–89 (1988). [https://doi.org/10.1016/0893-6080\(88\)90023-8](https://doi.org/10.1016/0893-6080(88)90023-8)
40. Guyon, I., Boser, B., Vapnik, V.: Automatic capacity tuning of very large VC-dimension classifiers. In: S. Hanson, J. Cowan, C. Giles (eds.) *Advances in Neural Information Processing Systems*, vol. 5. Morgan-Kaufmann (1992). https://proceedings.neurips.cc/paper_files/paper/1992/file/eaee339c4d89fc102edd9dbdb6a28915-Paper.pdf
41. Harder, F., Mehlitz, P., Wachsmuth, G.: Reformulation of the m-stationarity conditions as a system of discontinuous equations and its solution by a semismooth Newton method. *SIAM Journal on Optimization* **31**(2), 1459–1488 (2021). <https://doi.org/10.1137/20M1321413>

42. Harris, C.R., Millman, K.J., van der Walt, S.J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N.J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M.H., Brett, M., Haldane, A., del Río, J.F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., Oliphant, T.E.: Array programming with NumPy. *Nature* **585**(7825), 357–362 (2020). <https://doi.org/10.1038/s41586-020-2649-2>
43. Hastie, T., Tibshirani, R., Friedman, J.H., Friedman, J.H.: The elements of statistical learning: data mining, inference, and prediction, vol. 2. Springer (2009). <https://doi.org/10.1007/978-0-387-84858-7>
44. Hoheisel, T., Kanzow, C., Schwartz, A.: Theoretical and numerical comparison of relaxation methods for mathematical programs with complementarity constraints. *Mathematical Programming* **137**, 257–288 (2013). <https://doi.org/10.1007/s10107-011-0488-5>
45. Horst, A.M., Hill, A.P., Gorman, K.B.: palmerpenguins: Palmer Archipelago (Antarctica) penguin data (2020). <https://doi.org/10.5281/zenodo.3960218>
46. Hunter, J.D.: Matplotlib: A 2d graphics environment. *Computing in Science & Engineering* **9**(3), 90–95 (2007). <https://doi.org/10.1109/MCSE.2007.55>
47. Jiang, W., Siddiqui, S.: Hyper-parameter optimization for support vector machines using stochastic gradient descent and dual coordinate descent. *EURO Journal on Computational Optimization* **8**(1), 85–101 (2020). <https://doi.org/10.1007/s13675-019-00115-7>
48. Kim, Y., Leyffer, S., Munson, T.: MPEC methods for bilevel optimization problems. In: S. Dempe, A. Zemkoho (eds.) *Bilevel Optimization: Advances and Next Challenges*, pp. 335–360. Springer International Publishing, Cham (2020). https://doi.org/10.1007/978-3-030-52119-6_12
49. Klein, A., Falkner, S., Bartels, S., Hennig, P., Hutter, F.: Fast Bayesian optimization of machine learning hyperparameters on large datasets. In: A. Singh, J. Zhu (eds.) *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research*, vol. 54, pp. 528–536. PMLR (2017). <https://proceedings.mlr.press/v54/klein17a.html>
50. Kunapuli, G., Bennett, K., Hu, J., Pang, J.S.: Bilevel model selection for support vector machines. In: CRM proceedings and lecture notes, vol. 45, pp. 129–158 (2008)
51. Kunapuli, G., Bennett, K., Hu, J., Pang, J.S.: Classification model selection via bilevel programming. *Optimization Methods and Software* **23**(4), 475–489 (2008). <https://doi.org/10.1080/10556780802102586>
52. Lawphongpanich, S., Hearn, D.W.: An MPEC approach to second-best toll pricing. *Mathematical Programming* **101**(1), 33–55 (2004). <https://doi.org/10.1007/s10107-004-0536-5>
53. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998). <https://doi.org/10.1109/5.726791>
54. Lessmann, S., Stahlbock, R., Crone, S.F.: Optimizing hyperparameters of support vector machines by genetic algorithms. In: H.R. Arabnia, R. Joshua (eds.) *Proceedings of the 2005 International Conference on Artificial Intelligence, ICAI 2005, Las Vegas, Nevada, USA, June 27-30, 2005, Volume 1*, pp. 74–82. CSREA Press (2005)
55. Leyffer, S.: MacMPEC. <https://wiki.mcs.anl.gov/leyffer/index.php/MacMPEC>
56. Leyffer, S.: Complementarity constraints as nonlinear equations: Theory and numerical experience. In: S. Dempe, V. Kalashnikov (eds.) *Optimization with Multivalued Mappings: Theory, Applications, and Algorithms*, pp. 169–208. Springer US, Boston, MA (2006). https://doi.org/10.1007/0-387-34221-4_9
57. Leyffer, S., López-Calva, G., Nocedal, J.: Interior methods for mathematical programs with complementarity constraints. *SIAM Journal on Optimization* **17**(1), 52–77 (2006). <https://doi.org/10.1137/040621065>
58. Li, Q., Li, Z., Zemkoho, A.: Bilevel hyperparameter optimization for support vector classification: theoretical analysis and a solution method. *Mathematical Methods of Operations Research* **96**(3), 315–350 (2022). <https://doi.org/10.1007/s00186-022-00798-6>
59. Li, Z., Qian, Y., Li, Q.: A unified framework and a case study for hyperparameter selection in machine learning via bilevel optimization. In: 2022 5th International Conference on Data Science and Information Technology (DSIT), pp. 1–8. IEEE (2022)
60. Lin, G.H., Fukushima, M.: A modified relaxation scheme for mathematical programs with complementarity constraints. *Annals of Operations Research* **133**(1), 63–84 (2005). <https://doi.org/10.1007/s10479-004-5024-z>
61. Lincoff, G.H., Knopf, A.A.: Mushroom. *The Audubon Society Field Guide to North American Mushrooms* (1981). <https://doi.org/10.24432/C5959T>
62. Lohweg, V.: Banknote authentication. UCI Machine Learning Repository (2013). <https://doi.org/10.24432/C55P57>
63. Luo, Z.Q., Pang, J.S., Ralph, D.: *Mathematical Programs with Equilibrium Constraints*. Cambridge University Press (1996). <https://www.cambridge.org/core/books/mathematical-programs-with-equilibrium-constraints/03981C32ABDD55A4001BF58BA0C57444>
64. Mangasarian, O.L.: *Nonlinear programming*. McGraw-Hill, New York (1969)
65. Mantovani, R.G., Rossi, A.L., Vanschoren, J., Bischl, B., De Carvalho, A.C.: Effectiveness of random search in SVM hyper-parameter tuning. In: 2015 International Joint Conference on Neural Networks (IJCNN), pp. 1–8. Ieee (2015). <https://doi.org/10.1109/IJCNN.2015.7280664>
66. Marsland, S.: *Machine learning: an algorithmic perspective*. Chapman and Hall/CRC (2011). <https://github.com/hongzhonglu/machine-learning-books/blob/master/Machine%20Learning%20-%20An%20Algorithmic%20Perspective%202nd%20edition%202014.pdf>
67. Moore, G., Bergeron, C., Bennett, K.P.: Model selection for primal SVM. *Machine Learning* **85**, 175–208 (2011). <https://doi.org/10.1007/s10994-011-5246-7>
68. Moore, G.M.: *Bilevel programming algorithms for machine learning model selection*. Ph.D. thesis, Rensselaer Polytechnic Institute (2010). <https://www.proquest.com/dissertations-theses/bilevel-programming-algorithms-machine-learning/docview/860000550/se-2>

69. Moore, G.M., Bergeron, C., Bennett, K.P.: Nonsmooth bilevel programming for hyperparameter selection. In: 2009 IEEE International Conference on Data Mining Workshops, pp. 374–381. IEEE (2009). <https://doi.org/10.1109/ICDMW.2009.74>
70. Murray, W.: Analytical expressions for the eigenvalues and eigenvectors of the Hessian matrices of barrier and penalty functions. *Journal of Optimization Theory and Applications* **7**(3), 189–196 (1971). <https://doi.org/10.1007/BF00932477>
71. Nelder, J.A., Mead, R.: A Simplex Method for Function Minimization. *The Computer Journal* **7**(4), 308–313 (1965). <https://doi.org/10.1093/comjnl/7.4.308>
72. Nicole, D., Takeda, K., Wolton, I., Cox, S.: Southampton high performance computing centre. In: High-Performance Computing, pp. 33–41. Springer US, Boston, MA (1999). https://doi.org/10.1007/978-1-4615-4873-7_4
73. Nocedal, J., Wright, S.J.: Numerical Optimization. Springer (1999). <https://doi.org/10.1007/978-0-387-40065-5>
74. Okuno, T., Takeda, A., Kawana, A., Watanabe, M.: On lp-hyperparameter learning via bilevel nonsmooth optimization. *Journal of Machine Learning Research* **22**(245), 1–47 (2021). <http://jmlr.org/papers/v22/18-485.html>
75. Pang, J.S., Fukushima, M.: Complementarity constraint qualifications and simplified B-stationarity conditions for mathematical programs with equilibrium constraints. *Computational Optimization and Applications* **13**, 111–136 (1999). <https://doi.org/10.1023/A:1008656806889>
76. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Édouard Duchesnay: Scikit-learn: Machine learning in python. *Journal of Machine Learning Research* **12**(85), 2825–2830 (2011). <http://jmlr.org/papers/v12/pedregosa11a.html>
77. Pelikan, M., Goldberg, D.E., Cantú-Paz, E.: BOA: The Bayesian optimization algorithm. In: Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation - Volume 1, GECCO'99, p. 525–532. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1999). <https://dl.acm.org/doi/10.5555/2933923.2933973>
78. Qi, L., Wei, Z.: On the constant positive linear dependence condition and its application to SQP methods. *SIAM Journal on Optimization* **10**(4), 963–981 (2000). <https://doi.org/10.1137/S1052623497326629>
79. Qiu, S., Chen, Z.: Solving mathematical programs with equilibrium constraints as nonlinear programming: A new framework. *arXiv: Optimization and Control* (2015). <https://api.semanticscholar.org/CorpusID:119660481>
80. Rahman, M.S., Rahman, M.K., Kaykobad, M., Rahman, M.S.: isGPT: An optimized model to identify subgolgi protein types using SVM and random forest based feature selection. *Artificial intelligence in medicine* **84**, 90–100 (2018)
81. Ralph*, D., Wright, S.J.: Some properties of regularization and penalization schemes for MPECs. *Optimization Methods and Software* **19**(5), 527–556 (2004). <https://doi.org/10.1080/10556780410001709439>
82. Scheel, H., Scholtes, S.: Mathematical programs with complementarity constraints: Stationarity, optimality, and sensitivity. *Mathematics of Operations Research* **25**(1), 1–22 (2000). <https://doi.org/10.1287/moor.25.1.1.15213>
83. Scholtes, S.: Convergence properties of a regularization scheme for mathematical programs with complementarity constraints. *SIAM Journal on Optimization* **11**(4), 918–936 (2001). <https://doi.org/10.1137/S1052623499361233>
84. Schwartz, A.: Mathematical programs with complementarity constraints: Theory, methods and applications. Ph.D. thesis, Universität Würzburg (2011). https://tu-dresden.de/mn/math/numerik/schwartz/ressourcen/dateien/Schwartz_Dissertation.pdf
85. Sigillito, V., Wing, S., Hutton, L., Baker, K.: Ionosphere. UCI Machine Learning Repository (1989). <https://doi.org/10.24432/C5W01B>
86. Smith, J.W., Everhart, J.E., Dickson, W., Knowler, W.C., Johannes, R.S.: Using the ADAP learning algorithm to forecast the onset of diabetes mellitus. In: Proceedings of the annual symposium on computer application in medical care, p. 261. American Medical Informatics Association (1988)
87. Smith, N., Gales, M.: Speech recognition using SVMs. In: T. Dietterich, S. Becker, Z. Ghahramani (eds.) *Advances in Neural Information Processing Systems*, vol. 14. MIT Press (2001). https://proceedings.neurips.cc/paper_files/paper/2001/file/d8330f857a17c53d217014ee776bfd50-Paper.pdf
88. Snoek, J., Larochelle, H., Adams, R.P.: Practical Bayesian optimization of machine learning algorithms. In: F. Pereira, C. Burges, L. Bottou, K. Weinberger (eds.) *Advances in Neural Information Processing Systems*, vol. 25. Curran Associates, Inc. (2012). https://proceedings.neurips.cc/paper_files/paper/2012/file/05311655a15b75fab86956663e1819cd-Paper.pdf
89. Solodov, M.V.: Constraint Qualifications. John Wiley & Sons, Ltd (2011). <https://doi.org/10.1002/9780470400531.eorms0978>
90. Stackelberg, H.v.: Marktform und gleichgewicht. *The Economic Journal* **45**(178), 334–336 (1934). <https://doi.org/10.2307/2224643>
91. Steffensen, S., Ulbrich, M.: A new relaxation scheme for mathematical programs with equilibrium constraints. *SIAM Journal on Optimization* **20**(5), 2504–2539 (2010). <https://doi.org/10.1137/090748883>
92. Stone, M.: Cross-validatory choice and assessment of statistical predictions. *Journal of the royal statistical society: Series B (Methodological)* **36**(2), 111–133 (1974). <https://doi.org/10.1111/j.2517-6161.1974.tb00994.x>
93. Street, W.N., Wolberg, W.H., Mangasarian, O.L.: Nuclear feature extraction for breast tumor diagnosis. In: *Electronic imaging* (1993). <https://api.semanticscholar.org/CorpusID:14922543>

94. Tefas, A., Kotropoulos, C., Pitas, I.: Using support vector machines to enhance the performance of elastic graph matching for frontal face authentication. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **23**(7), 735–746 (2001). <https://doi.org/10.1109/34.935847>
95. Van Rossum, G., Drake, F.L.: *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA (2009). <https://dl.acm.org/doi/book/10.5555/1593511>
96. Vapnik, V.: *The nature of statistical learning theory*. Springer science & business media (1999). <https://doi.org/10.1007/978-1-4757-3264-1>
97. Wächter, A., Biegler, L.T.: On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming* **106**, 25–57 (2005). <https://doi.org/10.1007/s10107-004-0559-y>
98. Wainer, J., Fonseca, P.: How to tune the RBF SVM hyperparameters? an empirical evaluation of 18 search algorithms. *Artificial Intelligence Review* **54**(6), 4771–4797 (2021). <https://doi.org/10.1007/s10462-021-10011-5>
99. Waskom, M.L.: seaborn: statistical data visualization. *Journal of Open Source Software* **6**(60), 3021 (2021). <https://doi.org/10.21105/joss.03021>
100. Wendland, H.: *Scattered data approximation*, vol. 17. Cambridge university press (2004). <https://doi.org/10.1017/CB09780511617539>
101. Wolberg, W., Mangasarian, O., Street, N., Street, W.: *Breast cancer wisconsin (diagnostic)*. UCI Machine Learning Repository (1995). <https://doi.org/10.24432/C5DW2B>
102. Wolpert, D.H.: The lack of a priori distinctions between learning algorithms. *Neural Computation* **8**(7), 1341–1390 (1996). <https://doi.org/10.1162/neco.1996.8.7.1341>
103. Wright, M.H.: Some properties of the Hessian of the logarithmic barrier function. *Mathematical Programming* **67**(1), 265–295 (1994). <https://doi.org/10.1007/BF01582224>
104. Ye, J.J.: Necessary and sufficient optimality conditions for mathematical programs with equilibrium constraints. *Journal of Mathematical Analysis and Applications* **307**(1), 350–369 (2005). <https://doi.org/10.1016/j.jmaa.2004.10.032>
105. Zhang, F.: *The Schur complement and its applications*, vol. 4. Springer Science & Business Media (2006). <https://doi.org/10.1007/b105056>

A Table of notation

Parameters	
γ, δ	$:=$ Hyperparameters that determine the kernel.
C	$:=$ Regularisation hyperparameters for (Soft-Margin-SVM).
π^t	$:=$ Penalty parameter at iteration t .
τ^t	$:=$ Relaxation parameter at iteration t .
Training and Validation Data	
x_i	$:=$ Data point i .
x_i^k	$:=$ Training data point i in cross fold experiment k .
\bar{x}_i^k	$:=$ Validation data point i in cross fold experiment k .
y_i	$:=$ Label for data point i .
y_i^k	$:=$ Label for training data point i in cross fold experiment k .
\bar{y}_i^k	$:=$ Label for validation data point i in cross fold experiment k .
n	$:=$ Total number of data point.
n^k	$:=$ Number of training data point in cross fold experiment k .
\bar{n}^k	$:=$ Number of validation data point in cross fold experiment k .
$Q(\gamma)_{ij}$	$:=$ $y_i y_j \exp(-\gamma \ x_i - x_j\)$ the full RBF kernel matrix.
$Q(\gamma)_{ij}^k$	$:=$ $y_i^k y_j^k \exp(-\gamma \ x_i^k - x_j^k\)$ the training RBF kernel matrix.
$\bar{Q}(\gamma)_{ij}^k$	$:=$ $\bar{y}_i^k \bar{y}_j^k \exp(-\gamma \ \bar{x}_i^k - \bar{x}_j^k\)$ the validation RBF kernel matrix.
Decision Variables	
w, b	$:=$ Weights and bias for the primal program (Soft-Margin-SVM).
α, η	$:=$ Dual variables for the dual program (Dual-SVM).
\bar{v}, \bar{v}, u	$:=$ Multipliers in the KKT conditions of the dual SVM (29).
λ, μ, ν, ξ	$:=$ Lagrange multipliers for the MPCC (8).
Functions	
f	$:=$ Objective function.
g	$:=$ Inequality constraint function i.e. $g(z) \geq 0$.
h	$:=$ Equality constraint function i.e. $h(z) = 0$.
G, H	$:=$ Complementary constraint functions i.e. $0 \leq G(z) \perp H(z) \geq 0$.
F	$:=$ System of nonlinear equations corresponding to M-stationary points (58).
Z, θ, φ	$:=$ Specific constraint functions of model (\mathcal{M}) defined in (34).
Index Sets	
I^g, I^h	$:=$ Index sets of the active inequality and equality constraints; see Section 2.1.
I^G, I^H, I^{GH}	$:=$ Index sets of the complementary constraints; see (1).
I, \bar{I}	$:=$ Sets of indices (k, i) for fold k , example i in training and validation; see (33).
C_{12345}	$:=$ The five index cases used for the proof of theoretical properties in Section 4.3

B Further details of the dual SVM

We begin by repeating the definition of the soft margin SVM, which we will call the primal

$$\begin{aligned}
 & \underset{w \in \mathbb{R}^d, \xi \in \mathbb{R}^n, b \in \mathbb{R}}{\text{minimise}} && \frac{1}{2} w^\top w + C \frac{1}{n} \sum_{i=1}^n \xi_i \\
 & \text{subject to} && \xi_i \geq 1 - y_i (w^\top x_i + b) \quad \text{for } i = 1, \dots, n, \\
 & && \xi_i \geq 0 \quad \text{for } i = 1, \dots, n.
 \end{aligned} \tag{Soft-Margin-SVM}$$

B.1 Conversion from Primal to Dual SVM

Here is presented the conversion from (Soft-Margin-SVM) to its dual form (Dual-SVM). The first step is to write the Lagrangian function $\mathcal{L} : \mathbb{R}^{d+3n+1} \rightarrow \mathbb{R}$ of the primal program (Soft-Margin-SVM). This is the sum of the original objective function plus each constraint scaled by a corresponding Lagrange multipliers $\alpha \in \mathbb{R}^n$ and $\eta \in \mathbb{R}^n$. For the textbook theory on Lagrange functions, see Boyd and Vandenberghe [8, Section 5.1.2].

$$\mathcal{L}(w, \xi, b, \alpha, \eta) = \frac{1}{2} w^\top w + C \frac{1}{n} \sum_{i=1}^n \xi_i + \sum_{i=1}^n \alpha_i (1 - \xi_i - y_i (w^\top x_i + b)) - \sum_{i=1}^n \eta_i \xi_i \tag{46}$$

We note three properties of the primal problem (Soft-Margin-SVM). First, the objective function is convex (in particular, positive quadratic). Secondly, the constraints are each linear in w, ξ, b , and so the feasible set is an (unbounded) convex polyhedron. Thirdly the feasible set has a non-empty interior since we can pick $\xi_i = \max(0, 1 - y_i (w^\top x_i + b)) + 10^{-6}$ which strictly satisfies each inequality. Therefore the problem satisfies Slater's condition and strong duality. Any feasible point w, ξ, b is a local minimiser of (Soft-Margin-SVM) (equivalent to a global minimiser by convexity) only if there exist Lagrange multipliers α, η that satisfy the following first-order stationarity conditions

$$\nabla_w \mathcal{L}(w, \xi, b, \alpha, \eta) = w - \sum_{i=1}^n \alpha_i y_i \phi(x_i) = \mathbf{0}, \quad (47a)$$

$$\nabla_b \mathcal{L}(w, \xi, b, \alpha, \eta) = \sum_{i=1}^n \alpha_i y_i = 0, \quad (47b)$$

$$\nabla_{\xi_i} \mathcal{L}(w, \xi, b, \alpha, \eta) = C \mathbf{1} - \alpha - \eta = \mathbf{0}. \quad (47c)$$

We now manipulate the Lagrangian (46) by substituting in the stationary conditions (47a, 47b, 47c). This will give us a function of the Lagrangian in terms of α, η assuming optimal choices for w, ξ, b . The term $\sum_{i=1}^n (C - \alpha_i - \eta_i) \xi_i$ vanishes due to (47c). The term $\sum_{i=1}^n \alpha_i y_i b$ vanishes due to (47b). And we eliminate w by substituting in (47a). The conversion of the problem from primal space variables (w, ξ and b) to dual space variables (α and η) like this allows us to apply the kernel trick $\phi(x_i) \phi(x_j) = K_\gamma(x_i, x_j)$. Below, we work through the manipulation line by line.

$$\begin{aligned} & \mathcal{L}^*(\alpha, \eta) \\ &= \frac{1}{2} w^\top w + C \frac{1}{n} \sum_{i=1}^n \xi_i + \sum_{i=1}^n \alpha_i (1 - \xi_i - y_i (w \phi(x_i) + b)) - \sum_{i=1}^n \eta_i \xi_i \\ &= \frac{1}{2} w^\top w + \sum_{i=1}^n (C - \alpha_i - \eta_i) \xi_i + \sum_{i=1}^n \alpha_i - \sum_{i=1}^n \alpha_i y_i w \phi(x_i) - \sum_{i=1}^n \alpha_i y_i b \\ &= \frac{1}{2} w^\top w + 0 + \sum_{i=1}^n \alpha_i - \sum_{i=1}^n \alpha_i y_i w \phi(x_i) - 0 \\ &= \frac{1}{2} \left(\sum_{i=1}^n \alpha_i y_i \phi(x_i) \right)^\top \left(\sum_{j=1}^m \alpha_j y_j \phi(x_j) \right) + \sum_{i=1}^n \alpha_i - \sum_{i=1}^n \alpha_i y_i \left(\sum_{j=1}^m \alpha_j y_j \phi(x_j) \right) \phi(x_i) \\ &= \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^m \alpha_i \alpha_j y_i y_j \phi(x_i) \phi(x_j) + \sum_{i=1}^n \alpha_i - \sum_{i=1}^n \sum_{j=1}^m \alpha_i \alpha_j y_i y_j \phi(x_i) \phi(x_j) \\ &= -\frac{1}{2} \sum_{i=1}^n \sum_{j=1}^m \alpha_i \alpha_j y_i y_j K_\gamma(x_i, x_j) + \sum_{i=1}^n \alpha_i \end{aligned}$$

By considering the infimum of the $\{\mathcal{L}(w, \xi, b, \alpha, \eta)\}$ (46) We now write the dual of (Soft-Margin-SVM). For the textbook duality theory of convex programs, see Boyd and Vandenberghe [8, Section 5.2].

$$\begin{aligned} & \underset{\alpha \in \mathbb{R}^n, \eta \in \mathbb{R}^n}{\text{maximise}} \quad \inf_{w, \xi \in \mathbb{R}^n, b \in \mathbb{R}} \left\{ \mathcal{L}(w, \xi, b, \alpha, \eta) \right\} \\ & \text{subject to} \quad 0 \leq \alpha_i \quad \text{for } i = 1, \dots, n, \\ & \quad \quad \quad 0 \leq \eta_i \quad \text{for } i = 1, \dots, n. \end{aligned}$$

The above is impractical for us. We replace the infimum of the Lagrangian ($\inf \{\mathcal{L}(w, \xi, b, \alpha, \eta)\}$) with the expression $\mathcal{L}^*(\alpha, \eta)$ derived from the stationary conditions above. Next, we multiply by -1 to convert from maximisation to a minimisation problem. Finally, the constraints $\alpha \geq 0, \eta \geq 0$ can be rearranged to $0 \leq \alpha \leq C$ by substituting out $\eta_i = C - \alpha_i$ according to (47c). Thus, the dual problem can be written as (Dual-SVM), which we will work with throughout the paper.

$$\begin{aligned} & \underset{\alpha \in \mathbb{R}^n}{\text{minimise}} \quad \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j y_i y_j K_\gamma(x_i, x_j) - \sum_{i=1}^n \alpha_i \\ & \text{subject to} \quad 0 \leq \alpha_i \leq C \quad \text{for } i = 1, \dots, n, \\ & \quad \quad \quad \sum_{i=1}^n \alpha_i y_i = 0. \end{aligned} \quad (\text{Dual-SVM})$$

The Karush–Kuhn–Tucker (KKT) conditions for the (Dual-SVM) are

$$\begin{aligned}\nabla_{\alpha_i} \mathcal{L}(\alpha, \underline{v}, \bar{v}, u) &= \sum_{j=1}^m \alpha_j y_i y_j K_\gamma(x_i, x_j) - 1 - \underline{v}_i + \bar{v}_i + u y_i = 0 \\ 0 &\leq \alpha_i \perp \underline{v}_i \geq 0 && \text{for } i = 1, \dots, n, \\ 0 &\leq (C - \alpha_i) \perp \bar{v}_i \geq 0 && \text{for } i = 1, \dots, n, \\ \sum_{i=1}^n \alpha_i y_i &= 0,\end{aligned}$$

where

$$\mathcal{L}(\alpha, \underline{v}, \bar{v}, u) := \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j y_i y_j K_\gamma(x_i, x_j) - \sum_{i=1}^n \alpha_i - \sum_{i=1}^n \underline{v}_i \alpha_i + \sum_{i=1}^n \bar{v}_i (\alpha_i - C) + u \sum_{i=1}^n \alpha_i y_i.$$

B.2 Reconstruction of primal variables from dual solutions

The weights w and bias b are primal variables. In order to use the kernel trick, we will be solving the dual optimisation problem (Dual-SVM) in terms of dual variables α . This makes it less obvious what our final classifier should predict as we cannot explicitly calculate $x \mapsto \text{sign}(w^T \varphi(x) + b)$. From the stationary conditions (47a), we have

$$w = \sum_{i=1}^n \alpha_i y_i \phi(x_i). \quad (48)$$

Using this formula (48), we can substitute out w from the classifier. Our classifier should predict

$$x \mapsto \text{sign} \left(\sum_{i=1}^n \alpha_i y_i \varphi(x_i) \varphi(x) + b \right).$$

We must work harder to find an expression for the bias b . Below, we write the complementary slackness conditions for (Soft-Margin-SVM)

$$(1 - \xi_i - y_i(w^T \phi(x_i) + b)) \alpha_i = 0, \quad (49)$$

$$\xi_i \eta_i = 0. \quad (50)$$

Suppose $\alpha_i > 0$. Then by complementary slackness (49) we must have

$$(1 - \xi_i - y_i(w^T \phi(x_i) + b)) = 0 \quad \text{for } i \in \{1, \dots, n\} : \alpha_i > 0. \quad (51)$$

Suppose $\eta_i > 0$. Notice, this is equivalent to $\alpha_i < C$ by the stationary condition (47c). Again, by complementary slackness (50) we must have

$$\xi_i = 0 \quad \text{for } i \in \{1, \dots, n\} : \alpha_i < C. \quad (52)$$

Putting these two (51, 52) together we get

$$1 - y_i (w^T \phi(x_i) + b) = 0 \quad \text{for } i \in \{1, \dots, n\} : 0 < \alpha_i < C.$$

This gives us an important insight. If α_i takes a value between 0 and C then we can conclude $y_i(w^T \phi(x_i) + b) = 1$. Geometrically, this means that example i sits exactly on the edge of the margin. We call these “support vectors”, and they are our classifier’s most crucial data points. Now recall that $y_i \in \{-1, +1\}$ so we can divide through by $\frac{1}{y_i} \equiv y_i$. This leads to

$$b = y_i - w^T \phi(x_i) \quad \text{for } i \in \{1, \dots, n\} : 0 < \alpha_i < C.$$

We can now substitute out w use (47a) giving:

$$b = y_i - \sum_{j=1}^m \alpha_j y_j \phi(x_j)^T \phi(x_i) \quad \text{for } i \in \{1, \dots, n\} : 0 < \alpha_i < C. \quad (53)$$

Theoretically, we can reconstruct the bias according to (53) from any support vector i on the margin’s edge - identified by $0 < \alpha_i < C$. However, most libraries make the bias reconstruction calculation (53) for each such support vector and then take the average. With this in mind, let $\Omega(\alpha_i, C)$ be the indicator function that equals 1 if $0 < \alpha_i < C$ and 0 otherwise. Whenever we write the bias as a function of dual variables $b(\alpha, C)$ we mean the following reconstruction

$$b = \frac{1}{\sum_{i=1}^n \Omega(\alpha_i, C)} \sum_{i=1}^n \Omega(\alpha_i, C) \left(y_i - \sum_{j=1}^n \alpha_j y_j K(x_i, x_j) \right). \quad (54)$$




Fig. 14: A graph of the function $\Omega(\alpha_i, C)$

The algorithms we propose require a further adjustment. Firstly, because we want to avoid the combinatorial nature of a 0,1 indicator function. Secondly, it is hard to distinguish between α_i very close to zero and α_j that should be zero but carries some small numerical error. We, therefore, use an alternative indicator function as proposed by Anthony Dunn [14].

$$\Omega(\alpha_i, C) = \begin{cases} 1 - \left(\frac{2\alpha_i}{C} - 1\right)^2 + \tau & \text{if } 0 \leq \alpha_i \leq C, \\ 0 & \text{otherwise.} \end{cases} \quad (55)$$

Similar to the original indicator, this takes the value zero for $\alpha_i \notin (0, C)$ and a positive value for $\alpha_i \in (0, C)$. Importantly, it assigns a much smaller value to α_i close to 0 or C . With this new indicator, the reconstruction (54) still averages across the examples where $0 < \alpha_i < C$ and so is still theoretically correct but far more numerically stable. The derivatives of this function are detailed in section Appendix B.3.

B.3 Derivatives of the bias

Here we provide explicit expressions for the first and second order derivatives of the bias that are useful when implementing code. The first and second order derivatives of the function $\Omega(\alpha_i, C) = 1 - \left(\frac{2\alpha_i}{C} - 1\right)^2 + \tau$ are:

$$\begin{aligned} \frac{\partial}{\partial \alpha_i} \Omega(\alpha_i, C) &= \frac{4C - 8\alpha_i}{C^2}; \\ \frac{\partial^2}{\partial \alpha_i^2} \Omega(\alpha_i, C) &= \frac{-8}{C^2}; \\ \frac{\partial}{\partial C} \Omega(\alpha_i, C) &= \frac{8\alpha_i^2 - 4\alpha_i C}{C^3}; \\ \frac{\partial^2}{\partial C^2} \Omega(\alpha_i, C) &= \frac{8\alpha_i C - 24\alpha_i^2}{C^4}; \\ \frac{\partial^2}{\partial \alpha_i \partial C} \Omega(\alpha_i, C) &= \frac{16\alpha_i - 4C}{C^3}. \end{aligned}$$

Given functions $f, g, h : \mathbb{R}^2 \rightarrow \mathbb{R}$ in two variables where $h(x, y) = \frac{g(x, y)}{f(x, y)}$ we have the following rules. The first-order quotient rule:

$$\frac{\partial}{\partial x} h = \frac{\partial}{\partial x} \frac{f}{g} = \frac{\left(\frac{\partial}{\partial x} f\right) \left(\frac{1}{g}\right) - \left(\frac{f}{g^2}\right) \left(\frac{\partial}{\partial x} g\right)}{g^2}.$$

And the second-order quotient rule:

$$\begin{aligned} \frac{\partial^2}{\partial x \partial y} h &= \frac{\partial^2}{\partial x \partial y} \frac{f}{g} = \frac{\left(\frac{\partial^2}{\partial x \partial y} f\right) \left(\frac{1}{g}\right) - \left(\frac{\partial^2}{\partial x \partial y} g\right) \left(\frac{f}{g^2}\right) - \left(\frac{\partial}{\partial x} g\right) \left(\frac{\partial}{\partial y} \frac{f}{g^2}\right) - \left(\frac{\partial}{\partial y} g\right) \left(\frac{\partial}{\partial x} \frac{f}{g^2}\right)}{g^2}; \\ \frac{\partial^2}{\partial x^2} h &= \frac{\partial^2}{\partial x^2} \frac{f}{g} = \frac{\left(\frac{\partial^2}{\partial x^2} f\right) \left(\frac{1}{g}\right) - h \left(\frac{\partial^2}{\partial x^2} g\right) - 2 \left(\frac{\partial}{\partial x} g\right) \left(\frac{\partial}{\partial x} \frac{f}{g^2}\right)}{g}. \end{aligned}$$

B.3.1 Derivative of the bias with respect to α

We wish to take the derivative of $b = \frac{f(\alpha)}{\sum_i \Omega(\alpha_i)}$ wrt to α . Start by just considering the derivative of the numerator:

$$\begin{aligned} f(\alpha) &= \sum_i \Omega(\alpha_i) \left(y_i - \sum_j \alpha_j y_j \exp(-\gamma \|x_i - x_j\|^2) \right) \\ &= \Omega(\alpha_k) \left(y_k - \sum_j \alpha_j y_j \exp(-\gamma \|x_k - x_j\|^2) \right) \\ &\quad + \sum_{i \neq k} \Omega(\alpha_i) (-\alpha_k y_k \exp(-\gamma \|x_i - x_k\|^2)) \\ &\quad + (\text{Terms that do not include } \alpha_k). \end{aligned}$$

Take the first order derivative of the numerator:

$$\begin{aligned} \frac{\partial}{\partial \alpha} f(\alpha)_k &= \frac{\partial}{\partial \alpha_k} \Omega(\alpha_k) \left(y_k - \sum_j \alpha_j y_j \exp(-\gamma \|x_k - x_j\|^2) \right) \\ &\quad - \sum_i \Omega(\alpha_i) y_k \exp(-\gamma \|x_i - x_k\|^2). \end{aligned}$$

And the second-order derivative of the numerator:

$$\begin{aligned} \frac{\partial^2}{\partial \alpha^2} f(\alpha)_{kl} &= \frac{\partial^2}{\partial \alpha_k \alpha_l} \Omega(\alpha_k) \left(y_k - \sum_j \alpha_j y_j \exp(-\gamma \|x_k - x_j\|^2) \right) \\ &\quad - \frac{\partial}{\partial \alpha_k} \Omega(\alpha_k) (y_l \exp(-\gamma \|x_k - x_l\|^2)) \\ &\quad - \frac{\partial}{\partial \alpha_l} \Omega(\alpha_l) (y_k \exp(-\gamma \|x_l - x_k\|^2)). \end{aligned}$$

Use the quotient rule $\frac{d}{dx} \frac{f}{g} = \frac{f'g - fg'}{g^2}$ to calculate the first order derivative of the bias wrt α :

$$\frac{\partial}{\partial \alpha_k} b(\alpha, \gamma, C) = \frac{\left(\frac{\partial}{\partial \alpha} f(\alpha)_k \right) \left(\sum_i \Omega(\alpha_i) \right) - f(\alpha) \left(\frac{\partial}{\partial \alpha_k} \Omega(\alpha_k) \right)}{\left(\sum_i \Omega(\alpha_i) \right)^2}.$$

And then the second-order quotient rule:

$$\frac{\partial^2}{\partial^2 \alpha_k \alpha_l} b(\alpha, \gamma, C) = \frac{\left(\frac{\partial^2}{\partial \alpha^2} f(\alpha)_{kl} \right) - \left(\frac{\partial^2}{\partial \alpha_k \alpha_l} \Omega(\alpha_k) \right) b(\alpha) - \left(\frac{\partial}{\partial \alpha_l} \Omega(\alpha_l) \right) \left(\frac{\partial}{\partial \alpha_k} b(\alpha_k) \right) - \left(\frac{\partial}{\partial \alpha_l} b(\alpha_l) \right) \left(\frac{\partial}{\partial \alpha_k} \Omega(\alpha_k) \right)}{\sum_i \Omega(\alpha_i)}.$$

B.3.2 Derivative of the bias with respect to C

We wish to take the derivative of b wrt to α . Start by just considering the numerator:

$$\begin{aligned} f(C) &= \sum_{i=1}^n \Omega(\alpha_i, C) \left(y_i - \sum_{j=1}^n \alpha_j y_j \exp(-\gamma \|x_i - x_j\|^2) \right); \\ \frac{\partial}{\partial C} f(C) &= \sum_i \frac{\partial}{\partial C} \Omega(\alpha_i, C) \left(y_i - \sum_j \alpha_j y_j \exp(-\gamma \|x_i - x_j\|^2) \right); \\ \frac{\partial^2}{\partial C^2} f(C) &= \sum_i \frac{\partial^2}{\partial C^2} \Omega(\alpha_i, C) \left(y_i - \sum_j \alpha_j y_j \exp(-\gamma \|x_i - x_j\|^2) \right). \end{aligned}$$

Use the quotient rule $\frac{d}{dx} \frac{f}{g} = \frac{f'g - fg'}{g^2}$ to calculate the first order derivative of the bias wrt C :

$$\frac{\partial}{\partial C} b(\alpha, C, \gamma) = \frac{\left(\frac{\partial}{\partial C} f(C) \right) \left(\sum_i \Omega(\alpha_i, C) \right) - f(C) \left(\sum_i \frac{\partial}{\partial C} \Omega(\alpha_i, C) \right)}{\left(\sum_i \Omega(\alpha_i, C) \right)^2}.$$

And then the second-order quotient rule:

$$\frac{\partial^2}{\partial C^2} b(\alpha, C, \gamma) = \frac{\left(\frac{\partial^2}{\partial C^2} f(C) \right) - \left(\sum_i \frac{\partial^2}{\partial C^2} \Omega(\alpha_i, C) \right) (b(C)) - 2 \left(\sum_i \frac{\partial}{\partial C} \Omega(\alpha_i, C) \right) \left(\sum_i \frac{\partial}{\partial C} b(C) \right)}{\left(\sum_i \Omega(\alpha_i, C) \right)^2}.$$

B.3.3 Derivative of the bias with respect to gamma

We wish to take the derivative of b wrt to α . Start by just considering the numerator:

$$\begin{aligned} f(\gamma) &= \sum_{i=1}^n \Omega(\alpha_i, C) \left(y_i - \sum_{j=1}^n \alpha_j y_j \exp(-\gamma \|x_i - x_j\|^2) \right); \\ \frac{\partial}{\partial \gamma} f(\gamma) &= \sum_i \Omega(\alpha_i, C) \left(\sum_j \alpha_j y_j \|x_i - x_j\|^2 \exp(-\gamma \|x_i - x_j\|^2) \right); \\ \frac{\partial^2}{\partial \gamma^2} f(\gamma) &= \sum_i \Omega(\alpha_i, C) \left(\sum_j \alpha_j y_j (-) \|x_i - x_j\|^4 \exp(-\gamma \|x_i - x_j\|^2) \right). \end{aligned}$$

The first-order derivative of bias with respect to γ is:

$$\frac{\partial}{\partial \gamma} b(\alpha, C, \gamma) = \frac{\frac{\partial}{\partial \gamma} f(\gamma)}{\sum_i \Omega(\alpha_i, C)}.$$

The second-order derivative of bias with respect to γ is:

$$\frac{\partial^2}{\partial \gamma^2} b(\alpha, C, \gamma) = \frac{\frac{\partial^2}{\partial \gamma^2} f(\gamma)}{\sum_i \Omega(\alpha_i, C)}.$$

B.3.4 Mixed derivative of the bias

Take second order derivatives of the numerator (f) with respect to each pair of C, γ, α giving:

$$\begin{aligned} f(\alpha, C, \gamma) &= \sum_i \Omega(\alpha_i, C) \left(y_i - \sum_j \alpha_j y_j \exp(-\gamma \|x_i - x_j\|^2) \right); \\ \frac{\partial^2}{\partial C \partial \gamma} f &= \sum_i \frac{\partial}{\partial C} \Omega(\alpha_i, C) \left(\sum_j \alpha_j y_j \|x_i - x_j\|^2 \exp(-\gamma \|x_i - x_j\|^2) \right); \\ \frac{\partial^2}{\partial \alpha \partial C} f &= \frac{\partial^2}{\partial \alpha_k \partial C} \Omega(\alpha_k) \left(y_k - \sum_j \alpha_j y_j \exp(-\gamma \|x_k - x_j\|^2) \right) \\ &\quad - \sum_i \frac{\partial}{\partial C} \Omega(\alpha_i) y_k \exp(-\gamma \|x_i - x_k\|^2); \\ \frac{\partial^2}{\partial \alpha \partial \gamma} f &= \frac{\partial}{\partial \alpha_k} \Omega(\alpha_k) \left(\sum_j \alpha_j y_j \|x_k - x_j\|^2 \exp(-\gamma \|x_k - x_j\|^2) \right) \\ &\quad + \sum_i \Omega(\alpha_i) y_k \|x_k - x_j\|^2 \exp(-\gamma \|x_i - x_k\|^2). \end{aligned}$$

The derivative of bias with respect to both C and gamma is:

$$\frac{\partial^2}{\partial C \partial \gamma} b(\alpha, \gamma, C) = \frac{\left(\frac{\partial^2}{\partial C \partial \gamma} f \right) \left(\sum_i \Omega(\alpha_i, C) \right) - \left(\frac{\partial}{\partial \gamma} f \right) \left(\sum_i \frac{\partial}{\partial C} \Omega(\alpha_i, C) \right)}{\left(\sum_i \Omega(\alpha_i, C) \right)^2}.$$

Use the quotient rule $\frac{d}{dx} \frac{f}{g} = \frac{f'g - fg'}{g^2}$, we calculate the second order derivative of the bias wrt α and γ :

$$\frac{\partial^2}{\partial \alpha_k \partial \gamma} b(\alpha, \gamma, C) = \frac{\left(\frac{\partial^2}{\partial \alpha_k \partial \gamma} f \right) \left(\sum_i \Omega(\alpha_i, C) \right) - \left(\frac{\partial}{\partial \gamma} f \right) \left(\frac{d}{d \alpha_k} \Omega(\alpha_k, C) \right)}{\left(\sum_i \Omega(\alpha_i, C) \right)^2}.$$

The derivative of bias with respect to both α and C is:

$$\frac{\partial^2}{\partial \alpha_k \partial C} b(\alpha, \gamma, C) = \frac{\left(\frac{\partial^2}{\partial \alpha_k \partial C} f \right) - \left(\frac{\partial^2}{\partial \alpha_k \partial C} \Omega(\alpha_k, C) \right) (b) - \left(\frac{\partial}{\partial \alpha_k} \Omega(\alpha_k, C) \right) \left(\frac{\partial}{\partial C} b \right) - \left(\sum_i \frac{\partial}{\partial C} \Omega(\alpha_i, C) \right) \left(\frac{\partial}{\partial \alpha_k} b \right)}{\sum_i \Omega(\alpha_i, C)}.$$

C Explicit formula of the Lagrangian of (\mathcal{M}) and derivatives

Early in the paper, we dealt with a compact form of the MPCC problem (\mathcal{M}) . However, when it comes to implementing solution methods, we need accurate formulas for the Lagrangian and its first and second-order derivatives. We use the notation defined in equations (33) and (34). In this appendix, we present those explicit formulas in full, beginning with the Lagrangian

$$\begin{aligned} \mathcal{L}(z, \mu) = & \sum_{(k,i) \in \bar{I}} \zeta_i + -\mu^C C - \mu^\gamma \gamma - \sum_{(k,i) \in \bar{I}} \mu_{(k,i)}^\zeta \zeta_i^k - \sum_{(k,i) \in \bar{I}} \mu_{(k,i)}^Z Z(z)_i^k - \sum_{(k,i) \in I} \mu_{(k,i)}^\theta \theta(z)_i^k \\ & - \sum_{(k,i) \in I} \mu_{(k,i)}^{G_1} \alpha_i^k - \sum_{(k,i) \in I} \mu_{(k,i)}^{H_1} \bar{v}_i^k - \sum_{(k,i) \in I} \mu_{(k,i)}^{G_2} (C - \alpha_i^k) - \sum_{(k,i) \in I} \mu_{(k,i)}^{H_2} \bar{v}_i^k - \sum_{k=1, \dots, K} \mu_k^\varphi \varphi(z)^k; \end{aligned} \quad (56)$$

where the decision variables and Lagrange multipliers are

$$\begin{aligned} z &= [C, \gamma, \zeta, \alpha, \bar{v}, u] \in \mathbb{R}^{3(K-1)n+n+K+2}; \\ \mu &= [\mu^C, \mu^\gamma, \mu^\zeta, \mu^Z, \mu^\theta, \mu^{G_1}, \mu^{H_1}, \mu^{G_2}, \mu^{H_2}, \mu^\varphi] \in \mathbb{R}^{5(K-1)n+2n+K+2}; \end{aligned}$$

with dimensions in terms of the total number of data points n and the number of folds K .

The first order derivative with respect to all input variables is a vector to vector function $\nabla \mathcal{L} : \mathbb{R}^L \rightarrow \mathbb{R}^L$ where $L = 8(K-1)n + 3n + 2K + 4$. We define it in terms of each partial derivative.

$$\begin{aligned} \nabla_C \mathcal{L}(z, \mu) &= -\mu^C - \sum_{k=1}^K \sum_{i=1}^{\bar{n}^k} \mu_{(k,i)}^Z \bar{y}_i^k \nabla_C b - \sum_{k=1}^K \sum_{i=1}^{n^k} \mu_{(k,i)}^{G_2}; \\ \nabla_\gamma \mathcal{L}(z, \mu) &= -\mu^\gamma + \sum_{k=1}^K \sum_{i=1}^{\bar{n}^k} \mu_{(k,i)}^Z \left(\sum_{j=1}^{n^k} \alpha_j^k \left\| \bar{x}_i^k - x_j^k \right\|^2 \bar{Q}(\gamma)_{ij}^k - \bar{y}_i^k \nabla_\gamma b \right) \\ &\quad + \sum_{k=1}^K \sum_{i=1}^{n^k} \mu_{(k,i)}^\theta \left(\sum_{j=1}^m \alpha_j^k \left\| x_i^k - x_j^k \right\|^2 Q(\gamma)_{ij}^k \right); \\ \nabla_{\zeta_i^k} \mathcal{L}(z, \mu) &= 1 - \mu_{(k,i)}^\zeta - \mu_{(k,i)}^Z; \\ \nabla_{\alpha_j^k} \mathcal{L}(z, \mu) &= - \sum_{i=1}^{\bar{n}^k} \mu_{(k,i)}^Z \left(\bar{Q}(\gamma)_{ij}^k + \bar{y}_i^k \nabla_{\alpha_j^k} b \right) - \sum_{i=1}^{n^k} \mu_{(k,i)}^\theta Q(\gamma)_{ij}^k - \mu_{(k,j)}^{G_1} + \mu_{(k,j)}^{G_2}, -\mu_k^\varphi y_j^k; \\ \nabla_{\bar{v}_i^k} \mathcal{L}(z, \mu) &= +\mu_{(k,i)}^\theta - \mu_{(k,i)}^{H_1}; \\ \nabla_{\bar{v}_i^k} \mathcal{L}(z, \mu) &= -\mu_{(k,i)}^\theta - \mu_{(k,i)}^{H_2}; \\ \nabla_{u^k} \mathcal{L}(z, \mu) &= - \sum_{i=1}^{n^k} \mu_{(k,i)}^\theta y_i^k. \end{aligned}$$

The Hessian of the Lagrangian is:

$$\nabla_{zz}^2 \mathcal{L}(z, \mu) = \begin{bmatrix} \nabla_{CC}^2 \mathcal{L} & \nabla_{C\gamma}^2 \mathcal{L} & 0 & \nabla_{C\alpha}^2 \mathcal{L} & 0 & 0 & 0 \\ \nabla_{\gamma C}^2 \mathcal{L} & \nabla_{\gamma\gamma}^2 \mathcal{L} & 0 & \nabla_{\gamma\alpha}^2 \mathcal{L} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \nabla_{\alpha C}^2 \mathcal{L} & \nabla_{\alpha\gamma}^2 \mathcal{L} & 0 & \nabla_{\alpha\alpha}^2 \mathcal{L} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

The partial second-order derivatives can be written explicitly as:

$$\begin{aligned}
\nabla_{CC}^2 \mathcal{L}(z, \mu) &= - \sum_{k=1}^K \sum_{i=1}^{\bar{n}^k} \mu_{(k,i)}^Z \bar{y}_i^k \nabla_{CC}^2 b(\alpha^k, \gamma, C); \\
\nabla_{\gamma\gamma}^2 \mathcal{L}(z, \mu) &= - \sum_{k=1}^K \sum_{i=1}^{\bar{n}^k} \mu_{(k,i)}^Z \left(\sum_{j=1}^{\bar{n}^k} \alpha_j^k \left\| \bar{x}_i^k - x_j^k \right\|^4 \bar{Q}(\gamma)_{ij}^k + \bar{y}_i^k \nabla_{\gamma\gamma}^2 b \right) - \sum_{k=1}^K \sum_{i=1}^{\bar{n}^k} \mu_{(k,i)}^\theta \sum_{j=1}^m \alpha_j^k \left\| x_i^k - x_j^k \right\|^4 Q(\gamma)_{ij}^k; \\
\nabla_{\alpha_j^k \alpha_l^k}^2 \mathcal{L}(z, \mu) &= \sum_{i=1}^{\bar{n}^k} \mu_{(k,i)}^Z \bar{y}_i^k \nabla_{\alpha_j^k \alpha_l^k}^2 b(\alpha^k, \gamma, C); \\
\nabla_{C\gamma}^2 \mathcal{L}(z, \mu) &= - \sum_{k=1}^K \sum_{i=1}^{\bar{n}^k} \mu_{(k,i)}^Z \bar{y}_i^k \nabla_{C\gamma}^2 b(\alpha^k, \gamma, C); \\
\nabla_{C\alpha_l^k}^2 \mathcal{L}(z, \mu) &= - \sum_{i=1}^{\bar{n}^k} \mu_{(k,i)}^Z \bar{y}_i^k \nabla_{C\alpha_l^k}^2 b(\alpha^k, \gamma, C); \\
\nabla_{\gamma\alpha_j^k}^2 \mathcal{L}(z, \mu) &= \sum_{i=1}^{\bar{n}^k} \mu_{(k,i)}^Z \left(\left\| \bar{x}_i^k - x_j^k \right\|^2 \bar{Q}(\gamma)_{ij}^k - \bar{y}_i^k \nabla_{\alpha_j^k \gamma}^2 b \right) + \sum_{i=1}^{\bar{n}^k} \mu_{(k,i)}^\theta \left\| x_i^k - x_j^k \right\|^2 Q(\gamma)_{ij}^k.
\end{aligned}$$

D Semismooth Newton method

During our experiments, we compared to the Semi-smooth Newton method. This method writes the M-stationary conditions of problem (MPCC) as a system of semismooth equations, which can be solved with a Newton-style method. This framework is proposed in [41]. It requires, in the first instance, the use of an NCP function that, given inputs a, b returns zero if and only if $a, b \geq 0$ and $ab = 0$. Two examples of such functions are the minimum and Fischer-Burmeister functions defined by

$$\min(a, b) := \begin{cases} a & \text{if } a \leq b, \\ b & \text{otherwise;} \end{cases} \quad \text{fb}(a, b) := \sqrt{a^2 + b^2} - a - b.$$

The second step of building the framework consists of introducing the M-stationarity function $\text{msf} : \mathbb{R}^4 \rightarrow \mathbb{R}^2$, which has the property that its zero level set $\{(G_i, H_i, \nu_i, \xi_i) : \text{msf}(G_i, H_i, \nu_i, \xi_i) = 0\}$ is equal to the set of points that satisfy the M-stationary conditions; i.e.,

$$\text{msf}(G_i, H_i, \nu_i, \xi_i) = 0 \iff \begin{cases} 0 \leq G_i \perp H_i \leq 0, \\ G_i \nu_i = 0, \\ H_i \xi_i = 0, \\ (\nu_i, \xi_i \geq 0 \text{ or } \nu_i \xi_i = 0). \end{cases} \quad (57)$$

For an example of such a function, see [41, page 1470, Section 3.2]. The full system of equations to be solved to compute M-stationary points for problem (MPCC) becomes

$$F(z, \lambda, \mu, \nu, \xi) := \begin{bmatrix} \nabla_z \mathcal{L}(z, \lambda, \mu, \nu, \xi) \\ [\text{fb}(g_i(z), \lambda_i)]_{i=1, \dots, p} \\ [h_i(z)]_{i=1, \dots, q} \\ [\text{msf}(G_i(z), H_i(z), \nu_i, \xi_i)]_{i=1, \dots, r} \end{bmatrix} = 0 \quad (58)$$

with the function F being semismooth (see [41] and references therein for the definition), while \mathcal{L} is the Lagrangian function as defined in (8). Note that here, the fb function may be replaced by min or any other NCP function.

It can easily be shown that $(z^*, \lambda^*, \mu^*, \nu^*, \xi^*)$ solves equation (58) if and only if z^* is an M-stationary point of the (MPCC) with Lagrange multipliers $(\lambda^*, \mu^*, \nu^*, \xi^*)$. To solve the system, we use a semismooth Newton Method with an Amijo line search and the Newton derivative to address the involved non-smoothness, according to [41, Algorithm 5.1].




Fig. 15: The heat map shows the hyperparameter space for the dataset Ionosphere. Green regions correspond to (C, γ) combinations, which allow the SVM to achieve high validation accuracy. The heat map is repeated four times for the grid, random, Bayesian and pattern search methods. On each, blue crosses are plotted on all the (C, γ) combinations evaluated by that search method.

The system of equations, when applied to our SVM hyperparameter tuning model (\mathcal{M}) , is

$$F_{\mathcal{M}}(z, \mu) = \begin{bmatrix} \nabla_z \mathcal{L}(z, \mu) \\ \text{fb}(C, \mu^C) \\ \text{fb}(\gamma, \mu^\gamma) \\ \left[\text{fb} \left(c_i^k, \mu_{(k,i)}^{\zeta} \right) \right]_{(k,i) \in \bar{I}} \\ \left[\text{fb} \left(Z(z)_i^k, \mu_{(k,i)}^Z \right) \right]_{(k,i) \in \bar{I}} \\ \left[\theta(z)_i^k \right]_{(k,i) \in I} \\ \left[\varphi(z)_i^k \right]_{(k,i) \in I} \\ \left[\text{msf} \left(\alpha_i^k, \underline{v}_i^k, \mu_{(k,i)}^{G_1}, \mu_{(k,i)}^{H_1} \right) \right]_{(k,i) \in I} \\ \left[\text{msf} \left(C - \alpha_i^k, \bar{v}_i^k, \mu_{(k,i)}^{G_2}, \mu_{(k,i)}^{H_2} \right) \right]_{(k,i) \in I} \end{bmatrix} \quad (59)$$

where $\text{fb} : \mathbb{R}^2 \rightarrow \mathbb{R}$ is the Fischer Burmeister function and $\text{msf} : \mathbb{R}^4 \rightarrow \mathbb{R}^2$ is the M-stationarity function, the derivatives of the Lagrangian are written explicitly in Appendix C and the functions Z, θ, φ are written in (34a,b,c).

E Summary of derivative-free methods

Grid search [37, page 79] requires the data scientist to choose a set of plausible values for each hyperparameter. Given a combination of hyperparameters, the algorithm trains, evaluates and then discards K fresh SVMs according to K -fold cross-validation. This is repeated independently for every combination in the Cartesian product that constitutes the grid. In the end, the hyperparameter combination that results in the best mean validation accuracy across the K classifiers is returned.

As for random search [37, page 81][65], it works in much the same way as grid search. But, rather than a systematic search over a grid, at each iteration, the hyperparameters are drawn randomly and independently from predefined distributions. This has the advantage of being able to search in-between the grid points of a grid search.

Bayesian optimisation [34, 49, 77, 88] is a more sophisticated probabilistic method. It maintains a model approximation of the function \tilde{f} that maps hyperparameter selections to their corresponding mean validation accuracy and an acquisition function representing the expected improvement in \tilde{f} . At each iteration, the next choice of hyperparameters is determined by minimising the acquisition function - then the observed validation accuracy is used to update the approximation of \tilde{f} . Top-level parameters controlling the balance between exploration and exploitation and the smoothness assumptions of \tilde{f} must themselves be carefully chosen. Bayesian optimisation has much more computationally expensive iterations, but is often able to search points close to the global minimum.

Pattern search [73, Algorithm 9.2] differs from the previous three DFO methods in that it is a local rather than a global search. The algorithm maintains a current best hyperparameter selection (C, γ) and step size h . At each iteration, it evaluates new hyperparameters $\{(C \pm h, \gamma)\} \cup \{(C, \gamma \pm h)\}$ one step size away along one of the coordinate axes. If it finds a better objective value, it updates the current best selection and continues. Otherwise, it reduces its step size. Once the step size is sufficiently small, it will terminate—hopefully at a local minimum. This method can be made more sophisticated with pattern moves, momentum or a line search. Though it is less used in machine learning, it provides an interesting comparison to our bilevel approach, which is also a local method. All four methods are visualised in Figure 15.

F Further details of the classification datasets

Below, we list the datasets we use with a short description of any adjustments we make and a citation to the original publication. Each dataset is annotated with a tuple (n, d) where n is the number of observations and d is the dimension of the feature vectors.

1. **Banknote** (1372, 4) The four features of this data set are digitally extracted from images of banknotes. Each banknote is then labelled as either genuine or forged. It is mostly balanced with 762 genuine and 610 forged examples. The banknote identification dataset was collected by Volker Lohweg [62].
2. **BUPA Liver Disorders** (345, 6) Five features come from various blood tests relevant to liver disorders. A sixth feature comes from the patient’s response to the question, “How many half-pint equivalents of alcoholic beverages do you drink per day?”. If the patient is found to have any form of liver disorder, the patient is given a positive label. The data was published by Richard S. Forsyth and BUPA Medical Research Ltd. [29]
3. **Circles** (n, d) A synthetic dataset where half the examples are scattered uniformly over the surface of the d -dimensional unit ball and given a positive label. The other half are scattered on the surface of a smaller ball within the interior of the unit ball and given a negative label. Gaussian noise with mean zero and standard deviation σ^2 is added to each point. For our experiments in Section 5, we parametrise with $n = 54, d = 2, \sigma^2 = 0.2$. See figure 6, for an illustration.
4. **Cleveland** (297, 14) The Cleveland Heart Disease data set was introduced by Detrano et al. [20]. Its features represent the physical attributes of a patient, such as age, sex, blood pressure, cholesterol, and heart rate, as recorded by their doctors. The labels represent a positive or negative diagnosis of heart disease. Originally, the dataset had 303 examples, 76 attributes and five classes. We adjust the dataset in the following manner. First, we select only 14 features that are commonly used by other authors. Then, we remove examples with missing values. Finally, we replace the different classes of heart disease with a positive label and the absence of heart disease with a negative label so that we can train a binary classifier.
5. **Glass** (214, 9) Each sample is described in terms of its oxide content (sodium, magnesium, aluminium, etc.) and refractive index. The samples are labelled by their glass type, which can be useful to determine for criminal investigations. This data was collected by B. German [36] and is openly available on the UCI Machine Learning Repository. Since it is a multi-class dataset, we make ‘building windows float processed’ (70 examples) into the positive class and combine the vehicle, containers, tableware and lamp glass types (68 examples) into the negative class.
6. **Hastie** (100, 10) This is the simulated data set from equation 10.2 of “The Elements of Statistical Learning” [43]. It has ten features that are independently and identically sampled according to a unit-independent Gaussian distribution. The labels are detriment as $y_i = 1$ if $\|x_i\|^2 > 9.34$ otherwise $y_i = -1$.
7. **Ionosphere** (351, 34) This data was collected from radar antennas in the ionosphere in Goose Bay, Labrador [85] and used to classify structure in the ionosphere. It is somewhat unbalanced, with 225 positive labels and 126 negative labels.
8. **Iris** (100, 4) The famous Iris flower data set was introduced by Fisher and Fisher Anderson [1, 26]. Its four features represent the lengths and widths of the sepals and petals. The labels represent the species as either Setosa, Virginica or Versicolor. We discard all examples of the Versicolor class, so we end up with balanced binary labels of 50 Setosa and 50 Virginica.
9. **Linear** (n, d) A simple synthetic dataset where the data points are scattered randomly and uniformly over the space $[0, 1]^d$. A random hyperplane is chosen. The labels are then assigned as +1 if they lie above the hyperplane and -1 if they lie below. This is the only linearly separable dataset. In our experiments we choose $n = 100, d = 10$.
10. **MNIST** (600, 784) Low resolution, 28x28 pixel, images of handwritten digits 0, 1, 2, ..., 9. The dataset is massive with 60,000 examples. We subsample 300 examples of the digit 0 and 300 examples of the digit 1 in order to convert to a binary classification problem.[53]
11. **Mushrooms** (300, 117) This dataset details observations of gilled mushrooms in the Agaricus and Lepiota families [53]. All of the features are categorical, for example, the feature ‘cap shape’ could take the value bell, conical, convex, flat, knobbed or sunken. We add indicator variables for every category, which makes 117 features. Each example is labelled as either poisonous or edible.
12. **Moons** (54, 2) A synthetic two-dimensional data set where points of the different classes are scattered in interweaving half-circles around each other, which can only be separated by a very nonlinear ‘S’-shaped boundary.
13. **Palmer** (344, 5) Penguins had their features, such as the culmen length, culmen depth, flipper length, body mass and sex, recorded and were labelled by their species. We perform a few modifications: drop the rows with missing values; replace the sex with a binary 1,0 indicator; and replace the species with a positive label if Gentoo and a negative label otherwise. The data was collected by Palmer Station Antarctica LTER and K. Gorman [45] and is openly available on GitHub.

-
14. **Pima** (768, 8) The features of the Pima Indians Diabetes Database are measurements such as BMI, insulin level, and age from female patients of Pima Indian heritage. The labels represent the diagnosis of diabetes. It is quite unbalanced with 500 negative examples but only 268 positive. The dataset was provided by the National Institute of Diabetes and Digestive and Kidney Diseases and introduced to the literature by Smith et al. [86].
 15. **Sonar** (208, 60) This data was collected to improve mine detection in the ocean. The many features are those collected from the sonar signal. The labels are positive if the sonar signals bounced off a metal cylinder and negative if they bounced off rocks. Terry Sejnowski made the data available, and it was first published by Gorman and Sejnowski [39].
 16. **Swiss Roll** (n, d) Swiss Roll is a common synthetic dataset that can be tricky for an SVM as the two classes are rolled inside one another. It is presented in the book Machine Learning: An Algorithmic Perspective [66, Chapter 6, figure 6.13]. We use the Scikit-learn function `make_swiss_roll` to generate the data with Gaussian noise with standard deviation 0.1.
 17. **Wisconsin** (569, 30) This first breast cancer dataset was presented in the work of Street et al. [93]. Its thirty features are detailed measurements of the breast and tumour, such as radius, concavity, and compactness. Each example is then labelled as malignant or benign.
 18. **XOR** ($n, 2$) A two-dimensional synthetic data set where points are scattered uniformly at random across the plane $[-1, 1]^2$. They are then labelled as positive if the product of their coordinates is positive and negative otherwise. This leads to the top right and bottom left quadrants being the opposite class to the top left and bottom right quadrants. It can be a particularly problematic dataset for classifiers.