



An Efficient Active-Set method With Applications to Sparse Approximations and Risk Minimization

Spyridon Pougkakiotis¹ · Jacek Gondzio² · Dionysis Kalogierias³

Received: 7 May 2024 / Revised: 27 November 2024 / Accepted: 29 April 2025 /
Published online: 3 June 2025
© The Author(s) 2025

Abstract

In this paper we present an efficient active-set method for the solution of convex quadratic programming problems with general piecewise-linear terms in the objective, with applications to sparse approximations and risk-minimization. The algorithm is derived by combining a proximal method of multipliers (PMM) with a standard semismooth Newton method (SSN), and is shown to be globally convergent under minimal assumptions. Further local linear (and potentially superlinear) convergence is shown under standard additional conditions. The major computational bottleneck of the proposed approach arises from the solution of the associated SSN linear systems. These are solved using a Krylov-subspace method, accelerated by certain novel general-purpose preconditioners which are shown to be optimal with respect to the proximal penalty parameters. The preconditioners are easy to store and invert, since they exploit the structure of the nonsmooth terms appearing in the problem's objective to significantly reduce their memory requirements. We showcase the efficiency, robustness, and scalability of the proposed solver on a variety of problems arising in risk-averse portfolio selection, L^1 -regularized partial differential equation constrained optimization, quantile regression, and binary classification via linear support vector machines. We provide computational evidence, on real-world datasets, to demonstrate the ability of the solver to efficiently and competitively handle a diverse set of medium- and large-scale optimization instances.

Keywords Active-set methods · Proximal methods of multipliers · Semismooth Newton methods · Nonsmooth optimization · Convex programming · Krylov subspace methods · Preconditioning

Mathematics Subject Classification 65K05 · 90C06 · 90C25 · 49M15 · 49J52 · 65F10

✉ Spyridon Pougkakiotis
spyridon.pougkakiotis@kcl.ac.uk

Jacek Gondzio
j.gondzio@ed.ac.uk

Dionysis Kalogierias
dionysis.kalogierias@yale.edu

¹ Department of Mathematics, King's College London, London, UK

² School of Mathematics and Maxwell Institute for Mathematical Sciences, University of Edinburgh, Edinburgh, UK

³ Department of Electrical Engineering, Yale University, New Haven, USA

1 Introduction

1.1 Problem Formulation

In this paper we consider convex optimization problems of the following form:

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & \left\{ c^\top x + \frac{1}{2} x^\top Q x + \sum_{i=1}^l ((Cx + d)_i)_+ + \delta_{\mathcal{K}}(x) \right\}, \\ \text{s.t.} \quad & Ax = b, \end{aligned} \quad (\text{P})$$

where $c \in \mathbb{R}^n$, $Q \in \mathbb{R}^{n \times n}$ is a positive semi-definite matrix, $C \in \mathbb{R}^{l \times n}$, $d \in \mathbb{R}^l$, $A \in \mathbb{R}^{m \times n}$ is a linear constraint matrix with $b \in \mathbb{R}^m$ a given right-hand side, and $\mathcal{K}$ is a closed convex set $\mathcal{K} \triangleq \{x \in \mathbb{R}^n \mid x \in [a_l, a_u]\}$, with $a_l, a_u \in \mathbb{R}^n$, such that $(a_l)_i \in \mathbb{R} \cup \{-\infty\}$, $(a_u)_i \in \mathbb{R} \cup \{+\infty\}$, for all $i = 1, \dots, n$. Additionally, $(\cdot)_+ \equiv \max\{\cdot, 0\}$, while $\delta_{\mathcal{K}}(\cdot)$ is an indicator function for the set $\mathcal{K}$, that is, $\delta_{\mathcal{K}}(x) = 0$ if $x \in \mathcal{K}$, and $\delta_{\mathcal{K}}(x) = \infty$ otherwise.

For simplicity, we introduce some additional notation and equivalently write (P) as

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & \{f(x) + h(Cx + d) + \delta_{\mathcal{K}}(x)\}, \\ \text{s.t.} \quad & Ax = b. \end{aligned}$$

where we define $f(x) \triangleq c^\top x + \frac{1}{2} x^\top Q x$, and $h(w) \triangleq \sum_{i=1}^l (w_i)_+$.

Remark 1 Let us notice that the model in (P) can readily accommodate terms of the form of $\|Cx + d\|_1$ where $C \in \mathbb{R}^{l \times n}$, and $d \in \mathbb{R}^l$. Indeed, letting $c = c - C^\top \mathbb{1}_l$ and adding the constant term $-\mathbb{1}_l^\top d$ in the objective of (P), we notice that

$$\|Cx + d\|_1 \equiv -\mathbb{1}_l^\top (Cx + d) + \sum_{i=1}^l (2(Cx + d)_i)_+,$$

where $\mathbb{1}_l \triangleq (1, \dots, 1)^\top \in \mathbb{R}^l$. Similarly, any piecewise-linear term of the form

$$\sum_{i=1}^l \max\{(C_1 x + d_1)_i, (C_2 x + d_2)_i\},$$

where $C_1, C_2 \in \mathbb{R}^{l \times n}$ and $d_1, d_2 \in \mathbb{R}^l$, can also be readily modeled. Specifically, setting $c = c + C_2^\top \mathbb{1}_l$ and adding the constant term $d_2^\top \mathbb{1}_l$ in the objective yields

$$\mathbb{1}_l^\top (C_2 x + d_2) + \sum_{i=1}^l ((C_1 x + d_1 - C_2 x - d_2)_i)_+ \equiv \sum_{i=1}^l \max\{(C_1 x + d_1)_i, (C_2 x + d_2)_i\}.$$

Finally, it is important to note that model (P) allows for multiple piecewise-linear terms of the form $\max\{Cx + d, 0\}$, $\|Cx + d\|_1$ or $\max\{C_1 x + d_1, C_2 x + d_2\}$, since we can always adjust l to account for more than one term. Hence, one can observe that (P) is very general and can be used to model a plethora of very important problems that arise in practice.

Let $F \triangleq [A^\top - C^\top]^\top$ and $\hat{b} \triangleq [b^\top d^\top]^\top$. Using Fenchel duality, we can verify (see Appendix A.1) that the dual of (P) is

$$\begin{aligned} \max_{x \in \mathbb{R}^n, y \in \mathbb{R}^{m+l}, z \in \mathbb{R}^n} \quad & \hat{b}^\top y - \frac{1}{2} x^\top Q x - \delta_{\mathcal{K}}^*(z) - h^*(y_{m+1:m+l}), \\ \text{s.t.} \quad & c + Qx - F^\top y + z = 0_n. \end{aligned} \quad (\text{D})$$

Throughout the paper we make use of the following blanket assumption.

Assumption 1 Problems (P) and (D) are both feasible and A has full row rank.

If $C = 0_{0,n}$, using [1, Proposition 2.3.4] we note that Assumption 1 implies that there exists a primal-dual triple (x^*, y^*, z^*) solving (P)–(D). If the primal-dual pair (P)–(D) is feasible, it must remain feasible for any $C \in \mathbb{R}^{l \times n}$, since in this case (P) can be written as a convex quadratic problem by appending appropriate (necessarily feasible) linear inequality constraints. Thus, Assumption 1 suffices to guarantee that the solution set of (P)–(D) is nonempty.

In light of the discussion in Remark 1, we observe that **problem (P) models a plethora of very important problems arising in several application domains** spanning, among others, operational research, machine learning, data science, and engineering. More specifically, (P) can model linear and convex quadratic programming instances, regularized (fused) lasso instances (often arising in signal or image processing and machine learning, e.g., see [2–5]), or portfolio allocation problems [6] arising in finance. Furthermore, various optimal control problems can be tackled in the form of (P), such as those arising from L^1 -regularized partial differential equation (PDE) constrained optimization, assuming that a *discretize-then-optimize* strategy is adopted (e.g., see [7]). Additionally, various risk-minimization problems with linear random cost functions can be modeled by (P) (e.g., see [8–10]). Nonetheless, even risk-minimization problems with nonlinear random cost functions, which are typically solved via Gauss-Newton schemes (e.g., see [11]), often require the solution of sub-problems of the form of (P). Finally, continuous relaxations of integer programming problems with applications in operational research (e.g., [12]) often take the form of (P). Given the multitude of problems requiring easy access to (usually accurate) solutions of (P), the derivation of efficient, robust, and scalable solution methods is of paramount importance and has naturally attracted a lot of attention in the literature.

1.2 Existing Solution Methods

There is a large variety of first-order methods capable of finding an approximate solution to (P). For example, one could employ proximal (sub-)gradient (e.g., see [13]) or splitting schemes (e.g., see [14]). While such solution methods are very general, easy to implement, and require very little memory, they are usually able to find only highly approximate solutions, not exceeding 2- or 3-digits of accuracy. If a more accurate solution is needed, then one has to resort to an approach that utilizes second-order information.

There are three major classes of second-order methods for problems of the form of (P). Those include globalized (smooth, semismooth, quasi or proximal) Newton methods (e.g., see [15–17]), variants of the proximal point method (e.g., see [18–22]), or interior point methods (IPMs) applied to a quadratic programming (QP) reformulation of (P) (e.g., see [3, 23–25]).

Most globalized Newton-like approaches or proximal point variants studied in the literature are developed for composite programming problems in which either $C = 0_{0,n}$ (e.g., see [18, 22, 26–28]) or $\mathcal{K} = \mathbb{R}^n$ and the general $\max\{\cdot, 0\}$ terms are substituted by separable ℓ_1 terms (e.g., see [19, 29, 30]). Nonetheless, there have been developed certain globalized Newton-like schemes, specialized to the case of L^1 -regularized PDE-constrained optimization (see [31, 32]), in which both ℓ_1 term as well as box constraints are explicitly handled. We should notice, however, that globalized Newton-like schemes applied to (P) require additional assumptions (on top of Assumption 1) to guarantee convergence or stability of the

related Newton linear systems arising as sub-problems. Typically, superlinear convergence of Newton-like schemes is observed “close to a solution”. Under additional assumptions, global convergence can be achieved via appropriate line-search or trust-region strategies (e.g., see the developments in [15, 28, 33, 34] and the references therein).

The potential convergence and stability issues arising within Newton-like schemes can be alleviated by combining Newton-like methods with proximal point variants. Solvers based on the proximal point method can achieve superlinear convergence, assuming their penalty parameters increase indefinitely at a suitable rate (e.g., see [35, 36]). For problems (P)–(D), the sub-problems arising within proximal methods are (possibly nonsmooth) convex optimization instances, and are typically solved by means of semismooth Newton strategies. The associated SSN linear systems are better conditioned than their (possibly regularized) interior-point counterparts (e.g., see [3, 23, 24, 37, 38]); however, convergence is expected to be slower, as such algorithms do not enjoy the polynomial worst-case complexity of interior-point methods. Nevertheless, these better conditioned linear systems can, in certain cases, allow one to achieve better computational and/or memory efficiency, especially if the nonsmooth terms are appropriately handled.

Various such solvers have been developed and analyzed in the literature. For example, the authors in [39] developed a dual augmented Lagrangian scheme combined with an SSN method for the solution of semidefinite programming problems and obtained very promising results. Variants of this scheme were then utilized for the solution of linear programming problems in [22], and for lasso-regularized problems in [30]. Similar dual augmented Lagrangian-based solvers (ALM) were recently developed for square-root Lasso problems in [40] as well as for the estimation of sparse Gaussian graphical models in [41]. Nonetheless, those aforementioned dual ALM solvers are brittle, in that slight alterations in the structure of the problem can significantly hinder their employability and efficiency. For example, a computationally attractive extension of [22] to convex quadratic (or, generally, nonseparable convex) programming is by no means obvious and requires a different algorithmic strategy than the dualization approach followed by the authors. Furthermore, there is a **wide gap in the current literature of proximal-SSN methods concerning the efficient solution of the associated SSN linear systems**. While the convergence results appearing in the literature enable the use of inexact linear system solves (e.g., allowing the use of iterative linear algebra, such as Krylov subspace solvers), no general preconditioning strategies have been proposed or analyzed and the question of exploiting efficient linear algebra in this context is still open.

1.3 Contributions

In this paper, we develop an active-set method for (P)–(D) by employing an appropriate proximal method of multipliers (PMM) using a standard backtracking-linesearch semismooth Newton (SSN) strategy for solving the associated PMM sub-problems. The SSN linear systems are approximately solved by means of Krylov subspace methods, using certain novel general-purpose preconditioners. Unlike most proximal point methods given in the literature (e.g., see the primal approaches in [21, 42], the dual approaches in [22, 30] or the primal-dual approaches in [18, 27, 35]), the proposed method introduces proximal terms for each primal and dual variable of the problem, and this results in Newton linear systems which are easy to precondition and solve. We dualize each of the two nonsmooth terms of the objective in (P), which contributes to the simplification of the resulting SSN linear systems, and paves the way for readily analyzing the convergence of the proposed algorithm.

At this point, we should note the inherent difference of our approach compared to the dual augmented Lagrangian schemes developed in [22, 30, 39–41]. Indeed, by avoiding the complete dualization of the problem at hand, and instead focusing on a primal-dual formulation combined with the development of an associated linear algebra kernel, the proposed methodology enables the efficient solution of problems of the form of (P) (with nonseparable, smooth or nonsmooth, terms appearing in the objective function), by fully exploiting their structure.

Our contributions are summarized as follows:

- We are able to directly apply the theory developed in [22] to show that the proposed outer proximal method of multipliers is globally convergent under mere feasibility assumptions, and locally superlinearly convergent under standard additional assumptions. Importantly, we should mention that the theory in [22], although generally applicable to convex programming problems that satisfy an appropriate error bound condition of the form of Lemma 1, was applied in the said paper to analyze the convergence of a proximal dual ALM scheme designed for linear programming. Nonetheless, in this work we showcase that the proposed active-set scheme can be directly analyzed using this theory, and without requiring any additional assumptions other than those employed in [22] in the context of linear programming, despite the generality of model (P). Additionally, by introducing an appropriate augmented Lagrangian penalty function (see Section 2), we obtain continuously differentiable PMM sub-problems. As a result, we are able to apply a standard globally convergent (at a local superlinear rate) backtracking-linesearch semismooth Newton scheme (e.g., see [39]) for their solution. Thus, our active-set scheme is readily theoretically supported, and covers a very wide range of problems not currently addressed by similar methods in the literature, unless (P) is reformulated without exploiting its inherent structure.
- The price to pay for the strong theoretical properties of the proposed approach is that the difficulty arising from the nonsmooth terms is hidden in the solution of the associated SSN linear systems. Nonetheless, we identify a principled way of exploiting the resulting active-set structure, that occurs by utilizing the Clarke subdifferential within SSN, to create novel, robust, and highly efficient preconditioners for the solution of these linear systems by means of Krylov subspace solvers. Specifically, we derive and analyze highly effective preconditioning strategies which we show to be optimal with respect to the penalty parameters of the PMM; that is, the conditioning of the preconditioned linear systems does not deteriorate as we increase the penalty parameters of the PMM (which, in turn, accelerates the convergence of the PMM). Additionally, the preconditioners are versatile and general (i.e., not problem-dependent), and are very manageable in terms of memory requirements, giving a strong computational edge of the proposed approach compared to alternative second-order methods. The developed linear algebra kernel of our method showcases the advantages of utilizing active-set schemes compared to alternative (state-of-the-art) interior point methods for similar problems; such a study was missing from the field of augmented Lagrangian methods, and we argue that it should be capitalized more heavily in similar augmented Lagrangian-type approaches appearing in the literature. To the best of our knowledge, aside from the work in [32], which is specialized to the case of L^1 -regularized PDE constrained optimization, most proximal Newton-like schemes utilizing Krylov subspace methods do so without employing any (nontrivial) preconditioner (e.g., see [26, 30, 40, 42]).
- We provide extensive computational evidence to support that the proposed approach exhibits a significantly better practical convergence than that predicted by the theory,

over a very wide range of applications. Specifically, we showcase the robustness and the efficiency of our approach over a plethora of instances arising from risk-averse portfolio selection problems, L^1 -regularized partial differential equation constrained optimization, quantile regression as well as linear classification via support vector machines. In each of these cases, we compare our approach to a robust regularized interior point method (see [24, 43]) which is available on GitHub, and showcase its efficiency not only in terms of CPU time and memory requirements, but also in terms of total iterations required to obtain a sufficiently accurate solution. Indeed, we demonstrate that while our approach requires more iterations compared to the alternative polynomially convergent interior point method, the discrepancy in terms of iteration count is not particularly significant. This is reflected in the computational efficiency of our approach, in terms of CPU time. At the same time, we demonstrate that both second-order solvers are significantly more robust and efficient compared to a state-of-the-art alternating direction method of multipliers [44] (ADMM), since we are interested in finding sufficiently accurate solutions in our numerical experiments, which prevent first-order schemes from being competitive.

- The method is very versatile and can solve a plethora of very important applications (as we showcase in Section 5), by explicitly exploiting the structure of the nonsmooth terms appearing in the objective of (P). Additionally, it deals with general box constraints, without requiring the introduction of auxiliary variables to deal with upper and lower bounds separately, something that is required when utilizing conic-based solvers. As a result, the associated linear systems solved within SSN have significantly smaller dimensions, compared to linear systems arising within interior point methods suitable for the solution of problems like (P) (e.g., see [3, 24, 25]), potentially making the proposed approach a more attractive alternative for very large-scale instances.

Before closing this section, we should mention that the algorithm developed in this paper arose in an attempt to consolidate two of our previously developed methodologies laid out in two technical reports ([45, 46]); the former report studies (P) by assuming that the nonsmooth terms are separable, while the latter studies (P) in its full generality. However, these previous approaches dealt with certain nonsmooth terms of (P) explicitly. As a result, the associated PMM sub-problems of the developed algorithms were nonsmooth, and that resulted in limited theoretical support of the inner SSN scheme utilized to tackle these sub-problems. This issue was overcome in this consolidated work, by transferring the difficulty to the SSN linear systems, and then utilizing our versatile iterative linear algebra kernel (and associated preconditioning strategy) to ensure that this new and improved version of the active-set schemes developed in [45, 46] is equally scalable but also more robust and better theoretically supported.

To summarize, in Section 2 we derive a proximal method of multipliers and briefly discuss its convergence properties. Then, in Section 3 we present a well-studied globally convergent backtracking-linesearch semismooth Newton scheme used to approximately solve the PMM sub-problems, noticing its local superlinear convergence properties under no additional assumptions. Next, in Section 4, we propose *novel* general-purpose preconditioners for the associated SSN linear systems and analyze their effectiveness. In Section 5, our approach is extensively tested on a wide variety of optimization instances. Finally, we derive some conclusions in Section 6.

Notation. Given a vector x in $\mathbb{R}^n$, $\|x\|$ denotes the Euclidean norm. Letting $R \succ 0_{n,n}$ be a symmetric positive definite matrix, we denote $\|x\|_R^2 = x^\top R x$. Given a closed set $\mathcal{K} \subset \mathbb{R}^n$, we write $\Pi_{\mathcal{K}}(x) \triangleq \arg \min\{\|x - z\| \mid z \in \mathcal{K}\}$, while for any $R \succ 0_{n,n}$ we write $\text{dist}_R(z, \mathcal{K}) \triangleq \inf_{z' \in \mathcal{K}} \|z - z'\|_R$ (with $\text{dist}_I(z, \mathcal{K}) \equiv \text{dist}(z, \mathcal{K})$). Given a proper and closed convex function

$p: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$, we define $\mathbf{prox}_p(u) \triangleq \arg \min_x \{p(x) + \frac{1}{2}\|u - x\|^2\}$. Given an arbitrary rectangular matrix A , $\sigma_{\max}(A)$ denotes its maximum singular value. For an arbitrary square matrix B , $\lambda(B)$ is the set of eigenvalues of B while $\lambda_{\max}(B)$ (resp., $\lambda_{\min}(B)$) denotes its maximum (resp., minimum) eigenvalue. Given an index set $\mathcal{D}$, $|\mathcal{D}|$ denotes its cardinality. Given a rectangular matrix $A \in \mathbb{R}^{m \times n}$ and an index set $\mathcal{B} \subseteq \{1, \dots, n\}$, we denote the columns of A , the indices of which belong to $\mathcal{B}$, as $A_{\mathcal{B}}$. Given a square matrix $Q \in \mathbb{R}^{n \times n}$, we denote the subset of columns and rows of Q , the indices of which belong to $\mathcal{B}$, as $(Q)_{\mathcal{B}, \mathcal{B}}$. We denote by $\text{Diag}(Q)$ the diagonal matrix with diagonal elements equal to those of Q . Given a matrix $A \in \mathbb{R}^{m \times n}$, and two index sets $\mathcal{B}_1 \subseteq \{1, \dots, m\}$ and $\mathcal{B}_2 \subseteq \{1, \dots, n\}$, we denote the subset of rows and columns of A , the indices of which belong to $\mathcal{B}_1, \mathcal{B}_2$ respectively, as $(A)_{\mathcal{B}_1, \mathcal{B}_2}$. Finally, given an arbitrary vector d with n components as well as some indices $1 \leq i_1 \leq i_2 \leq n$, we denote by $d_{i_1:i_2}$ the vector $(d_{i_1}, d_{i_1+1}, \dots, d_{i_2})$. To avoid confusion, indices i are always used to denote entries of vectors or matrices, while k and j are reserved to denote iteration counters (outer (PMM) and inner (SSN), respectively).

2 A Primal-Dual Proximal Method of Multipliers

In what follows, we derive the proximal augmented Lagrangian penalty function corresponding to the primal problem (P). Using the latter, we derive a primal-dual PMM for solving the pair (P)–(D). The convergence of this PMM scheme is subsequently analyzed, assuming that we are able to find sufficiently accurate solutions to its associated sub-problems. In the next section, we briefly present a standard semismooth Newton scheme for the solution of these sub-problems, with global and local superlinear convergence guarantees.

2.1 Derivation of the PMM

We begin by deriving the Lagrangian associated to (P). First, we define the function

$$\varphi(x) \triangleq f(x) + h(Cx + d) + \delta_{\mathcal{K}}(x) + \delta_{\{0_m\}}(b - Ax).$$

Following the dualization strategy proposed in [47, Chapter 11], we let

$$\hat{\varphi}(x, u', w', v') \triangleq f(x) + h(Cx + d + w') + \delta_{\mathcal{K}}(x + v') + \delta_{\{0_m\}}(b - Ax + u'),$$

for which it holds that $\varphi(x) = \hat{\varphi}(x, 0_m, 0_l, 0_n)$. Then, by letting $y = [y_1^\top \ y_2^\top]^\top$, the Lagrangian associated to (P) reads:

$$\begin{aligned} \ell(x, y, z) &\triangleq \inf_{u', w', v'} \left\{ \hat{\varphi}(x, u', w', v') - y_1^\top u' - y_2^\top w' - z^\top v' \right\} \\ &= f(x) - \sup_{u'} \left\{ y_1^\top u' - \delta_{\{0_m\}}(b - Ax + u') \right\} \\ &\quad - \sup_{w'} \left\{ y_2^\top w' - h(Cx + d + w') \right\} - \sup_{v'} \left\{ z^\top v' - \delta_{\mathcal{K}}(x + v') \right\} \\ &= f(x) - y_1^\top (Ax - b) + y_2^\top (Cx + d) - h^*(y_2) + z^\top x - \delta_{\mathcal{K}}^*(z), \end{aligned}$$

where we used the definition of the Fenchel conjugate. Before we proceed with the derivation of the augmented Lagrangian associated to $\ell(\cdot)$, let us observe that the Fenchel conjugate

of $h(\cdot)$ is $h^*(u) = \delta_C(u)$, where $C \triangleq \{u \in \mathbb{R}^l \mid v_i \in [0, 1], i \in \{1, \dots, l\}\}$. Given a penalty parameter $\beta > 0$, the augmented Lagrangian corresponding to (P) reads:

$$\begin{aligned} \mathcal{L}_\beta(x; y, z) \triangleq & f(x) - y_1^\top (Ax - b) + \frac{\beta}{2} \|Ax - b\|^2 \\ & + (Cx + d)^\top \Pi_C(y_2 + \beta(Cx + d)) - \frac{1}{2\beta} \|y_2 - \Pi_C(y_2 + \beta(Cx + d))\|^2 \\ & - \frac{1}{2\beta} \|z\|^2 + \frac{1}{2\beta} \|z + \beta x - \beta \Pi_{\mathcal{K}}(\beta^{-1}z + x)\|^2. \end{aligned} \quad (2.1)$$

A detailed derivation of the augmented Lagrangian penalty function given in (2.1) can be found in Appendix A.2.

Assume that at iteration $k \geq 0$ we have the estimates (x_k, y_k, z_k) as well as the penalty parameters β_k, ρ_k , such that $\rho_k \triangleq \frac{\beta_k}{\tau_k}$, where $\{\tau_k\}_{k=0}^\infty$ is a nonincreasing positive sequence, i.e. $\tau_k > 0$ for all $k \geq 0$. We define the continuously differentiable function

$$\phi(x) \equiv \phi_{\rho_k, \beta_k}(x; x_k, y_k, z_k) \triangleq \mathcal{L}_{\beta_k}(x; y_k, z_k) + \frac{1}{2\rho_k} \|x - x_k\|^2. \quad (2.2)$$

Using the previous notation, we need to find x^* such that

$$(\nabla \phi(x^*))^\top (x - x^*) \geq 0, \quad \forall x \in \mathbb{R}^n,$$

where it can be shown (see Appendix A.3) that

$$\begin{aligned} \nabla \phi(x) = & \nabla f(x) - A^\top y_{1k} + \beta_k A^\top (Ax - b) + C^\top \Pi_C(y_{2k} + \beta_k(Cx + d)) \\ & + (z_k + \beta_k x) - \beta_k \Pi_{\mathcal{K}}(\beta_k^{-1}z_k + x) + \rho_k^{-1}(x - x_k). \end{aligned}$$

We now describe the primal-dual PMM in Algorithm PD-PMM.

Algorithm PD-PMM Primal-dual proximal method of multipliers

Input: $(x_0, y_0, z_0) \in \mathbb{R}^n \times \mathbb{R}^{m+l} \times \mathbb{R}^n$, $\beta_0, \beta_\infty, \tau_\infty > 0$, $\{\tau_k\}_{k=0}^\infty$ such that $\tau_k \searrow \tau_\infty > 0$.

Choose a sequence of positive numbers $\{\epsilon_k\}$ such that $\epsilon_k \rightarrow 0$.

for $(k = 0, 1, 2, \dots)$ **do**

Find x_{k+1} such that:

$$\|\nabla \phi(x_{k+1})\| \leq \epsilon_k. \quad (2.3)$$

$$y_{k+1} = \left[\begin{array}{c} y_{1k} - \beta_k(Ax_{k+1} - b) \\ \Pi_C(y_{2k} + \beta_k(Cx_{k+1} + d)) \end{array} \right]. \quad (2.4)$$

$$z_{k+1} = (z_k + \beta_k x_{k+1}) - \beta_k \Pi_{\mathcal{K}}(\beta_k^{-1}z_k + x_{k+1}). \quad (2.5)$$

$$\beta_{k+1} \nearrow \beta_\infty \leq \infty, \quad \rho_{k+1} = \frac{\beta_{k+1}}{\tau_{k+1}}. \quad (2.6)$$

end for

return (x_k, y_k, z_k) .

Notice that we allow step (2.3) to be computed inexactly. In Section 2.2 we will provide precise conditions on the error sequence guaranteeing that Algorithm PD-PMM achieves

global convergence, and additional conditions for achieving a local linear or superlinear rate (where the local superlinear convergence requires that $\beta_k \rightarrow \infty$). Further conditions on the starting point and on the starting penalty parameter β_0 , required to guarantee a global linear convergence rate, are also discussed.

2.2 Convergence Analysis

In this section we provide conditions on the error sequence $\{\epsilon_k\}$ in (2.3) that guarantee the convergence of Algorithm PD-PMM, potentially at a global linear or local superlinear rate. The analysis is derived by applying the results shown in [22, Section 2] or [48] (or by an extension of the analyses in [35, 36]) after connecting Algorithm PD-PMM to an appropriate proximal point iteration. As before, we let $F = [A^\top - C^\top]^\top$ and $\hat{b} = [b^\top d^\top]^\top$ and define the maximal monotone operator $T_\ell: \mathbb{R}^{2n+m+l} \rightrightarrows \mathbb{R}^{2n+m+l}$, associated to (P)–(D):

$$T_\ell(x, y, z) \triangleq \left\{ (u', v', w') \mid u' \in \nabla f(x) - F^\top y + z, v' \in Fx - \hat{b} + \begin{bmatrix} 0_m \\ \partial h^*(y_2) \end{bmatrix}, w' + x \in \partial \delta_{\mathcal{K}}^*(z) \right\}.$$

The inverse of this operator reads

$$T_\ell^{-1}(u', v', w') \triangleq \arg \max_{y, z} \min_x \left\{ \ell(x, y, z) + u'^\top x - v'^\top y - w'^\top z \right\}.$$

Notice that Assumption 1 implies that $T_\ell^{-1}(0) \neq \emptyset$. Following the result in [22], and since f is convex quadratic while h is piecewise linear, we note that T_ℓ is in fact a polyhedral multifunction (see [49] for a detailed discussion on the properties of such multifunctions). In light of this property of T_ℓ we note that the following specialized *metric subregularity* condition (see [50] for a definition) holds automatically without any additional assumptions.

Lemma 1 *For any $r > 0$, there exists $\kappa > 0$ such that*

$$\text{dist}(p, T_\ell^{-1}(0)) \leq \kappa \text{dist}(0, T_\ell(p)), \quad \forall p \in \mathbb{R}^{2n+m}, \text{ with } \text{dist}(p, T_\ell^{-1}(0)) \leq r. \quad (2.7)$$

Proof The reader is referred to [22, Lemma 2.4] as well as [49]. $\square$

Next, let some sequence of positive definite matrices $\{R_k\}_{k=0}^\infty$ be given with $R_k \triangleq \tau_k I_n \oplus I_{m+l} \oplus I_n$, for all $k \geq 0$, where τ_k is defined in Algorithm PD-PMM and $\oplus$ denotes the direct sum of two matrices. We define the single-valued proximal operator $P_k: \mathbb{R}^{2n+m+l} \mapsto \mathbb{R}^{2n+m+l}$, associated with T_ℓ :

$$P_k \triangleq (R_k + \beta_k T_\ell)^{-1} R_k. \quad (2.8)$$

In particular, under our assumptions on the matrices R_k , we have that (e.g. see [36]) for all $(u_1, v_1, w_1), (u_2, v_2, w_2) \in \mathbb{R}^{2n+m+l}$, the following inequality (nonexpansiveness) holds

$$\|(u_1, v_1, w_1) - P_k(u_2, v_2, w_2)\|_{R_k} \leq \|(u_1, v_1, w_1) - (u_2, v_2, w_2)\|_{R_k}. \quad (2.9)$$

Obviously, we can observe that if $(x^*, y^*, z^*) \in T_\ell^{-1}(0)$, then $P_k(x^*, y^*, z^*) = (x^*, y^*, z^*)$. We are now able to connect Algorithm PD-PMM with the proximal point iteration produced by (2.8).

Proposition 1 Let $\{(x_k, y_k, z_k)\}_{k=0}^\infty$ be a sequence of iterates produced by Algorithm PD-PMM. Then, for every $k \geq 0$ we have that

$$\|(x_{k+1}, y_{k+1}, z_{k+1}) - P_k(x_k, y_k, z_k)\|_{R_k} \leq \frac{\beta_k}{\min\{\sqrt{\tau_k}, 1\}} \|\nabla\phi(x_{k+1})\|. \quad (2.10)$$

Proof From (2.2), we observe that

$$\nabla\phi(x_{k+1}) = \nabla\mathcal{L}_{\beta_k}(x_{k+1}; y_k, z_k) + \frac{1}{2\rho_k}(x_{k+1} - x_k).$$

Thus, by utilizing [35, Proposition 7], we deduce

$$(\nabla\phi(x_{k+1}), 0_{m+l}, 0_n) + \beta_k^{-1}(\tau_k(x_k - x_{k+1}), y_k - y_{k+1}, z_k - z_{k+1}) \in T_\ell(x_{k+1}, y_{k+1}, z_{k+1}).$$

Subtracting $P_k(x_k, y_k, z_k)$ from both sides, taking norms, and using the nonexpansiveness of P_k (see (2.9)), yields (2.10) and concludes the proof. $\square$

Now that we have established the connection of Algorithm PD-PMM with the proximal point iteration governed by the operator P_k defined in (2.8), we can directly provide conditions on the error sequence in (2.3), to guarantee global (possibly linear) and local linear (potentially superlinear) convergence of Algorithm PD-PMM. To that end, we will make use of certain results, as reported in [22, Section 2]. Firstly, we provide the global convergence result for the algorithm.

Theorem 1 Let Assumption 1 hold. Let $\{(x_k, y_k, z_k)\}_{k=0}^\infty$ be generated by Algorithm PD-PMM. Furthermore, assume that we choose a sequence $\{\epsilon_k\}_{k=0}^\infty$ in (2.3), such that

$$\epsilon_k \leq \frac{\min\{\sqrt{\tau_k}, 1\}}{\beta_k} \delta_k, \quad 0 \leq \delta_k, \quad \sum_{k=0}^\infty \delta_k < \infty. \quad (2.11)$$

Then, $\{(x_k, y_k, z_k)\}_{k=0}^\infty$ is bounded and converges to a primal-dual solution of (P)–(D).

Proof The proof is omitted since it is a direct application of [22, Theorem 2.3]. $\square$

Next, we discuss local linear (and potentially superlinear) convergence of Algorithm PD-PMM. To that end, let $r > \sum_{k=0}^\infty \delta_k$, where δ_k is defined in (2.11). Then, from Lemma 1 we know that there exists $\kappa > 0$ associated with r such that

$$\text{dist}((x, y, z), T_\ell^{-1}(0)) \leq \kappa \text{dist}(0, T_\ell(x, y, z)), \quad (2.12)$$

for all $(x, y, z) \in \mathbb{R}^{2n+m}$ such that $\text{dist}((x, y, z), T_\ell^{-1}(0)) \leq r$.

Theorem 2 Let Assumption 1 hold, and assume that (x_0, y_0, z_0) is chosen to satisfy

$$\text{dist}_{R_0}((x_0, y_0, z_0), T_\ell^{-1}(0)) \leq r - \sum_{k=0}^\infty \delta_k,$$

where $\{\delta_k\}_{k=0}^\infty$ is given in (2.11). Let also κ be given as in (2.12) and assume that we choose a sequence $\{\epsilon_k\}_{k=0}^\infty$ in (2.3) such that

$$\epsilon_k \leq \frac{\min\{\sqrt{\tau_k}, 1\}}{\beta_k} \min\{\delta_k, \delta'_k \|(x_{k+1}, y_{k+1}, z_{k+1}) - (x_k, y_k, z_k)\|_{R_k}\}, \quad (2.13)$$

where $0 \leq \delta_k$, $\sum_{k=0}^{\infty} \delta_k < \infty$, and $0 \leq \delta'_k < 1$, $\sum_{k=0}^{\infty} \delta'_k < \infty$. Then, for all $k \geq 0$ we have that

$$\text{dist}_{R_{k+1}} \left((x_{k+1}, y_{k+1}, z_{k+1}, T_\ell^{-1}(0)) \right) \leq \mu_k \text{dist}_{R_k} \left((x_k, y_k, z_k, T_\ell^{-1}(0)) \right),$$

$$\mu_k \triangleq (1 - \delta'_k)^{-1} \left(\delta'_k + (1 + \delta'_k) \frac{\kappa \gamma_k}{\sqrt{\beta_k^2 + \kappa^2 \gamma_k^2}} \right), \quad \lim_{k \rightarrow \infty} \mu_k = \mu_\infty \triangleq \frac{\kappa \gamma_\infty}{\sqrt{\beta_\infty^2 + \kappa^2 \gamma_\infty^2}},$$

where $\gamma_k \triangleq \max\{\tau_k, 1\}$ and $\gamma_\infty = \max\{\tau_\infty, 1\}$ (noting that $\mu_\infty = 0$, if $\beta_\infty = \infty$).

Proof The proof is omitted since it follows by direct application of [22, Theorem 2.5] (see also [36, Theorem 2]). $\square$

Remark 2 Following [22, Remarks 2, 3], we can choose a nonincreasing sequence $\{\delta'_k\}_{k=0}^{\infty}$ and a large enough β_0 such that $\mu_0 < 1$, which in turn implies that $\mu_k \leq \mu_0 < 1$, yielding a global linear convergence of both $\text{dist}((x_k, y_k, z_k), T_\ell^{-1}(0))$ as well as $\text{dist}_{R_k}((x_k, y_k, z_k), T_\ell^{-1}(0))$, assuming that the starting point of the algorithm satisfies the assumption stated in Theorem 2. On the other hand, as is implicitly mentioned in Theorem 2, if β_k is forced to increase indefinitely, we obtain a local superlinear convergence rate (notice that $\mu_\infty = 0$ if $\beta_\infty = \infty$).

3 Semismooth Newton Method

In this section we briefly present a standard semismooth Newton (SSN) scheme suitable for the solution of problem (2.3), appearing in Algorithm PD-PMM. To that end, we fix the estimates (x_k, y_k, z_k) as well as the penalty parameters β_k, ρ_k . We set $x_{k_0} = x_k$, and at every iteration j of SSN, we approximately solve a system of the following form:

$$J_{k_j} d_x = -\nabla \phi(x_{k_j}), \quad (3.1)$$

where $J_{k_j} \in \partial_x^C(\nabla \phi(x_{k_j}))$. The symbol $\partial_x^C(\cdot)$ denotes the *Clarke subdifferential* of a function (see [51]) with respect to x , which can be obtained as the convex hull of the *Bouligand subdifferential* ([51]). Any element of the Clarke subdifferential is a *Newton derivative* (see [52, Chapter 13]), since $\nabla \phi(x_{k_j})$ is a *piecewise continuously differentiable and regular function*. Using [52, Theorem 14.8], we obtain that for any interval $[l, u]$ on the real line (i.e., $l < u$):

$$\partial_w^C(\Pi_{[l,u]}(w_i)) = \begin{cases} \{1\}, & \text{if } w \in (l, u), \\ \{0\}, & \text{if } w \notin [l, u], \\ [0, 1], & \text{if } w \in \{l, u\}. \end{cases}$$

For computational reasons, we always choose matrices M_{k_j} from the Bouligand subdifferential. This choice will become apparent in the Section 4. On the other hand, it is well-known (see [16, Theorem 4]) that an inexact semismooth Newton scheme using the Bouligand subdifferential converges at a local linear rate (assuming that the linear systems are solved up to an appropriate accuracy), if the equation $\nabla \phi(x) = 0_n$ is *BD-regular* at the optimum (x_k^*) (that is, each element of the Bouligand subdifferential of $\nabla \phi(x_k^*)$ is nonsingular). Note, however, that since we employ the semismooth Newton scheme to solve the sub-problems arising from Algorithm PD-PMM, we obtain that the resulting nonsmooth equations are indeed BD-regular for every outer iteration $k \geq 0$. Thus, assuming that the associated linear systems are solved up to a sufficient accuracy (see [16, Theorem 4]), we obtain local linear

convergence rate of the resulting inexact semismooth Newton scheme. If, additionally, the solution accuracy of the associated linear systems is increased at a suitable rate, the resulting local rate can be superlinear.

We complete the derivation of the SSN by introducing backtracking line-search to ensure global convergence. Then, under no additional assumptions, one can show that SSN is globally convergent. Algorithm SSN outlines a semismooth Newton method for the solution of (2.3). We assume that the associated linear systems are approximately solved by means of a Krylov subspace method. An analysis of the effect of errors arising from the use of Krylov methods within SSN applied to nonsmooth equations can be found in [26].

Algorithm SSN Semismooth Newton method

Input: $\epsilon_k > 0$, $\mu \in (0, \frac{1}{2})$, $\delta \in (0, 1)$, $\gamma \in (0, 1]$, $\bar{\eta} \in (0, 1)$, $x_{k_0} = x_k$, $y_{k_0} = y_k$.

for ($j = 0, 1, 2, \dots$) **do**

 Choose $J_{k_j} \in \partial_x^C (\nabla \phi(x_{k_j}))$, and solve

$$J_{k_j} d_x = -\nabla \phi(x_{k_j}),$$

$$\text{such that } \|J_{k_j} d_x + \nabla \phi(x_{k_j})\| \leq \min \left\{ \bar{\eta}, \|\nabla \phi(x_{k_j})\|^{1+\gamma} \right\}.$$

(Line-search) Set $\alpha_j = \delta^{m_j}$, where m_j is the first nonnegative integer for which:

$$\phi(x_{k_j} + \delta^{m_j} d_x) \leq \phi(x_{k_j}) + \mu \delta^{m_j} (\nabla \phi(x_{k_j}))^\top d_x$$

$$x_{k_{j+1}} = x_{k_j} + \alpha_j d_x.$$

if ($\|\nabla \phi(x_{k_{j+1}})\| \leq \epsilon_k$) **then**

return $x_{k_{j+1}}$.

end if

end for

Below we provide a theorem characterizing the convergence of Algorithm SSN.

Theorem 3 *Let $\{x_{k_j}\}_{j=0}^\infty$ be the sequence of iterates generated by Algorithm SSN. Then, this sequence converges to the unique optimal solution x_k^* of $\min_x \phi(x)$, and $\|x_{k_{j+1}} - x_k^*\| = \mathcal{O}(\|x_{k_j} - x_k^*\|^{1+\gamma})$.*

Proof The proof is omitted since it follows by a trivial extension of the results given in [39, Proposition 3.3, Theorems 3.4–3.5]. $\square$

4 The SSN Linear Systems

The major bottleneck of the proposed active-set scheme is the approximate solution of the linear systems given in (3.1). Since Algorithm SSN does not require an exact solution, we can utilize preconditioned Krylov subspace solvers for the efficient solution of such systems. In what follows we will derive an equivalent reformulation of (3.1), and we will subsequently propose and analyze a robust preconditioner for the acceleration of the associated Krylov subspace method employed for its solution.

4.1 Derivation of the Linear Systems

Before we proceed with the derivation of the linear systems, we first need to note that our implementation treats separable nonsmooth terms differently for improved memory and computational efficiency. Specifically, if the problem given in (P) contains a term $\sum_{i=1}^n ((Dx + d)_i)_+$, where D is a diagonal ("weight") matrix, we will discuss why and how this is treated differently. To that end, we assume that the matrix C appearing in (P) has the form $C = [\hat{C}^\top D]^\top$, where $\hat{C} \in \mathbb{R}^{s \times n}$ is some rectangular matrix, and $D \in \mathbb{R}^{n \times n}$ is a diagonal square matrix, with $s + n = l$. While this structure might seem rather specific, it is very common as it appears when dealing with sparse approximation problems, in which the objective function contains a term of the form $\|Dx\|_1$.

Let $k, j \geq 0$ be some arbitrary iterations of Algorithm PD-PMM, and SSN, respectively. Given the assumed structure of matrix C , let us re-write the gradient of (2.2):

$$\begin{aligned} \nabla \phi(x) = & \nabla f(x) - A^\top y_{1k} + \beta_k A^\top (Ax - b) + \hat{C}^\top \Pi_{C_s} \left((y_{2k})_{1:s} + \beta_k (\hat{C}x + (d)_{1:s}) \right) \\ & + D \Pi_{C_n} \left((y_{2k})_{s+1:l} + \beta_k (Dx + (d)_{s+1:l}) \right) \\ & + (z_k + \beta_k x) - \beta_k \Pi_K (\beta_k^{-1} z_k + x) + \rho_k^{-1} (x - x_k), \end{aligned}$$

where $C_q = \{u \in \mathbb{R}^q \mid u_i \in [0, 1], i \in \{1, \dots, q\}\}$. We notice that any element

$$B_{\delta, k_j} \in \partial_w^C (\Pi_K(w)) \big|_{w=\beta_k^{-1} z_k + x_{k_j}}$$

yields a Newton derivative (see [52, Theorem 14.8]). The same applies for any

$$B_{h, k_j} \in \partial_w^C (\Pi_C(w)) \big|_{w=y_{2k} + \beta_k (Cx_{k_j} + d)}.$$

We choose $B_{\delta, k_j}, B_{h, k_j}$ from the Bouligand subdifferential to improve computational efficiency. We set $B_{\delta, k_j} \in \mathbb{R}^{n \times n}$ as the diagonal matrix with

$$(B_{\delta, k_j})_{i,i} \triangleq \begin{cases} 1, & \text{if } \beta_k^{-1} z_{k,i} + x_{k_j,i} \in (a_{l_i}, a_{u_i}), \\ 0, & \text{otherwise,} \end{cases} \quad (4.1)$$

for $i \in \{1, \dots, n\}$. On the other hand, we separate two cases for $B_{h, k_j} \in \mathbb{R}^{l \times l}$. Specifically, we set

$$\begin{aligned} (B_{h, k_j})_{i,i} &\triangleq \begin{cases} 1, & \text{if } (y_k)_i + \beta_k (Cx + d)_i \in (0, 1), \\ 0, & \text{otherwise,} \end{cases} \quad \text{for } i \in \{1, \dots, s\}, \text{ and} \\ (B_{h, k_j})_{i,i} &\triangleq \begin{cases} 1, & \text{if } (y_k)_i + \beta_k (Cx + d)_i \in [0, 1], \\ 0, & \text{otherwise,} \end{cases} \quad \text{for } i \in \{s+1, \dots, l\}. \end{aligned} \quad (4.2)$$

Having defined these, we can now formulate the chosen Newton derivative of $\nabla \phi(x_{k_j})$, as

$$J_{k_j} = Q + \beta_k A^\top A + \beta_k C^\top B_{h, k_j} C + \beta_k I_n - \beta_k B_{\delta, k_j} + \rho_k^{-1} I_n.$$

However, by letting $B_{h, k_j}^1 \triangleq (B_{h, k_j})_{1:s, 1:s}$ and $B_{h, k_j}^2 \triangleq (B_{h, k_j})_{s+1:l, s+1:l}$, and by introducing two auxiliary variables w_1, w_2 , we observe that the semismooth Newton system in (3.1) can equivalently be written as

$$\begin{bmatrix} -(Q + \rho_k^{-1} I_n + \beta_k (I_n - B_{\delta, k_j}) + \beta_k D B_{h, k_j}^2 D) & A^\top & -\hat{C}^\top \\ A & \beta_k^{-1} I_m & 0_{m,s} \\ -B_{h, k_j}^1 \hat{C} & 0_{s,m} & \beta_k^{-1} I_s \end{bmatrix} \begin{bmatrix} dx \\ w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} \xi_1 \\ \xi_2 \\ \xi_3 \end{bmatrix}, \quad (4.3)$$

where

$$\begin{bmatrix} \xi_1 \\ \xi_2 \\ \xi_3 \end{bmatrix} \triangleq \begin{bmatrix} c + Qx_{k_j} + D\Pi_{C_n}((y_{2k})_{s+1:l} + \beta_k(Dx_{k_j} + (d)_{s+1:l})) + \tilde{z}_{k_j} + \rho_k^{-1}(x_{k_j} - x_k) \\ \beta_k^{-1}y_{1k} - (Ax_{k_j} - b) \\ \beta_k^{-1}\Pi_{C_s}((y_{2k} + \beta_k(Cx_{k_j} + d))_{1:s}) \end{bmatrix},$$

with $\tilde{z}_{k_j} \triangleq (z_k + \beta_k x_{k_j}) - \beta_k \Pi_K(\beta_k^{-1} z_k + x_{k_j})$. Derived from (4.2), we define four index sets:

$$\mathcal{B}_{h,k_j}^1 \triangleq \left\{ i \in \{1, \dots, s\} \mid (B_{h,k_j}^1)_{i,i} = 1 \right\}, \quad \mathcal{N}_{h,j}^1 \triangleq \{1, \dots, s\} \setminus \mathcal{B}_{h,j}^1,$$

and

$$\mathcal{B}_{h,k_j}^2 \triangleq \left\{ i \in \{1, \dots, n\} \mid (B_{h,k_j}^2)_{i,i} = 1 \right\}, \quad \mathcal{N}_{h,j}^2 \triangleq \{1, \dots, n\} \setminus \mathcal{B}_{h,j}^2.$$

Observe that

$$B_{h,k_j}^1 \hat{C} = \mathcal{P} \begin{bmatrix} (\hat{C})_{\mathcal{B}_{h,k_j}^1, 1:n} \\ 0_{|\mathcal{N}_{h,k_j}^1|, n} \end{bmatrix},$$

where $\mathcal{P}$ is an appropriate permutation matrix. Then, from (4.3) we obtain

$$(w_2)_{\mathcal{N}_{h,k_j}^1} = (\Pi_{C_s}((y_{2k} + \beta_k(Cx + d))_{1:s}))_{\mathcal{N}_{h,k_j}^1},$$

and (4.3) reduces to

$$\underbrace{\begin{bmatrix} -(Q + \rho_k^{-1}I_n + \beta_k(I_n - B_{\delta,k_j}) + \beta_k DB_{h,k_j}^2 D) & A^\top & -(\hat{C}^\top)_{\mathcal{B}_{h,k_j}^1} \\ A & \beta_k^{-1}I_m & 0_{m, |\mathcal{B}_{h,k_j}^1|} \\ -(\hat{C})_{\mathcal{B}_{h,k_j}^1, 1:n} & 0_{|\mathcal{B}_{h,k_j}^1|, m} & \beta_k^{-1}I_{|\mathcal{B}_{h,k_j}^1|} \end{bmatrix}}_{M_{k_j}} \begin{bmatrix} d_x \\ w_1 \\ (w_2)_{\mathcal{B}_{h,k_j}^1} \end{bmatrix} \quad (4.4)$$

$$= \begin{bmatrix} \xi_1 \\ \xi_2 \\ (\xi_3)_{\mathcal{B}_{h,k_j}^1} \end{bmatrix} + \begin{bmatrix} (\hat{C}^\top)_{\mathcal{N}_{h,k_j}^1} (w_2)_{\mathcal{N}_{h,k_j}^1} \\ 0_m \\ 0_{|\mathcal{B}_{h,k_j}^1|} \end{bmatrix},$$

the coefficient matrix of which is *quasi-definite* (see [53]), invertible, and its conditioning can be directly controlled by the regularization parameters (β_k, ρ_k) of Algorithm PD-PM.

4.2 Preconditioning and Iterative Solution Via Krylov Subspace Methods

Next, we would like to construct an effective preconditioner for M_{k_j} (i.e., the coefficient matrix of (4.4)). Before delving into this, let us observe that in the derivation of this linear system, we utilized the active-set $\mathcal{B}_{h,k_j}^1$ to reduce the size of the resulting coefficient matrix. As before, let us define

$$\mathcal{B}_{\delta,k_j} \triangleq \left\{ i \in \{1, \dots, n\} \mid (B_{\delta,k_j})_{i,i} = 1 \right\}, \quad \mathcal{N}_{\delta,k_j} \triangleq \{1, \dots, n\} \setminus \mathcal{B}_{\delta,k_j}.$$

In what follows, we will make use of the active sets $\mathcal{B}_{h,k_j}^2$ and $\mathcal{B}_{\delta,k_j}$ to derive cheap and effective preconditioners for the iterative solution of (4.4) by means of Krylov subspace solvers. For simplicity of notation, let us define

$$H_{k_j} \triangleq Q + \rho_k^{-1} I_n + \beta_k (I_n - B_{\delta,k_j}) + \beta_k D B_{h,k_j}^2 D,$$

and consider an approximation of this matrix that reads

$$\tilde{H}_{k_j} \triangleq \text{Diag}(Q) + \rho_k^{-1} I_n + \beta_k (I_n - B_{\delta,k_j}) + \beta_k D B_{h,k_j}^2 D.$$

Let us also define the compound matrix

$$G \triangleq \begin{bmatrix} A \\ -(\hat{C})_{\mathcal{B}_{h,k_j}^1, 1:n} \end{bmatrix},$$

where we drop the iteration dependence for simplicity, and two sets of interest

$$\mathcal{B}_{k_j} \triangleq \mathcal{B}_{\delta,k_j} \cap \mathcal{N}_{h,k_j}^2, \quad \mathcal{N}_{k_j} \triangleq \{1, \dots, n\} \setminus \mathcal{B}_{k_j}.$$

We are now ready to define a positive definite preconditioner for the coefficient matrix in (4.4), as follows:

$$P_{k_j} \triangleq \begin{bmatrix} \tilde{H}_{k_j} & 0_{n, m+|\mathcal{B}_{h,k_j}^1|} \\ 0_{m+|\mathcal{B}_{h,k_j}^1|, n} & \left(G E_{k_j} G^\top + \beta_k^{-1} I_{m+|\mathcal{B}_{h,k_j}^1|} \right) \end{bmatrix}, \quad (4.5)$$

with $E_{k_j} \in \mathbb{R}^{n \times n}$ the diagonal matrix defined as

$$(E_{k_j})_{(i,i)} \triangleq \begin{cases} \frac{1}{(\tilde{H}_{k_j})_{i,i}}, & \text{if } i \in \mathcal{B}_{k_j}, \\ 0, & \text{otherwise.} \end{cases} \quad (4.6)$$

Thus, we utilize a block diagonal preconditioner in which we drop all columns of G that belong to $\mathcal{N}_{k_j}$. This provides a clear motivation behind our choice of matrices in (4.1) and (4.2).

The preconditioner in (4.5) is an extension of the preconditioner proposed in [37] for the solution of linear systems arising from the application of a regularized interior point method to convex quadratic programming. Being a diagonal matrix, E_{k_j} yields a sparse approximation of the Schur complement of the saddle-point matrix in (4.5). This approximation is then used to construct a positive definite block-diagonal preconditioner (i.e. P_{k_j}), which can be used within a symmetric Krylov solver, like the minimum residual (MINRES) method (see [54]). The (2, 2) block of P_{k_j} can be inverted via a Cholesky decomposition.

Next, we analyze the spectral properties of the preconditioned matrix $(P_{k_j})^{-1} M_{k_j}$. Let

$$S_{P_{k_j}} \triangleq \left(G \tilde{H}_{k_j}^{-1} G^\top + \beta_k^{-1} I_m \right), \quad S_{k_j} \triangleq \left(G E_{k_j} G^\top + \beta_k^{-1} I_m \right).$$

In the following lemma, we bound the eigenvalues of the preconditioned matrix $S_{P_{k_j}}^{-1} S_{k_j}$. This is subsequently used to analyze the spectrum of $P_{k_j}^{-1} M_{k_j}$.

Lemma 2 *For any iterates k and j of Algorithms PD-PMM and SSN, respectively, we have*

$$1 \leq \lambda \leq 1 + \sigma_{\max}^2(\hat{F}) \left(\frac{1}{2 + \beta_{\infty}^{-2} \tau_{\infty}} \right),$$

where $\lambda \in \lambda(S_{P_{k_j}}^{-1} S_{k_j})$, $\hat{F} = [A^\top - \hat{C}^\top]^\top$, and $\beta_\infty, \tau_\infty$ are defined in Algorithm PD-PMM.

Proof Consider the preconditioned matrix $S_{P_{k_j}}^{-1} S_{k_j}$, and let (λ, u) be its eigenpair. Then, λ must satisfy the following equation:

$$\lambda = \frac{u^\top \left(G_{\mathcal{B}_{k_j}} D_{\mathcal{B}_{k_j}} G_{\mathcal{B}_{k_j}}^\top + \beta_k^{-1} I_m + G_{\mathcal{N}_{k_j}} D_{\mathcal{N}_{k_j}} G_{\mathcal{N}_{k_j}}^\top \right) u}{u^\top \left(G_{\mathcal{B}_{k_j}} D_{\mathcal{B}_{k_j}} G_{\mathcal{B}_{k_j}}^\top + \beta_k^{-1} I_m \right) u},$$

where $D_{\mathcal{B}_{k_j}} = \left(\tilde{H}_{k_j}^{-1} \right)_{(\mathcal{B}_{k_j}, \mathcal{B}_{k_j})}$, $D_{\mathcal{N}} = \left(\tilde{H}_{k_j}^{-1} \right)_{(\mathcal{N}_{k_j}, \mathcal{N}_{k_j})}$. The above equality holds since $\tilde{H}_{k_j}$ is a diagonal matrix, and $(E_{k_j})_{(i,i)} = 0$ for every $i \in \mathcal{N}_{k_j}$ (indeed, see the definition in (4.6)). Hence, from positive semi-definiteness of Q , we obtain

$$\begin{aligned} 1 \leq \lambda &= 1 + \frac{u^\top \left(G_{\mathcal{N}_{k_j}} D_{\mathcal{N}_{k_j}} G_{\mathcal{N}_{k_j}}^\top \right) u}{u^\top \left(G_{\mathcal{B}_{k_j}} D_{\mathcal{B}_{k_j}} G_{\mathcal{B}_{k_j}}^\top + \beta_k^{-1} I_m \right) u} \leq 1 + \beta_k \sigma_{\max}^2(G_{\mathcal{N}_{k_j}}) \left(2\beta_k + \rho_k^{-1} \right)^{-1} \\ &\leq 1 + \sigma_{\max}^2(\hat{F}) \left(\frac{1}{2 + \beta_k^{-2} \tau_k} \right) \leq 1 + \sigma_{\max}^2(\hat{F}) \left(\frac{1}{2 + \beta_\infty^{-2} \tau_\infty} \right), \end{aligned}$$

where we used that $\rho_k = \beta_k / \tau_k$, $\beta_k \leq \beta_\infty$, and $\tau_k \geq \tau_\infty$. $\square$

Given Lemma 2, we are now able to invoke [37, Theorem 3] to characterize the spectral properties of the preconditioned matrix $P_{k_j}^{-1} M_{k_j}$. Let

$$\begin{aligned} \bar{S}_{k_j} &\triangleq \left(S_{P_{k_j}} \right)^{-\frac{1}{2}} S_{k_j} \left(S_{P_{k_j}} \right)^{-\frac{1}{2}}, & \bar{H}_{k_j} &\triangleq \tilde{H}_{k_j}^{-1/2} H_{k_j} \tilde{H}_{k_j}^{-1/2}, \\ \alpha_{NE} &\triangleq \lambda_{\min}(\bar{S}_{k_j}), & \beta_{NE} &\triangleq \lambda_{\max}(\bar{S}_{k_j}), & \alpha_H &\triangleq \lambda_{\min}(\bar{H}_{k_j}), & \beta_H &\triangleq \lambda_{\max}(\bar{H}_{k_j}). \end{aligned}$$

Notice that Lemma 2 yields upper and lower bounds for α_{NE} and β_{NE} . From the definition of $\bar{H}_{k_j}$ we can also obtain that $\alpha_H \leq 1 \leq \beta_H$ (see [37]). We are now ready to state the spectral properties of the preconditioned matrix $P_{k_j}^{-1} M_{k_j}$.

Theorem 4 Let k and j be some arbitrary iterates of Algorithms PD-PMM and SSN, respectively. Then, the eigenvalues of $P_{k_j}^{-1} M_{k_j}$ lie in the union of the following intervals:

$$I_- \triangleq \left[-\beta_H - \sqrt{\beta_{NE}}, -\alpha_H \right], \quad I_+ \triangleq \left[\frac{1}{1 + \beta_H}, 1 + \sqrt{\beta_{NE} - 1} \right].$$

Proof The proof follows by direct application of [37, Theorem 3]. $\square$

Remark 3 By combining Lemma 2 with Theorem 4, we can observe that the eigenvalues of the preconditioned matrix $P_{k_j}^{-1} M_{k_j}$ are not deteriorating as $\beta_k \rightarrow \infty$. In other words, the preconditioner is robust with respect to the penalty parameters β_k, ρ_k of Algorithm PD-PMM. Furthermore, our choices of $B_{h,k_j}, B_{\delta,k_j}$ in (4.1) and (4.2) serve the purpose of further sparsifying the preconditioner in (4.5), thus potentially further sparsifying its Cholesky decomposition.

We note that the preconditioner in (4.5) is a very efficient choice if $n \geq m + s$, which is the case in several of the experiments considered in Section 5. However, any of the preconditioners given in [24] can be directly applied for systems appearing in the proposed inner-outer

scheme. For example, in the case where $m + s \geq n$ one could adapt the preconditioner given in [24, Section 3.2] (using the developments of this section), which would (potentially) be a more efficient choice when compared to that given in (4.5). Additionally, we should mention that for problems considered within this work, a diagonal approximation of the Hessian (within the preconditioner) seems sufficient to deliver very good performance. Indeed, this is the case for a wide range of problems. However, in certain instances, one might consider nondiagonal approximations of the Hessian. In that case, the preconditioner in (4.5) can be readily generalized and analyzed, following the developments in [24, Section 3.1]. The aforementioned extensions are omitted here for brevity of exposition, but should be considered in a production implementation of the proposed method.

Before closing this section, let us briefly explain the difference between the preconditioning approach proposed in this manuscript and the preconditioner proposed in [37] in the context of regularized IPMs. Specifically, both preconditioners belong to the class of column-dropping preconditioners for primal-dual regularized saddle-point systems (as specified in [24]). However, both [24] and [37] study this class in the context of regularized interior point methods. Thus, the core difference between this preconditioner and the one employed in the present paper lies in the strategy employed to decide which columns of the (2,1) and (3,1) blocks of the Newton system (4.3) to drop when building our preconditioner. Specifically, in the context of regularized IPMs, one has to employ a heuristic mechanism of deciding which columns should be dropped by inspecting the diagonal values of the (1,1) block of the Newton system and subsequently by comparing them to the current value of the barrier parameter of the IPM. As a result, the conditioning of the associated preconditioned normal equations, in the IPM case, depends on some dynamically chosen constant C_k (see [37, Theorem 1]) which in practice needs to be continually adjusted based on the observed Krylov iterations required to reach a desired accuracy in the previous IPM iteration(s). This can result in added computational times, and potential loss of robustness (which in certain cases might lead to a repeated solution of the Newton system with multiple re-computations of the preconditioner).

This is completely alleviated in the setting of the present paper, which gives a straightforward rule for deciding on which columns of the (2,1) and (3,1) blocks of the Newton matrix should be dropped, based on a simple membership condition determined by the associated active-sets at each outer-inner iteration k_j (see (4.6)). In other words, we showcase that column dropping preconditioning strategies, like the one developed in this work, are much better adapted to active-set methods (assuming that such active-set methods are appropriately designed to accommodate for the efficient application of such preconditioners; indeed, this fact has informed our algorithmic design in Sections 2, 3), despite their original conception in the context of regularized IPMs. This is also reflected on the associated spectral analysis given in Lemma 2. Indeed, the eigenvalues of $S_{P_{k_j}}^{-1} S_{k_j}$ are not only independent of the penalty parameters of the PMM, but can also be tightly bounded by a constant independent of the PMM-SNN iteration (something that does not hold in the regularized IPM case; see [37, Theorem 1]).

4.3 A Short Comparison to Similar Primal-Dual AL Methods For QP

In an attempt to place our method in the literature, we note that the most closely related primal-dual augmented Lagrangian algorithms are those developed in [18, 20] in the context of convex quadratic programming (noting that the latter also enables the solution of nonconvex

quadratic problems). Both these methods are relatively similar to each other, so we focus on [18] for simplicity. In the aforementioned paper, the authors develop a primal-dual proximal point algorithm for the solution of convex quadratic programs of the form

$$\min_{x \in \mathbb{R}^n} c^\top x + \frac{1}{2} x^\top Q x + \delta_{\mathcal{K}}(z), \quad \text{s.t. } Ax = z,$$

observing that this convex QP formulation is noticeably different from the formulation utilized in this work. Their method employs the primal-dual proximal augmented Lagrangian merit function (see [27]), that reads (given some primal-dual estimate $(\bar{x}, \bar{y})$):

$$\begin{aligned} \mathcal{M}(x, z, y; \bar{x}, \bar{y}) = & c^\top x + \frac{1}{2} x^\top Q x + \delta_{\mathcal{K}}(z) + y^\top (Ax - z) + \frac{1}{2\mu} \|Ax - z\|^2 \\ & + \frac{1}{2\mu} \|z - Ax + \mu(y - \bar{y})\|^2 + \frac{\sigma}{2} \|x - \bar{x}\|^2. \end{aligned}$$

Then, by applying SSN on this merit function, the resulting SSN linear systems take the form (after appropriate reductions):

$$Mw = r, \quad M = \begin{bmatrix} Q + \sigma_k I_n & A^\top (I - P_k) \\ (I - P_k)A & -\mu(I - P_k/2) \end{bmatrix},$$

where $I - P_k$ is a diagonal “projection” matrix (with 1’s or 0’s in the diagonal, depending on the activity of each inequality constraint) arising from the utilization of the semismooth Newton method, and r is an appropriate right hand side. Note that, unlike our method, the active-set projection matrices are restricted to the (2,1) Newton block (i.e., they are applied to matrix A) and there is not an obvious way to transfer them to the (1,1) block (as in (4.3), where we chose beforehand which projection matrices will appear on the (3,1) and which on the (1,1) block of the SSN linear system). Thus, for example, any separable nonsmooth term in the objective (such as an ℓ_1 -norm regularizer or a separable inequality constraint) would not yield any significant memory reductions in the factorization of such a SSN linear system matrix (after reformulation of the problem in the convex QP form discussed in this subsection and the utilization of the algorithm presented in [18]). Instead, our method enables us to treat separable nonsmooth terms specially, which results in certain projection matrices appearing in the (1,1) block of our SSN linear systems. Then, by utilizing our preconditioning strategy, we can exploit these projection matrices associated to such separable nonsmooth terms to meaningfully reduce the memory requirements of our method (as is demonstrated in Section 4.2).

Let us now briefly only focus on inequality constraints (ignoring piecewise-linear terms). Specifically, we should mention that the way that inequality constraints are handled in this work and in [18, 20] are complementary and can also be used at the same time for separable and nonseparable inequalities, respectively. Our methodology enables (additionally) a principled mechanism of exploiting either separable or nonseparable piecewise-linear terms in the objective (in distinct ways) to reduce the memory required to effectively solve the associated SSN linear systems.

5 Applications and Numerical Results

In this section we present various applications that can be modeled by problem (P). Specifically, we consider portfolio optimization, L^1 -regularized partial differential equation (PDE) constrained optimization, quantile regression, and binary classification problems. We first

discuss (and numerically demonstrate) the effectiveness of the approach for the solution of single-period mean-risk portfolio optimization problems, where risk is measured via the *conditional value at risk* or the *mean absolute semi-deviation*. Subsequently, we apply the proposed scheme to solve L^1 -regularized Poisson optimal control problems. Then, we test the proposed approach on quantile regression problems, and finally on binary classification problems via linear support vector machines.

Implementation details. Before proceeding to the numerical results, we mention certain implementation details of the proposed algorithm. Following the discussion in Section 4, our implementation takes problems of the following form as input:

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & \left\{ c^\top x + \frac{1}{2} x^\top Q x + \sum_{i=1}^l ((Cx + d)_i)_+ + \|Dx\|_1 + \delta_K(x) \right\}, \\ \text{s.t.} \quad & Ax = b, \end{aligned}$$

where $D \in \mathbb{R}^{n \times n}$ is a diagonal weight matrix. As discussed in Section 4, the separable nonsmooth term $\|Dx\|_1$ is treated separately as shown in the derivation of the SSN linear systems. The MATLAB implementation can be found on GitHub¹. The experiments are run on a PC with a 2.3GHz Intel core i7-12700H processor, 32GB of RAM, using the Windows 11 operating system.

We initialize the PMM penalty parameters as $\beta_0 = 50$, and $\rho_0 = 100$. These are increased in subsequent iterations, at a rate that depends on the residual reduction achieved in that particular PMM iteration. When solving the PMM sub-problems using Algorithm SSN, we use a predictor-corrector-like heuristic in which the first iteration is accepted without line-search and then line-search is activated for subsequent iterations. Algorithm PD-PMM is allowed to run for at most 200 iterations, while in each outer iteration Algorithm SSN is run for at most 40 steps. We utilize the minimum residual (MINRES) method [54] for the solution of the associated SSN linear systems, which is allowed to run for at most 150 iterations. At the beginning of the optimization process the proposed scheme does not utilize the preconditioner given in (4.5), unless MINRES reaches 100 iterations (without a preconditioner). Unless stated otherwise, a primal-dual iterate is accepted as optimal if the conditions given in (A.3; see also Appendix A.4 for additional details) are satisfied for the tolerance specified by the user. All other implementation details follow directly from the developments in Sections 2–4. We refer the reader to the implementation on GitHub for additional details.

In the presented experiments, we compare the proposed approach against the robust regularized interior point solver given in [24, 43] (known as IP-PMM; *interior point-proximal method of multipliers*), the exact and inexact implementations of which can be found on GitHub^{2,3}. There are several reasons why the comparison against IP-PMM is meaningful. Firstly, both methods can be classified as second-order solvers. Secondly, both solvers scale well with the problem size, since the inexact variant of IP-PMM also utilizes iterative linear algebra for the solution of its associated linear systems. In that case, both methods rely on a similar preconditioning strategy (we note that the preconditioner analyzed in (4.5) is a variant inspired from the preconditioner proposed in [24, Section 3.1], which is used in the benchmark inexact IP-PMM code). Finally, both methods are implemented in MATLAB. To ensure a fair comparison, we run both the exact and the inexact variants of IP-PMM and report the statistics of the best performing variant. To showcase the efficiency of both second-order

¹ https://github.com/spougkakiotis/Krylov_SSN_PMM

² https://github.com/spougkakiotis/IP_PMM

³ https://github.com/spougkakiotis/IP-PMM_QP_Solver

solvers (active-set and IP-PMM) against potential first-order alternatives, in the first application (i.e., on the portfolio optimization instances) we will also report the results obtained by applying the well-known ADMM-based OSQP solver given in [44], the implementation of which can also be found on GitHub⁴.

5.1 Portfolio Optimization

We first consider the mean-risk portfolio selection problem (originally proposed in [55]), where we minimize some convex risk measure, while keeping the expected return of the portfolio above some desirable level. A variety of models for this problem have been extensively analyzed and solved in the literature (e.g., see [6, 9, 56, 57]). The departure from the variance as a measure of risk often allows for great flexibility in the decision making of investors, enabling them to follow potential regulations as well as to better control the risk associated with an “optimal” portfolio. Optimality conditions and existence of solutions for several general deviation measures have been characterized in [58]. Additionally, the comparison of portfolios obtained by minimizing different risk measures has also been considered multiple times (e.g. see [10, 59–61]).

The method presented in this paper is general and not related to a specific model choice. Indeed, we would like to showcase the efficiency of our approach for obtaining accurate solutions to portfolio optimization problems with two distinct risk measures of practical interest. To that end, we focus on the solution of a standard portfolio selection model that has been used in the literature. All numerical results are obtained on real-world datasets. As a result the problems are of medium-scale. Nonetheless, even for such medium-scale problems, we will be able to demonstrate the efficiency of the proposed approach, when compared to a robust interior point method employed in the literature [3] for similar problems. We also note that both second-order solvers (i.e., IPM and active-set) significantly outperform OSQP on these instances, however the latter was included in the comparison for completeness. Some large-scale instances will be tackled in the following subsections.

We consider the case where a decision vector $x \in \mathbb{R}^n$ represents a portfolio of n financial instruments, such that

$$x_i \in [a_{li}, a_{ui}], \quad \text{with } a_{li} = 0, \quad a_{ui} \leq 1, \quad \text{for all } i = 1, \dots, n, \quad \text{and} \quad \sum_{i=1}^n x_i = 1.$$

This requirement indicates that no short positions are allowed, and no more than a_{ui} % of the total wealth can be invested in instrument i . Let $\xi \in \mathbb{R}^n$ denote a random vector, the i -th entry of which represents the random return of the i -th instrument. Then, the negative of the random return of a given portfolio x is given by $g(x, \xi) \triangleq -x^\top \xi$. In this paper we assume that ξ follows some continuous distribution $p(\xi)$, as well as that there is a one-to-one correspondence between percentage return and monetary value (as in [9, Section 3]). Additionally, given some expected benchmark return r (e.g., the *market index*), we only consider portfolios that yield an expected return above a certain threshold, i.e. $\mathbb{E}[-g(x, \xi)] \geq r$.

⁴ <https://github.com/osqp/osqp-matlab>

Table 1 Portfolio optimization datasets

Name	# of assets	# of data points	Timeline
DowJones [DJ]	28	1363	Feb. 1990–Apr. 2016
NASDAQ100 [NDQ]	82	596	Nov. 2004–Apr. 2016
FTSE100 [FT]	83	717	Jul. 2002–Apr. 2016
FF49Industries [FF]	49	2325	Jul. 1969–Jul. 2015
SP500 [SP]	442	595	Nov. 2004–Apr. 2016
NASDAQComp [NDC]	1203	685	Feb. 2003–Apr. 2016

Finally, given the previously stated constraints, we would like to minimize some convex risk measure $\varrho(\cdot)$ of interest. By putting everything together, the model reads as

$$\begin{aligned}
 & \min_x \varrho(g(x, \xi)), \\
 & \text{s.t. } \sum_{i=1}^n x_i = 1, \\
 & \quad \mathbb{E}[-g(x, \xi)] \geq r, \\
 & \quad x_i \in [a_{li}, a_{ui}], \quad i = 1, \dots, n.
 \end{aligned} \tag{5.1}$$

There are several methods for solving such stochastic problems; let us mention two important variants. There is the *parametric approach* (e.g., as in [6, 9]), where one assumes that the returns follow some known distribution which is subsequently sampled to yield finite-dimensional optimization problems, and the *sampling approach* (e.g., as in [56]), where one obtains a finite number of samples (without assuming a specific distribution). Such samples are often obtained by historical observations, and this approach is also followed in this paper. It is well-known that historical data cannot fully predict the future (see [56]), however it is a widely-used practice. The reader is referred to [62] for an extensive study on probabilistic models for portfolio selection problems, with practical considerations.

Additional soft or hard constraints can be included when solving a portfolio selection problem. Such constraints can either be incorporated directly via the use of a model (e.g., see [63, Section 2]) as hard constraints or by including appropriate penalty terms in the objective (soft constraints). It is important to note that the model given in (P) is quite general and as a result has great expressive power, allowing one to incorporate various important risk measures (and their combinations), as well as modeling constraints of interest.

Real-world datasets. In what follows, we solve two different instances of problem (5.1). In particular, we consider two risk measures; the conditional value at risk (e.g., see [9]), as well as the mean absolute semi-deviation (e.g., see [64, Section 6.2.2.], noting that this is in fact equivalent to the mean absolute deviation originally proposed in [65]).

We showcase the effectiveness of the proposed approach on 6 real datasets taken from [66]. Each dataset contains time series for weekly asset returns and market indexes for different major stock markets, namely, DowJones, NASDAQ100, FTSE100, SP500, NASDAQComp, and FF49Industries. In the first 5 markets the authors in [66] provide the market indexes, while for the last dataset the uniform allocation strategy is considered as a benchmark. Additional information on the datasets is collected in Table 1. We note that the authors in [66] disregard stocks with less than 10 years of observations.

5.1.1 Conditional Value at Risk

First, we consider portfolio optimization problems that seek a solution minimizing the conditional value at risk; a measure which is known to be *coherent* (see [67] for a definition of coherent risk measures). In particular using the notation introduced earlier, we consider the following optimization problem

$$\min_{x \in \mathbb{R}^n} \{ \text{CVaR}_\alpha (g(x, \xi)) + \delta_{\mathcal{K}}(x) \}, \quad \text{s.t. } Ax = b, \quad (5.2)$$

where $g(x, \xi)$ is the random cost function, $A \in \mathbb{R}^{m \times n}$ models the linear constraint matrix of problem (5.1) (where an auxiliary variable has been introduced to transform the inequality constraint involving the discretized expectation into an equality), and $\mathcal{K} \triangleq [a_l, a_u]$. In the above $1 - \alpha \in (0, 1)$ is the confidence level. It is well-known ([9]) that given a continuous random variable X , the conditional value at risk can be computed as an infimal convolution

$$\text{CVaR}_\alpha (X) \triangleq \min_{t \in \mathbb{R}} \{ t + \alpha^{-1} \mathbb{E} [(X - t)_+] \}.$$

Thus, we can write problem (5.2) in the following equivalent form:

$$\min_{(x, t) \in \mathbb{R}^n \times \mathbb{R}} \left\{ t + \frac{1}{l\alpha} \sum_{i=1}^l \left(-\xi_i^\top x - t \right)_+ + \delta_{\mathcal{K}}(x) \right\}, \quad \text{s.t. } Ax = b,$$

where all expectations have been substituted by summations since we assume the availability of a dataset $\{\xi_1, \dots, \xi_l\}$.

We solve problem (5.2) using the proposed active-set method (AS), the interior point-proximal method of multipliers (IP-PMM) given in [43], and OSQP ([44]). We set $\text{tol} = 10^{-5}$, and run the three methods for each of the datasets described in Table 1 for varying confidence level. We report the confidence parameter α , the number of PMM, SSN, and Krylov iterations required by the active-set scheme, the number IP-PMM and Krylov iterations required by the interior point scheme (with Krylov iterations reported only if the inexact IP-PMM variant performs better), and the number of OSQP iterations, the CPU time needed by each of the three schemes, as well as the number of factorizations used within the active set (PMM-SSN) scheme. As already mentioned earlier, the AS scheme does not utilize a preconditioner until it is deemed necessary. Additionally, it often happens that the active-set is not altered from one iteration to the next, and the factorization of the preconditioner does not need to be re-computed. The results are collected in Table 2. In all the numerical results that follow, the lowest running time exhibited by a solver, assuming it successfully converged, is presented in bold.

From Table 2 we observe that for the smaller instances (DowJones, NASDAQ100, FTSE100, FF49Industries) both second-order methods perform well and are comparable in terms of CPU time. Nonetheless, even in this case, the proposed active-set solver requires less computations per iteration. Indeed, this can be seen by the fact that the two methods achieve similar times but the active-set scheme is performing significantly more SSN iterations. On the other hand, for the larger instances (SP500, NASDAQComp) the proposed active-set scheme outperforms the interior point method significantly. This is mostly due to the efficiency gained from the lower memory requirements and cheaper factorizations of the preconditioner associated with the active-set solver. We additionally observe that both second-order methods are robust with respect to the confidence level and consistently outperform OSQP, which struggles to find a 5-digit accurate solution. We note that OSQP could potentially be competitive for smaller tolerances (e.g., for finding a 3-digit accurate

Table 2 CVaR portfolio selection: varying confidence level ($\tau=10^{-5}$)

Dataset	α	Iterations			Time (s)		
		PM(MSSN){Fact.} {Krylov}	IP-PM(M(Krylov)	OSQP	AS	IP-PM(M	OSQP
DJ	0.05	29(78) 15 {3841}	22(-) ¹	32,575	0.23	0.91	10.94
	0.10	33(88) 24 {4116}	22(-)	42,625	0.40	0.84	13.61
	0.15	36(103) 35 {3954}	19(-)	43,100	0.54	0.75	14.45
NDQ	0.05	26(80) 20 {3722}	17(-)	30,575	0.28	0.51	7.92
	0.10	29(104) 27 {4781}	18(-)	34,700	0.46	0.55	8.81
	0.15	32(100) 33 {4352}	17(-)	40,500	0.45	0.53	10.46
FE	0.05	30(108) 73 {2208}	21(-)	21,825	0.82	1.02	10.96
	0.10	31(101) 46 {3266}	16(-)	35,825	0.52	0.72	17.70
	0.15	34(99) 54 {2798}	18(-)	36,550	0.65	0.93	18.24
FF	0.05	38(118) 31 {3491}	27(-)	50,000 [‡] 2	0.60	3.80	59.16 [‡]
	0.10	37(132) 20 {2777}	23(-)	50,000 [‡]	0.76	3.21	57.92 [‡]
	0.15	38(108) 19 {4169}	20(-)	50,000 [‡]	0.88	2.82	51.25 [‡]
SP	0.05	29(117) 93 {2753}	24(699)	21,075	1.87	6.57	65.70
	0.10	32(119) 70 {2861}	27(877)	32,375	1.70	7.79	99.39
	0.15	31(106) 64 {3069}	23(-)	39,225	1.76	13.49	119.95
NDC	0.05	32(225) 154 {4538}	26(593)	16,100	6.89	23.65	143.82
	0.10	31(168) 123 {3441}	27(830)	31,275	6.91	27.54	291.97
	0.15	34(179) 124 {3538}	28(925)	16,100	8.65	31.43	403.75

Notes: – in the Krylov iterations for IP-PM(M indicates that the exact variant of IP-PM(M was used. ‡ indicates that the solver reached the maximum number of iterations

Table 3 MASD portfolio selection ($\tau_{01} = 10^{-5}$)

Dataset	Iterations			Time (s)		
	PMM(SSN)[Fact.]{Krylov}	IP-PMM(Krylov)	OSQP	AS	IP-PMM	OSQP
DJ	48(105)[31][2686]	14(–)	14,050	0.79	0.54	4.60
NDQ	41(119)[37][3436]	13(–)	40,550	0.44	0.28	10.72
FE	42(133)[53][3020]	12(–)	47,125	0.69	0.30	23.29
FF	52(100)[0][4289]	18(–)	12,050	0.77	2.70	14.24
SP	42(180)[65][4313]	19(–)	41,275	2.06	4.09	128.09
NDC	46(176)[75][4569]	14(–)	41,000	6.74	10.57	367.32

solution); however, the application under consideration dictates that an accurate solution is needed, since a small improvement in the portfolio output can translate into considerable profits in practice. Concerning IP-PMM, we note that the inexact variant is only elected for some of the largest instances (i.e., SP500 and NASDAQComp). This is expected, since these problems are small- to medium-scale. Nonetheless, large-scale instances will be tackled in the following subsections.

5.1.2 Mean Absolute Semi-Deviation

Next, we consider portfolio optimization problems that seek a solution minimizing the mean absolute semi-deviation, which is also coherent. We consider the following optimization problem

$$\min_{x \in \mathbb{R}^n} \{ \text{MASD}(g(x, \xi)) + \delta_{\mathcal{K}}(x) \}, \quad \text{s.t. } Ax = b, \quad (5.3)$$

where, given a continuous random variable X , the associated risk is defined as

$$\text{MASD}(X) \triangleq \mathbb{E}[(X - \mathbb{E}[X])_+] \equiv \frac{1}{2} \mathbb{E}|X - \mathbb{E}[X]|,$$

where the equivalence follows from [64, Proposition 6.1]. Given a dataset $\{\xi_1, \dots, \xi_l\}$, problem (5.3) can be written as

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \left\{ \frac{1}{l} \sum_{i=1}^l \left(-\xi_i^\top x + \mu^\top x \right)_+ + \delta_{\mathcal{K}}(x) \right\}, \\ \text{s.t. } Ax = b, \end{aligned}$$

where $\mu \triangleq \frac{1}{l} \sum_{i=1}^l \xi_i^\top x$. Note that this model is in the form of (P). We fix $\tau_{01} = 10^{-5}$ and run the three methods on the 6 datasets. The results are collected in Table 3.

From Table 3 we observe that both second-order schemes are, once again, robust and comparable for all instances. In this case, the larger instances (SP500, NASDAQComp) were solved in comparable time by both solvers. Nevertheless, it is important to note that the active-set scheme is the better choice, since it has significantly less memory requirements (this is due to solving a smaller problem formulation while also only utilizing information from the active sets when constructing its associated preconditioners). Notably the proposed method is robust (i.e., converges reliably to an accurate solution). Finally, as in the case of the CVaR instances, OSQP struggled to find accurate solutions at a reasonable number of

iterations, making it a less efficient choice for such problems. In subsequent experiments, we will drop OSQP from our benchmarks since it performed significantly worse compared to both second-order solvers.

5.2 PDE-Constrained Optimization

Let us now test the proposed methodology on some optimization problems with partial differential equation constraints. We consider optimal control problems of the following form:

$$\begin{aligned} \min_{y,u} \quad & \frac{1}{2} \|y - \bar{y}\|_{L^2(\Omega)}^2 + \frac{\alpha_1}{2} \|u\|_{L^1(\Omega)}^2 + \frac{\alpha_2}{2} \|u\|_{L^2(\Omega)}^2, \\ \text{s.t.} \quad & \text{Dy}(\mathbf{x}) + u(\mathbf{x}) = g(\mathbf{x}), \quad u_a(\mathbf{x}) \leq u(\mathbf{x}) \leq u_b(\mathbf{x}), \end{aligned} \quad (5.4)$$

where $(y, u) \in H^1(\Omega) \times L^2(\Omega)$, D is some linear differential operator, $\mathbf{x}$ is a 2-dimensional spatial variable, and $\alpha_1, \alpha_2 \geq 0$ are the regularization parameters of the control variable. The problem is considered on a given compact spatial domain Ω , where $\Omega \subset \mathbb{R}^2$ has boundary $\partial\Omega$, and is equipped with Dirichlet boundary conditions. The algebraic inequality constraints are assumed to hold a.e. on Ω , while u_a and u_b may take the form of constants or functions of the spatial variables.

We solve problem (5.4) via a discretize-then-optimize strategy. We employ the Q1 finite element discretization implemented in IFISS⁵ (see [68, 69]) which yields a sequence of ℓ_1 -regularized convex quadratic programming problems in the form of (P). We note that the discretization of the smooth parts of problem (5.4) follows a standard Galerkin approach (e.g., see [70]), while the L^1 term is discretized by the *nodal quadrature rule* as in [71] (an approximation that achieves a first-order convergence—see [7]).

In what follows, we consider two-dimensional L^1/L^2 -regularized Poisson optimal control problems. The problem is posed on $\Omega = (0, 1)^2$. Following [25, Section 5.1], we set constant control bounds $u_a = -2$, $u_b = 1.5$, and the desired state as $\bar{y} = \sin(\pi x_1) \sin(\pi x_2)$. In order to ensure that both solvers exhibit a consistent behaviour over PDE-constrained optimization problems of arbitrary size, we enforce stopping conditions with *absolute* tolerance $\text{tol} = 10^{-4}$. This is important in this case, since IP-PMM solves a smooth reformulation of the problem, and its relative tolerance criteria differ significantly from those of the active-set method when the discretized PDE size increases. In Table 4, we fix the L^2 regularization parameter to the value $\alpha_2 = 10^{-2}$, and present the runs of the two methods (i.e. active-set, IP-PMM) for varying L^1 regularization (i.e. α_1) as well as grid size. We report the size of the resulting discretized problems (before any reformulation), the value of α_1 , the number iterations required by each solver, as well as the total time to convergence.

From Table 4 we observe that both the active-set and IP-PMM are competitive, with the active-set method achieving faster convergence in terms of CPU time. Let us note that the number of Krylov iterations required to obtain an absolute 4-digit solution by both solvers is comparable. Thus, the active-set approach is more efficient, since it requires less memory (as it solves smaller and better conditioned linear systems, and thus the inversion of the associated preconditioners is less resource-consuming). On the other hand, we can observe that both AS and IP-PMM have a very consistent behaviour for a wide range of values of the associated problem parameters.

Next, we fix $\alpha_1 = 10^{-4}$ and the problem size, and run the two solvers on PDE-constrained optimization instances with varying L^2 regularization. The results are collected in Table 5,

⁵ <https://personalpages.manchester.ac.uk/staff/david.silvester/ifiss/default.htm>

Table 4 Poisson control: varying grid size and L^1 regularization ($\alpha_2 = 10^{-2}$)

n	α_1	Iterations		Time (s)	
		PMM(SSN)[Fact.]{Krylov}	IP-PMM(Krylov)	AS	IP-PMM
$8.45 \cdot 10^3$	10^{-2}	14(39)[28]{1691}	13(1096)	3.14	4.36
	10^{-4}	14(39)[23]{1687}	14(1253)	2.67	4.50
	10^{-6}	14(39)[23]{1686}	14(1256)	2.49	4.11
$3.32 \cdot 10^4$	10^{-2}	16(34)[28]{1295}	15(1147)	9.98	18.97
	10^{-4}	16(34)[24]{1291}	15(1379)	9.98	22.55
	10^{-6}	16(34)[24]{1291}	15(1381)	9.82	23.40
$1.32 \cdot 10^5$	10^{-2}	18(31)[29]{1149}	14(876)	43.98	67.10
	10^{-4}	18(31)[26]{1143}	15(1111)	42.56	83.88
	10^{-6}	18(31)[26]{1143}	15(1116)	42.69	82.23
$5.26 \cdot 10^5$	10^{-2}	18(32)[30]{1251}	15(1166)	246.08	450.25
	10^{-4}	18(32)[30]{1250}	15(1171)	218.24	743.47
	10^{-6}	18(32)[30]{1250}	15(1174)	259.12	728.32
$2.11 \cdot 10^6$	10^{-2}	21(38)[35]{1535}	15(909)	1199.34	2676.53
	10^{-4}	21(38)[36]{1534}	16(1068)	1342.57	1878.91
	10^{-6}	21(38)[36]{1534}	16(1068)	1452.42	1612.56

Table 5 Poisson control: varying L^2 regularization ($n, \alpha_1 = (1.32 \cdot 10^5, 10^{-4})$).

α_2	Iterations		Time (s)	
	PMM(SSN)[Fact.]{Krylov}	IP-PMM(Krylov)	AS	IP-PMM
10^{-1}	18(31)[26]{1155}	18(3979)	44.06	423.02
10^{-4}	18(31)[26]{1142}	15(661)	77.29	106.84
10^{-6}	18(31)[26]{1141}	15(607)	60.52	103.41
0	18(31)[26]{1144}	15(499)	70.45	73.72

‡ indicates that the solver reached the maximum number of iterations

and showcase the robustness of the proposed active-set solver with respect to the problem parameters. Indeed, one can observe that the active-set scheme exhibits more consistent convergence behaviour compared to IP-PMM as we alter the problem parameters.

At this point, we would like to stress the importance of utilizing preconditioning (and, additionally, the effectiveness of the proposed preconditioning approach). Specifically, if we run the proposed method on any of the problems tested in this section, while at the same time turning-off the preconditioning capabilities, then the scheme does not converge for any of these instances. The current implementation allows at most 150 Krylov iterations for each SSN iteration. The above statement holds true even if we increased the maximum Krylov iterations to 1000.

As we have already mentioned earlier, the implemented approach enables preconditioning only when the number of Krylov iterations reaches 100 for one SSN iteration, at which point preconditioning is turned on. Looking at the numerical results presented so far, only

one problem converged without utilizing preconditioning; that is, the FF instance in the MASD portfolio selection tests. Nonetheless, most quantile regression and SVM classification instances also enabled the method to converge without the utilization of a preconditioner (see Sections 5.3, 5.4). However, the numerical results of this section demonstrate that the proposed preconditioning strategy is integral in ensuring the efficiency and reliability of the proposed method.

5.3 Penalized Quantile Regression

Let us now consider linear regression models of the following form

$$y_i = \beta_0 + \xi_i^\top \beta + \epsilon_i, \quad i \in \{1, \dots, l\}$$

where ξ_i is a d -dimensional vector of covariates, (β_0, β) are the regression coefficients and ϵ_i is some random error. A very popular problem in statistics is the estimation of the optimal coefficients, in the sense of minimizing a model of the following form:

$$\min_{(\beta_0, \beta) \in \mathbb{R} \times \mathbb{R}^d} \left\{ \frac{1}{l} \sum_{i=1}^l \ell(y_i - \beta_0 - \xi_i^\top \beta) + \lambda p(\beta) \right\}, \quad (5.5)$$

where $\ell(\cdot)$ is some loss function and $p(\cdot)$ is a penalty function with an associated regularization parameter $\lambda \geq 0$. Following [5], we consider the elastic-net penalty,

$$p(\beta) \equiv p_\tau(\beta) \triangleq \tau \|\beta\|_1 + \frac{1-\tau}{2} \|\beta\|_2^2, \quad 0 \leq \tau \leq 1.$$

For the loss function, we employ the quantile loss

$$\ell(w) \equiv \rho_\alpha(w) \triangleq (1-\alpha)w_- + \alpha w_+ = \frac{1}{2}(|w| + (2\alpha - 1)w), \quad 0 < \alpha < 1, \quad (5.6)$$

where $w \in \mathbb{R}$. Notice that the case $\alpha = \frac{1}{2}$ yields the absolute loss. Letting $x = [\beta_0 \ \beta^\top]^\top$, and using Remark 1, we can re-write problem (5.5) in the form of (P), as

$$\min_{x \in \mathbb{R}^{1+d}} \left\{ (\alpha - 1) \mathbf{1}_l^\top (Cx + d) + \frac{1}{2} x^\top Qx + \sum_{i=1}^l ((Cx + d)_i)_+ + \|Dx\|_1 \right\},$$

where

$$C_{:,i} = -\frac{1}{l} [1 \ \xi_i^\top], \quad d_i = \frac{1}{l} y_i, \quad i = 1, \dots, l, \quad Q = \begin{bmatrix} 0 & 0_{1,d} \\ 0_{d,1} & \lambda(1-\tau)I_d \end{bmatrix}, \quad D = \begin{bmatrix} 0 & 0_{1,d} \\ 0_{d,1} & \lambda\tau I_d \end{bmatrix}.$$

Real-world datasets. In what follows, we solve several instances of problem (5.5). We showcase the effectiveness of the proposed approach on 5 regression problems taken from the LIBSVM library (see [72]). Additional information on the datasets is collected in Table 6.

We fix $\text{tol} = 10^{-4}$, $\lambda = 10^{-2}$, and $\tau = 0.5$, and run the two methods (active-set (AS), IP-PMM) on the 5 instances for varying quantile level α . The results are collected in Table 7.

From Table 7 we observe that the two approaches are comparable for the smaller instances (space_ga, abalone, and cpusmall). The active-set scheme significantly outperforms IP-PMM on the cadata problem, mainly due to its superior numerical stability and better conditioning of its linear systems. Indeed, this problem is highly ill-conditioned; this is reflected in the increased number of IP-PMM iterations needed to obtain a 4-digit accurate solution. The

Table 6 Quantile regression datasets

Name	# of data points	# of features
space_ga	3107	6
abalone	4177	8
cpusmall	8192	12
cadata	20,640	8
E2006-tfidf	16,087	150,360

Table 7 Quantile regression: varying confidence level ($\epsilon_{\text{tol}} = 10^{-4}$, $\lambda = 10^{-2}$, $\tau = 0.5$)

Problem	α	Iterations		Time (s)	
		PMM(SSN)[Fact.]{Krylov}	IP-PMM(Krylov)	AS	IP-PMM
space_ga	0.50	32(36)[0]{623}	18(–)	0.15	0.31
	0.65	32(37)[0]{636}	20(–)	0.13	0.38
	0.80	30(34)[0]{583}	16(–)	0.12	0.25
	0.90	29(31)[0]{536}	22(–)	0.11	0.35
abalone	0.50	27(35)[0]{603}	34(–)	0.17	1.39
	0.65	27(31)[0]{534}	31(–)	0.15	1.13
	0.80	25(30)[0]{492}	22(–)	0.12	0.85
	0.90	24(24)[0]{367}	20(–)	0.09	0.74
cpusmall	0.50	28(32)[0]{538}	25(–)	0.26	2.02
	0.65	28(33)[0]{577}	25(–)	0.29	2.01
	0.80	29(30)[0]{518}	22(–)	0.24	1.81
	0.90	29(31)[0]{487}	25(–)	0.23	2.13
cadata	0.50	4(46)[0]{860}	49(–)	0.27	11.40
	0.65	5(49)[0]{933}	42(–)	0.31	10.51
	0.80	5(50)[0]{861}	45(–)	0.30	10.60
	0.90	29(1405)[214]{5722}	$\overline{\wedge}^1$	8.95	$\overline{\wedge}$
E2006-tfidf	0.50	37(37)[0]{363}	$\dagger^2$	60.91	$\dagger$
	0.65	36(39)[0]{405}	$\dagger$	111.66	$\dagger$
	0.80	36(39)[0]{399}	$\dagger$	118.28	$\dagger$
	0.95	35(37)[0]{388}	$\dagger$	103.19	$\dagger$

¹ $\overline{\wedge}$ indicates that the solver incorrectly identified the problem as infeasible. ² $\dagger$ indicates that the solver ran out of memory

better conditioning of the linear systems associated with the active-set (compared to those appearing within IP-PMM) can also be observed from the fact that it did not utilize the preconditioner in most instances, enabling it to successfully converge matrix-free. Finally, we observe that for the large-scale instance (E2006-tfidf), both the exact and the inexact variants of IP-PMM crashed due to memory requirements.

Next, we fix $\epsilon_{\text{tol}} = 10^{-4}$, and $\alpha = 0.8$, and run the two methods for varying regularization parameters τ and λ . The results are collected in Table 8. We observe that both methods are robust for a wide range of parameter values. As before, the active-set scheme consistently

solved the large-scale instance (E2006-tfidf) for a wide range of parameter values, without running into any memory issues.

5.4 Binary Classification Via Linear Support Vector Machines

Finally, we are interested in training a binary linear classifier using regularized soft-margin linear support vector machines (SVMs), [73]. More specifically, given a training dataset $\{(y_i, \xi_i)\}_{i=1}^l$, where $y_i \in \{-1, 1\}$ are the *labels* and $\xi_i \in \mathbb{R}^d$ are the *feature vectors* (with d the number of features), we would like to solve the following optimization problem

$$\min_{(\beta_0, \beta) \in \mathbb{R} \times \mathbb{R}^d} \left\{ \frac{1}{l} \sum_{i=1}^l \left(1 - y_i \left(\xi_i^\top \beta - \beta_0 \right) \right)_+ + \lambda \left(\tau_1 \|\beta\|_1 + \frac{\tau_2}{2} \|\beta\|_2^2 \right) \right\}, \quad (5.7)$$

where $\lambda > 0$ is a regularization parameter, and $\tau_1, \tau_2 > 0$ are the weights for the ℓ_1 and ℓ_2 regularizers, respectively. The standard Euclidean-norm regularization is traditionally used as a trade-off between the margin of the classifier (the larger the better) and the correct classification of ξ_i , for all $i \in \{1, \dots, l\}$ (e.g., see [74]). However, this often leads to a dense estimate for β , which has led researchers in the machine learning community into considering the ℓ_1 regularizer instead, in order to encourage sparsity (e.g., see [75]). It is well-known that both regularizers can be combined to deliver the benefits of each of them, using the elastic-net penalty (see, for example, [76]), assuming that τ_1 and τ_2 are appropriately tuned.

Real-world datasets. In what follows we consider elastic-net SVM instances of the form of (5.7). We showcase the effectiveness of the proposed active-set scheme on 3 large-scale binary classification datasets taken from the LIBSVM library ([72]). The problem names, as well as the number of features and training points are collected in Table 9.

In the experiments to follow, we only consider large-scale problems that neither the exact nor the inexact version of IP-PMM can solve on a personal computer, due to excessive memory requirements. Nonetheless, we should note that, by following the developments in [3], IP-PMM can be specialized to problems of this form in order to be able to tackle such large-scale instances. However, the same applies to the proposed active-set method, and thus specialization is avoided for the sake of generality. We fix $\epsilon_{\text{tol}} = 10^{-5}$, $\lambda = 10^{-2}$, and run the active-set solver for all datasets given in Table 9 for varying values of the regularization parameters τ_1 , and τ_2 . The results are collected in Table 10.

We observe that the solver was consistently able to solve these instances, without running into memory issues. In this case, the behaviour of the active-set solver was affected by the parameters τ_1 and τ_2 (indeed, see the number of SSN iterations for different regularization values), however, we consistently obtain convergence in a very reasonable amount of time. We observe that no preconditioning was required for most of these problems, except from one instance originating from the news20 dataset (with $\tau_1 = 0.2$ and $\tau_2 = 0.8$). On the one hand, even when the preconditioner was required, the method did not run into any memory issues. On the other hand, we observe that the associated linear systems of the proposed method are relatively well-conditioned, allowing us to avoid the inversion of the preconditioner in many instances, which enables its matrix-free utilization in several settings.

Overall, we observe that the proposed algorithm is very general and can be applied in a plethora of very important applications arising in practice. We were able to showcase that the active-set nature of the method allows one to solve large-scale instances on a personal computer. The proposed method strikes a good balance between first-order methods, which are fast but unreliable, and second-order interior point methods, which are extremely robust but can struggle with the problem size, ill-conditioning, and memory requirements. We have

Table 8 Quantile regression: varying regularization ($\tau_01 = 10^{-4}, \alpha = 0.8$)

Problem	(τ, λ)	Iterations		Time (s)	
		PMM(SSN)[Fact.][Krylov}	IP-PMM(Krylov)	AS	IP-PMM
space_ga	$(0.2, 5 \cdot 10^{-2})$	30(36)[0][585}	26(-)	0.13	0.40
	$(0.4, 1 \cdot 10^{-2})$	30(33)[0][568}	20(-)	0.12	0.32
	$(0.6, 5 \cdot 10^{-3})$	30(33)[0][654}	18(-)	0.14	0.29
	$(0.8, 1 \cdot 10^{-3})$	30(32)[0][695}	22(-)	0.12	0.37
abalone	$(0.2, 5 \cdot 10^{-2})$	26(31)[0][422}	20(-)	0.12	0.76
	$(0.4, 1 \cdot 10^{-2})$	25(30)[0][485}	27(-)	0.14	0.90
	$(0.6, 5 \cdot 10^{-3})$	26(26)[0][474}	34(-)	0.12	1.27
	$(0.8, 1 \cdot 10^{-3})$	26(26)[0][622}	25(-)	0.13	0.74
cpusmall	$(0.2, 5 \cdot 10^{-2})$	28(35)[0][440}	19(-)	0.26	1.58
	$(0.4, 1 \cdot 10^{-2})$	29(29)[0][487}	17(-)	0.24	1.37
	$(0.6, 5 \cdot 10^{-3})$	29(30)[0][652}	17(-)	0.28	1.47
	$(0.8, 1 \cdot 10^{-3})$	29(29)[0][871}	19(-)	0.30	1.54
cadata	$(0.2, 5 \cdot 10^{-2})$	34(940)[265][4436}	50(-)	5.93	11.34
	$(0.4, 1 \cdot 10^{-2})$	6(56)[0][928}	50(-)	0.34	11.55
	$(0.6, 5 \cdot 10^{-3})$	6(69)[0][1274}	45(-)	0.40	10.18
	$(0.8, 1 \cdot 10^{-3})$	12(187)[0][3621}	33(-)	1.05	7.34
E2006-tfidf	$(0.2, 5 \cdot 10^{-2})$	34(43)[0][406}	$\dagger$ 1	61.64	$\dagger$
	$(0.4, 1 \cdot 10^{-2})$	36(39)[0][395}	$\dagger$	57.64	$\dagger$
	$(0.6, 5 \cdot 10^{-3})$	36(40)[0][430}	$\dagger$	60.10	$\dagger$
	$(0.8, 1 \cdot 10^{-3})$	35(36)[0][447}	$\dagger$	84.41	$\dagger$

$\dagger$ indicates that the solver ran out of memory

Table 9 Binary classification datasets

Name	# of training points	# of features
rcv1	20,242	47,236
real-sim	72,309	20,958
news20	19,996	1,355,191

Table 10 Elastic-net SVMs: varying regularization ($\tau_0 = 10^{-5}$, $\lambda = 10^{-2}$)

Problem	τ_1	τ_2	Iterations PMM(SSN)[Fact.]{Krylov}	Time (s) AS
rcv1	0.2	0.2	46(64)[0]{1497}	9.28
	0.8	0.2	42(50)[0]{642}	6.90
	0.2	0.8	44(68)[0]{1228}	8.98
	5	5	30(20)[0]{138}	3.74
real-sim	0.2	0.2	57(88)[0]{2382}	34.99
	0.8	0.2	38(22)[0]{180}	10.42
	0.2	0.8	55(84)[0]{1740}	31.21
	5	5	37(26)[0]{168}	10.66
news20	0.2	0.2	36(43)[0]{629}	70.22
	0.8	0.2	29(26)[0]{206}	40.29
	0.2	0.8	31(110)[31]{857}	142.18
	5	5	33(26)[0]{154}	39.94

demonstrated that convex quadratic problems with piecewise-linear terms can be solved very efficiently using the proposed active-set scheme, and we conjecture that the algorithm can be readily extended to deal with significantly more general problems, either by appropriately extending its associated theory, or by adjusting its versatile numerical linear algebra kernel.

6 Conclusions

We presented an active-set method for the solution of convex quadratic optimization problems with piecewise-linear terms in the objective, with applications to sparse approximations and risk minimization. After showing the convergence of the proposed algorithm under minimal assumptions, we derived and analyzed a robust numerical linear algebra kernel for the solution of its associated linear systems. We derive preconditioners suitable for Krylov subspace methods that are robust with respect to the penalty parameters of our method and can, in general, pave the way for the widespread utilization of iterative linear algebra in the context of active-set schemes. We were able to numerically demonstrate the efficiency, scalability and robustness of the proposed approach over a wide range of instances arising from risk-averse portfolio selection, L^1 -regularized partial differential equation constrained optimization, quantile regression and linear classification via support vector machines, showcasing its competitiveness against alternative state-of-the-art methods. We have demonstrated the generality and versatility of the proposed solver and its associated linear algebra kernel,

which make it a very competitive alternative to first- or second-order schemes appearing in the literature.

A Appendix

A.1 Derivation of the Dual Problem

Assume that notation introduced in Section 2.1. Let $y = [y_1^\top \ y_2^\top]^\top \in \mathbb{R}^{m+l}$. The dual of (P) is $\sup_{y,z} \inf_x \{\ell(x, y, z)\}$. We have

$$\begin{aligned} \inf_x \{\ell(x, y, z)\} &= -\sup_x \left\{ \left(F^\top y - z \right)^\top x - f(x) \right\} + y^\top \hat{b} - \delta_{\mathcal{K}}^*(z) - h^*(y_2) \\ &= -f^*(F^\top y - z) + y^\top \hat{b} - \delta_{\mathcal{K}}^*(z) - h^*(y_2) \\ &= -\frac{1}{2} x^\top Qx - \delta_{\{0\}} \left(c + Qx - F^\top y + z \right) + y^\top \hat{b} - \delta_{\mathcal{K}}^*(z) - h^*(y_2), \end{aligned}$$

where we used the definition of the convex conjugate and the definition of $f(\cdot)$, which yields $f^*(F^\top y - z) = \frac{1}{2} x^\top Qx + \delta_{\{0\}}(c + Qx - F^\top y + z)$. In turn, we observe that this is exactly the dual problem (D).

A.2 Derivation of the Augmented Lagrangian

Let $\beta > 0$ be some penalty parameter. The augmented Lagrangian corresponding to (P) reads:

$$\begin{aligned} \mathcal{L}_\beta(x; y, z) &\triangleq \sup_{u', v'} \left\{ \ell(x, u', v') - \frac{1}{2\beta} \|u' - y\|^2 - \frac{1}{2\beta} \|v' - z\|^2 \right\} \\ &= f(x) - \inf_{u'} \left\{ u'^\top (Fx - \hat{b}) + h^*(u_{m+1:m+l}) + \frac{1}{2\beta} \|u' - y\|^2 \right\} \\ &\quad - \inf_{v'} \left\{ -v'^\top x + \delta_{\mathcal{K}}^*(v') + \frac{1}{2\beta} \|v' - z\|^2 \right\} \\ &= f(x) - y_1^\top (Ax - b) + \frac{\beta}{2} \|Ax - b\|^2 \\ &\quad - \inf_{w'} \left\{ -w'^\top (Cx + d) + \delta_{\mathcal{C}}(w') + \frac{1}{2\beta} \|w' - y_2\|^2 \right\} \\ &\quad - \inf_{v'} \left\{ -v'^\top x + \delta_{\mathcal{K}}^*(v') + \frac{1}{2\beta} \|v' - z\|^2 \right\}. \end{aligned}$$

where $\mathcal{C} \triangleq \{w \in \mathbb{R}^l \mid w_i \in [0, 1], \ j = 1, \dots, l\}$. Let us now solve each of the two minimization problems appearing in the last expression. For the first expression, we use the fact that if $p_1(x) = p_2(x) + r^\top x$, where $p_1, \ p_2$ are two proper closed convex functions, then $\text{prox}_{p_1}(x) = \text{prox}_{p_2}(x - r)$. Thus, we obtain

$$\arg \min_{w'} \left\{ -w'^\top (Cx + d) + \delta_{\mathcal{C}}(w') + \frac{1}{2\beta} \|w' - y_2\|^2 \right\} = \Pi_{\mathcal{C}}(y_2 + \beta(Cx + d)).$$

For the second minimization problem, we proceed as before to obtain that

$$\arg \min_{w'} \left\{ -v'^{\top} x + \delta_{\mathcal{K}}^*(v') + \frac{1}{2\beta} \|v' - z\|^2 \right\} = \mathbf{prox}_{\beta \delta_{\mathcal{K}}^*}(z + \beta x).$$

Then, by using the *Moreau Identity* (see [77]):

$$\mathbf{prox}_{\beta p}(u') + \beta \mathbf{prox}_{\beta^{-1} p^*}(\beta^{-1} u') = u', \quad (\text{A.1})$$

and by performing some manipulations, we arrive at the final expression for the augmented Lagrangian given in (2.1).

A.3 Differentiability of the Augmented Lagrangian

In order to find the gradient of the augmented Lagrangian, we need to use the gradient expression of the Moreau envelope of a given proper closed convex function p . Specifically, given some $\beta > 0$, the Moreau envelope of p reads

$$e_{\beta p}(x) = \inf_w \left\{ p(w) + \frac{1}{2\beta} \|w - x\|^2 \right\},$$

and it can be shown that $\nabla e_{\beta p}(x) = \frac{1}{\beta} (x - \mathbf{prox}_{\beta p}(x))$. We proceed by computing the gradient of

$$\inf_{w'} \left\{ -w'^{\top} (Cx + d) + \delta_{\mathcal{C}}(w') + \frac{1}{2\beta} \|w' - y_2\|^2 \right\},$$

noting that the gradient of the second Moreau envelope (implicitly) appearing in the augmented Lagrangian can be computed similarly (using, additionally, the Moreau identity). We notice that

$$\begin{aligned} & \inf_{w'} \left\{ -w'^{\top} (Cx + d) + \delta_{\mathcal{C}}(w') + \frac{1}{2\beta} \|w' - y_2\|^2 \right\} \\ &= \inf_{w'} \left\{ \delta_{\mathcal{C}}(w') + \frac{1}{2\beta} \|w' - (y_2 + \beta(Cx + d))\|^2 \right\} - \frac{\beta}{2} \|Cx + d\|^2 - y_2^{\top} (Cx + d). \end{aligned}$$

Thus, using the expression of the gradient of the second (shifted) Moreau envelope alongside the chain rule, we obtain that

$$\begin{aligned} & \nabla_x \left(\inf_{w'} \left\{ \delta_{\mathcal{C}}(w') + \frac{1}{2\beta} \|w' - (y_2 + \beta(Cx + d))\|^2 \right\} - \frac{\beta}{2} \|Cx + d\|^2 - y_2^{\top} (Cx + d) \right) \\ &= -C^{\top} \Pi_{\mathcal{C}}(y_2 + \beta(Cx + d)). \end{aligned}$$

Proceeding in a similar manner for the second Moreau envelope, we arrive at the expression for the gradient of the proximal augmented Lagrangian penalty function given in Section 2.1.

A.4 Termination criteria

Let $\mathcal{C} = \{x \in \mathbb{R}^l \mid x_i \in [0, 1], i \in \{1, \dots, l\}\}$. We write the optimality conditions for (P)–(D) as

$$\begin{aligned} c + Qx - A^{\top} y_1 + C^{\top} y_2 + z &= 0_n, & Ax &= b, \\ y_2 &= \Pi_{\mathcal{C}}(y_2 + Cx + d), & x &= \Pi_{\mathcal{K}}(x + z), \end{aligned} \quad (\text{A.2})$$

and the termination criteria for Algorithm PD-PMM (given a tolerance $\epsilon > 0$) are set as

$$\max \left\{ \frac{\|c + Qx - A^\top y_1 + C^\top y_2 + z\|}{1 + \|c\|}, \frac{\|y_2 - \Pi_C(y_2 + Cx + d)\|}{1 + \left\| \frac{Ax - b}{d} \right\|}, \|x - \Pi_K(x + z)\| \right\} \leq \epsilon. \quad (\text{A.3})$$

Funding Dionysis Kalogieras is supported by the US National Science Foundation (NSF) under Grant 2242215.

Data Availability The portfolio optimization data can be accessed in <https://www.sciencedirect.com/science/article/pii/S2352340916303997>. The code generating the PDE-constrained optimization data can be accessed in <https://personalpages.manchester.ac.uk/staff/david.silvester/ifiss/default.htm>. The regression and classification data can be accessed in <https://www.csie.ntu.edu.tw/~protect/unhbox\voidb@x\penalty\@M\cjin/libsvmtools/datasets/>. The MATLAB code can be accessed in GitHub (https://github.com/spougkakiotis/Krylov_SSN_PMM). The numerical results can be reproduced by following the instructions given in GitHub, which include specific instructions on downloading and storing the publicly available datasets used within this work.

Declarations

Conflict of Interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Bertsekas, D.P., Nedic, A., Ozdaglar, E.: Convex Analysis and Optimization. Athena Scientific (2003)
2. Chen, S.S., Donoho, D.L., Saunders, M.A.: Atomic decomposition by basis pursuit. *SIAM Review* **43**(1), 129–159 (2001). <https://doi.org/10.1137/S003614450037906X>
3. De Simone, V., di Serafino, D., Gondzio, J., Pougkakiotis, S., Viola, M.: Sparse approximations with interior point methods. *SIAM Review* **64**(4), 954–988 (2022). <https://doi.org/10.1137/21M1401103>
4. Vapnik, V.N.: Statistical Learning Theory. John Wiley & Sons, New York (1998)
5. Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **67**(2), 301–320 (2005). <https://doi.org/10.1111/j.1467-9868.2005.00503.x>
6. Alexander, S., Coleman, T.F., Li, Y.: Minimizing CVaR and VaR for a portfolio of derivatives. *Journal of Banking & Finance* **30**(2), 583–605 (2006). <https://doi.org/10.1016/j.jbankfin.2005.04.012>
7. Wachsmuth, G., Wachsmuth, D.: Convergence and regularization results for optimal control problems with sparsity functional. *ESAIM: Control, Optimisation and Calculus of Variations* **17**(3), 858–886 (2011). <https://doi.org/10.1051/cocv/2010027>
8. Li, D., Ng, W.L.: Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical Finance* **10**(3), 387–406 (2000). <https://doi.org/10.1111/1467-9965.00100>
9. Rockafellar, R.T., Uryasev, S.: Optimization of conditional value-at-risk. *Journal of Risk* **2**(3), 21–41 (2000). <https://doi.org/10.21314/JOR.2000.038>
10. Silva, L.P., Alem, D., Carvalho, F.L.: Portfolio optimization using mean absolute deviation (MAD) and conditional value-at-risk (CVaR). *Production* **27** (2017). <https://doi.org/10.1590/0103-6513.208816>
11. Burke, J.V., Ferris, M.C.: A Gauss-Newton method for convex composite optimization. *Mathematical Programming* **71**, 179–194 (1995). <https://doi.org/10.1007/BF01585997>

12. Lazimy, R.: Mixed-integer quadratic programming. *Mathematical Programming* **22**, 332–349 (1982). <https://doi.org/10.1007/BF01581047>
13. Beck, A.: First-Order Methods in Optimization. MOS-SIAM Series on Optimization. SIAM & Mathematical Optimization Society, Philadelphia (2017). <https://doi.org/10.1137/1.9781611974997>
14. Deng, W., Yin, W.: On the global and linear convergence of the generalized alternating direction method of multipliers. *Journal of Scientific Computing* **66**, 889–916 (2016). <https://doi.org/10.1007/s10915-015-0048X>
15. Han, S.P., Pang, J.S., Rangaraj, N.: Globally convergent Newton methods for nonsmooth equations. *Mathematics of Operations Research* **17**(3), 586–607 (1992). <https://doi.org/10.1287/moor.17.3.586>
16. Martínez, J., Qi, L.: Inexact Newton methods for solving nonsmooth equations. *Journal of Computational and Applied Mathematics* **60**(1–2), 127–145 (1995). [https://doi.org/10.1016/0377-0427\(94\)00088-1](https://doi.org/10.1016/0377-0427(94)00088-1)
17. Stella, L., Themelis, A., Patrinos, P.: Forward-backward quasi-Newton methods for nonsmooth optimization problems. *Computational Optimization and Applications* **67**, 443–487 (2017). <https://doi.org/10.1007/s10589-017-9912-y>
18. De Marchi, A.: On a primal-dual Newton proximal method for convex quadratic programs. *Computational Optimization and Applications* (2022). <https://doi.org/10.1007/s10589-021-00342-y>
19. Dhingra, N.K., Khong, S.Z., Jovanović, M.R.: A second order primal-dual algorithm for nonsmooth convex composite optimization. In: 2017 IEEE 56th Annual Conference on Decision and Control (CDC), pp. 2868–2873 (2017). <https://doi.org/10.1109/CDC.2017.8264075>
20. Hermans, B., Themelis, A., Patrinos, P.: QPALM: a proximal augmented Lagrangian method for nonconvex quadratic programs. *Mathematical Programming Computation* **14**, 497–541 (2022). <https://doi.org/10.1007/s12532-022-00218-0>
21. Lee, J.D., Sun, Y., Saunders, M.A.: Proximal Newton-type methods for minimizing composite functions. *SIAM Journal on Optimization* **24**(3), 1420–1443 (2014). <https://doi.org/10.1137/130921428>
22. Li, X., Sun, D., Toh, K.C.: An asymptotically superlinearly convergent semismooth Newton augmented Lagrangian method for linear programming. *SIAM Journal on Optimization* **30**(3), 2410–2440 (2020). <https://doi.org/10.1137/19M1251795>
23. Fountoulakis, K., Gondzio, J., Zhlobich, P.: Matrix-free interior point method for compressed sensing problems. *Mathematical Programming Computation* **6**, 1–31 (2014). <https://doi.org/10.1007/s12532-013-0063-6>
24. Gondzio, J., Pougkakiotis, S., Pearson, J.W.: General-purpose preconditioning for regularized interior point methods. *Computational Optimization and Applications* **83**, 727–757 (2022). <https://doi.org/10.1007/s10589-022-00424-5>
25. Pearson, J.W., Porcelli, M., Stoll, M.: Interior-point methods and preconditioning for PDE-constrained optimization problems involving sparsity terms. *Numerical Linear Algebra with Applications* **27**(2), e2276 (2019). <https://doi.org/10.1002/nla.2276>
26. Chen, J., Qi, L.: Globally and superlinearly convergent inexact Newton-Krylov algorithms for solving nonsmooth equations. *Numerical Linear Algebra with Applications* **17**, 155–174 (2010). <https://doi.org/10.1002/nla.673>
27. Gill, P.E., Robinson, D.P.: A primal-dual augmented Lagrangian. *Computational Optimization and Applications* **15**, 1–25 (2012). <https://doi.org/10.1007/s10589-010-9339-1>
28. Ito, K., Kunisch, K.: On a semi-smooth Newton method and its globalization. *Mathematical Programming* **118**, 347–370 (2009). <https://doi.org/10.1007/s10107-007-0196-3>
29. Hans, E., Raasch, T.: Global convergence of damped semismooth Newton methods for ℓ_1 Tikhonov regularization. *Inverse Problems* **31**, 025,005 (2015). <https://doi.org/10.1088/0266-5611/31/2/025005>
30. Li, X., Sun, D., Toh, K.C.: A highly efficient semismooth Newton augmented Lagrangian method for solving Lasso problems. *SIAM Journal on Optimization* **28**(1), 433–458 (2018). <https://doi.org/10.1137/16M1097572>
31. Mannel, F., Rund, A.: A hybrid semismooth quasi-Newton method for nonsmooth optimal control with PDEs. *Optimization and Engineering* **22**, 2087–2125 (2021). <https://doi.org/10.1007/s11081-020-09523-w>
32. Porcelli, M., Simoncini, V., Stoll, M.: Preconditioning PDE-constrained optimization with l^1 -sparsity and control constraints. *Computers & Mathematics with Applications* **74**(5), 1059–1075 (2017). <https://doi.org/10.1016/j.camwa.2017.04.033>
33. Christof, C., De Los Reyes, H.C., Meyer, C.: A nonsmooth trust-region method for locally Lipschitz functions with applications to optimization problems constrained by variational inequalities. *SIAM Journal on Optimization* **30**, 2163–2196 (2020). <https://doi.org/10.1137/18M1164925>
34. Themelis, A., Stella, L., Patrinos, P.: Forward-backward envelope for the sum of two nonconvex functions: Further properties and nonmonotone linesearch algorithms. *SIAM Journal on Optimization* **28**(3), 2274–2303 (2018). <https://doi.org/10.1137/16M1080240>

35. Rockafellar, R.T.: Augmented Lagrangians and applications of the proximal point algorithm in convex programming. *Mathematics of Operations Research* **1**(2), 97–116 (1976). <https://doi.org/10.1287/moor.1.2.97>
36. Rockafellar, R.T.: Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization* **14**(5), 877–898 (1976). <https://doi.org/10.1137/0314056>
37. Bergamaschi, L., Gondzio, J., Martínez, A., Pearson, J.W., Pougkakiotis, S.: A new preconditioning approach for an interior point-proximal method of multipliers for linear and convex quadratic programming. *Numerical Linear Algebra with Applications* **28**(4), e2361 (2021). <https://doi.org/10.1002/nla.2361>
38. Waltz, R.A., Morales, J.L., Nocedal, J., Orban, D.: An interior algorithm for nonlinear optimization that combines line search and trust region steps. *Mathematical Programming* **107**(3, Ser. A), 391–408 (2006). <https://doi.org/10.1007/s10107-004-0560-5>
39. Zhao, X.Y., Sun, D.F., Toh, K.C.: A Newton-CG augmented Lagrangian method for semidefinite programming. *SIAM Journal on Optimization* **20**(4), 1737–1765 (2010). <https://doi.org/10.1137/080718206>
40. Wang, C., Tang, P.: A dual semismooth Newton based augmented Lagrangian method for large-scale linearly constrained sparse group square-root Lasso problems. *Journal of Scientific Computing* **96** (2023). <https://doi.org/10.1007/s10915-023-02271-w>
41. Lin, M., Sun, D., Toh, K.C., Wang, C.: Estimation of sparse Gaussian graphical models with hidden clustering structure. *Journal of Machine Learning Research* **25**(256), 1–36 (2024). <http://jmlr.org/papers/v25/20-354.html>
42. Patrinos, P., Bemporad, A.: Proximal Newton methods for convex composite optimization. In: 52nd IEEE Conference on Decision and Control, pp. 2358–2363 (2013). <https://doi.org/10.1109/CDC.2013.6760233>
43. Pougkakiotis, S., Gondzio, J.: An interior point-proximal method of multipliers for convex quadratic programming. *Computational Optimization and Applications* **78**, 307–351 (2021). <https://doi.org/10.1007/s10589-020-00240-9>
44. Stellato, B., Banjac, G., Goulart, P., Bemporad, A., Boyd, S.: OSQP: an operator splitting solver for quadratic programs. *Mathematical Programming Computation* **12**(4), 637–672 (2020). <https://doi.org/10.1007/s12532-020-00179-2>
45. Pougkakiotis, S., Gondzio, J., Kalogieras, D.: An active-set method for sparse approximations. Part I: Separable ℓ_1 terms. [arXiv:2201.10211v2](https://arxiv.org/abs/2201.10211v2) (2023)
46. Pougkakiotis, S., Gondzio, J., Kalogieras, D.: An active-set method for sparse approximations. Part II: General piecewise-linear terms. [arXiv:2302.14497v1](https://arxiv.org/abs/2302.14497v1) (2023)
47. Rockafellar, R.T., Wets, R.J.B.: Variational Analysis, *Grundlehren der mathematischen Wissenschaften*, vol. 317. Springer-Verlag Berlin Heidelberg (1998). <https://doi.org/10.1007/978-3-642-02431-3>
48. Rockafellar, R.T.: Generalizations of the proximal method of multipliers in convex optimization. *Computational Optimization and Applications* **87**, 219–247 (2024)
49. Robinson, S.M.: Some continuity properties of polyhedral multifunctions. In: H. König, B. Korte, K. Ritter (eds.) *Mathematical Programming at Oberwolfach, Mathematical Programming Studies*, vol. 14, pp. 206–214. Springer, Berlin, Heidelberg (1981). <https://doi.org/10.1007/BFb0120929>
50. Dontchev, A.L., Rockafellar, R.T.: Implicit Functions and Solution Mappings. Springer Series in Operations Research and Financial Engineering. Springer, New York, NY (2014). <https://doi.org/10.1007/978-1-4939-1037-3>
51. Clarke, F.: Optimization and Nonsmooth Analysis. Classics in Applied Mathematics. John Wiley and Sons, New York (1990). <https://doi.org/10.1137/1.9781611971309>
52. Clason, C., Valkonen, T.: Introduction to Nonsmooth Analysis and Optimization. [arXiv:2001.00216v4](https://arxiv.org/abs/2001.00216v4) (2023)
53. Vanderbei, R.J.: Symmetric quasidefinite matrices. *SIAM Journal on Optimization* **5**(1), 100–113 (1993). <https://doi.org/10.1137/0805005>
54. Paige, C.C., Saunders, M.A.: Solution of sparse indefinite systems of linear equations. *SIAM Journal on Numerical Analysis* **12**(4), 617–629 (1975). <https://doi.org/10.1137/0712047>
55. Markowitz, H.M.: Portfolio selection. *The Journal of Finance* **7**(1), 77–91 (1952). <https://doi.org/10.2307/2975974>
56. Krokmal, P., Uryasev, S., Palmquist, J.: Portfolio optimization with conditional value-at-risk objective and constraints. *Journal of Risk* **4**(2), 43–68 (2002). <https://doi.org/10.21314/JOR.2002.057>
57. Lwin, K.T., Qu, R., MacCarthy, B.L.: Mean-VaR portfolio optimization: A nonparametric approach. *European Journal of Operational Research* **260**(2), 751–766 (2017). <https://doi.org/10.1016/j.ejor.2017.01.005>
58. Rockafellar, R.T., Uryasev, S., Zabarankin, M.: Optimality conditions in portfolio analysis with general deviation measures. *Mathematical Programming* **108**, 515–540 (2006). <https://doi.org/10.1007/s10107-006-0721-9>

59. Gilbert, E., Meiklejohn, L.: A comparative analysis of risk measures: A portfolio optimisation approach. *Investment Analysts Journal* **48**(3) (2019). <https://doi.org/10.1080/10293523.2019.1643128>
60. Ramos, H., Righi, M., Guedes, C., Müller, F.: A comparison of risk measures for portfolio optimization with cardinality constraints (2022). <https://doi.org/10.2139/ssrn.4141301>
61. Righi, M.B., Borenstein, D.: A simulation comparison of risk measures for portfolio optimization. *Finance Research Letters* **24**, 105–112 (2018). <https://doi.org/10.1016/j.frl.2017.07.013>
62. Jorion, P.: Risk management lessons from long-term capital management. *European Financial Management* **6**(3), 277–300 (2000). <https://doi.org/10.1111/1468-036X.00125>
63. Boyd, S., Busseti, E., Diamond, S., N., K.R., Koh, K., Nystrup, P., Speth, J.: Multi-period trading via convex optimization. *Foundations and Trends in Optimization* **3**(1), 1–76 (2017). <https://doi.org/10.1561/24000000023>
64. Shapiro, A., Dentcheva, D., Ruszczyński, A.: *Lectures on Stochastic Programming: Modeling and Theory*. MOS-SIAM Series on Optimization. SIAM & Mathematical Optimization Society, Philadelphia (2009). <https://doi.org/10.1137/1.9780898718751>
65. Konno, H., Yamazaki, H.: Mean-absolute deviation portfolio optimization model and its applications to Tokyo stock market. *Management Science* **37**(5), 519–531 (1991). <https://doi.org/10.1287/mnsc.37.5.519>
66. Bruni, R., Cesarone, F., Scozzari, A., Tardella, F.: Real-world datasets for portfolio selection and solutions of some stochastic dominance portfolio models. *Data in Brief* **8**, 858–862 (2016). <https://doi.org/10.1016/j.dib.2016.06.031>
67. Artzner, P., Delbaen, F., Eber, J.M., Heath, D.: Coherent measures of risk. *Mathematical Finance* **9**(3), 203–228 (1999). <https://doi.org/10.1111/1467-9965.00068>
68. Elman, H.C., Ramage, A., Silvester, D.J.: Algorithm 866: IFISS, a Matlab toolbox for modelling incompressible flow. *ACM Transactions on Mathematical Software* **33**(2), 14 (2007). <https://doi.org/10.1145/1236463.1236469>
69. Elman, H.C., Ramage, A., Silvester, D.J.: IFISS: A computational laboratory for investigating incompressible flow problems. *SIAM Review* **52**(2), 261–273 (2014). <https://doi.org/10.1137/120891393>
70. Tröltzsch, F.: *Optimal Control of Partial Differential Equations: Theory, Methods and Applications, Graduate Studies in Mathematics*, vol. 112. American Mathematical Society (2010). <https://doi.org/10.1090/gsm/112>
71. Song, X., Chen, B., Yu, B.: An efficient duality-based approach for PDE-constrained sparse optimization. *Computational Optimization and Applications* **69**, 461–500 (2018). <https://doi.org/10.1007/s10589-017-9951-4>
72. Chang, C.C., Lin, C.J.: LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology* **2**, 27:1–27:27 (2011). Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>
73. Vapnik, V.N.: *The Nature of Statistical Learning Theory*. Information Science and Statistics. Springer, New York, NY (2000). <https://doi.org/10.1007/978-1-4757-3264-1>
74. Cortes, C., Vapnik, V.N.: Support-vector networks. *Machine Learning* **20**, 273–297 (1995). <https://doi.org/10.1007/BF00994018>
75. Bradley, P.S., Mangasarian, O.L.: Feature selection via concave minimization and support vector machines. In: *Proceedings of the Fifteenth International Conference on Machine Learning, ICML '98*, pp. 82–90. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1998)
76. Wang, L., Zhu, J., Zou, H.: The doubly regularized support vector machine. *Statistica Sinica* **16**(2), 589–615 (2006). <http://www.jstor.org/stable/24307560>
77. Moreau, J.J.: Proximité et dualité dans un espace Hilbertien. *Bulletin de la Société Mathématique de France* **93**, 273–299 (1965). <https://doi.org/10.24033/bsmf.1625>