



An active-set algorithm for norm constrained quadratic problems

Nikitas Rontsis¹ · Paul J. Goulart¹ · Yuji Nakatsukasa^{2,3}

Received: 9 June 2019 / Accepted: 4 January 2021 / Published online: 4 March 2021
© The Author(s) 2021

Abstract

We present an algorithm for the minimization of a nonconvex quadratic function subject to linear inequality constraints and a two-sided bound on the 2-norm of its solution. The algorithm minimizes the objective using an active-set method by solving a series of trust-region subproblems (TRS). Underpinning the efficiency of this approach is that the global solution of the TRS has been widely studied in the literature, resulting in remarkably efficient algorithms and software. We extend these results by proving that nonglobal minimizers of the TRS, or a certificate of their absence, can also be calculated efficiently by computing the two rightmost eigenpairs of an eigenproblem. We demonstrate the usefulness and scalability of the algorithm in a series of experiments that often outperform state-of-the-art approaches; these include calculation of high-quality search directions arising in Sequential Quadratic Programming on problems of the CUTEst collection, and Sparse Principal Component Analysis on a large text corpus problem (70 million nonzeros) that can help organize documents in a user interpretable way.

Keywords Nonconvex optimization · Trust-region · Active-set · Sequential quadratic programming · Dimensionality reduction

Mathematics Subject Classification 90C26 · 65F15 · 90C90

✉ Nikitas Rontsis
nrontsis@gmail.com

Paul J. Goulart
paul.goulart@eng.ox.ac.uk

Yuji Nakatsukasa
nakatsukasa@maths.ox.ac.uk

¹ Department of Engineering Science, University of Oxford, Oxford OX1 3PN, UK

² Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

³ National Institute of Informatics, Chiyoda City, Tokyo 101-8430, Japan

1 Introduction

Optimization with spherical $\|x\|_2 = 1$ or norm constraints $r_{\min} \leq \|x\|_2 \leq r_{\max}$ is used in a number of scientific fields. Many optimization problems, such as eigenvalue problems [29], dimensionality reduction [21] and compressed sensing [5], are naturally posed on or in the unit sphere, while norm constraints are often useful for regularization, e.g. in general nonlinear [7] or robust optimization [20]. In this paper we consider the solution of norm constrained problems in the form

$$\begin{aligned} & \text{minimize } f(x) := \frac{1}{2}x^T Px + q^T x \\ & \text{subject to } r_{\min} \leq \|x\|_2 \leq r_{\max} \\ & \quad Ax \leq b \end{aligned} \tag{P}$$

where $x \in \mathbb{R}^n$ is the decision variable, $P \in \mathbb{S}^n$, i.e. an $n \times n$ symmetric matrix, $A = [a_1^T \dots a_m^T] \in \mathbb{R}^{m \times n}$, $q \in \mathbb{R}^n$, $b = [b_1 \dots b_m]^T \in \mathbb{R}^m$ and r_{\min}, r_{\max} non-negative scalars.

It is generally difficult to solve problems of the form (P) due to the nonconvexity of the lower bounding norm constraint and the potential indefiniteness of the matrix P , which renders many, but certainly not all [20], of these problems intractable. Even finding a feasible point for (P) or testing if a first-order critical point of (P) is a local minimizer can be intractable (see Proposition 4 and [16]). As a result, we restrict our attention to the search for first order critical points and we assume that a feasible point is given or can be computed.

A specific tractable variation of (P) is the famous trust-region subproblem¹:

$$\begin{aligned} & \text{minimize } \frac{1}{2}x^T Px + q^T x \\ & \text{subject to } \|x\|_2 = r \\ & \quad Ax = b. \end{aligned} \tag{TRS}$$

Perhaps surprisingly, the TRS can be solved to global optimality despite its non-convexity, and global solution of the TRS has been widely studied in the literature. Specific approaches to its solution include exact Semidefinite Reformulations, Gradient Steps, Truncated Conjugate Gradient Steps, Truncated Lanczos Steps and Newton root-finding, with associated software packages for its solution. The reader can find details in [7] and [27, Chapter 4].

We suggest an active-set algorithm for the solution of (P) that, as we proceed to describe, takes advantage of the existing efficient global solution methods for the TRS. Using such an approach, (P) is optimized by solving a series of equality constrained subproblems. Each subproblem is a modification of (P) where a subset of its inequalities (called the *working-set*) is replaced by equalities while the rest are ignored. If a subproblem does not include a norm constraint then it is called an *Equality-constrained Quadratic Problem* (EQP); otherwise it is a *Trust Region Subproblem* (TRS). As described earlier, the global solution for the TRS, and also for the EQP [14], has been

¹ Note that the TRS is typically defined in the ball $\|x\| \leq r$ rather than the sphere $\|x\| = r$. We consider only the boundary solution as this fits better the needs of this paper.

widely studied in the literature. Thus we can efficiently address the global solution of each subproblem. However, a complication exists for the TRS because, unlike EQP, it can exhibit *local-nonglobal* minimizers i.e. local minimizers that are not globally optimal, which complicate the analysis. Indeed, if the optimal working-set includes a norm constraint and the global solutions of the respective TRS subproblem are infeasible for (P), then the solution of (P) must be obtained from a local-nonglobal optimizer of the TRS subproblem.

Thus, an algorithm for computing local-nonglobal minimizers of the TRS is required for our active-set algorithm. Unfortunately, algorithms for detecting the presence of/computing local-nonglobal solutions of (TRS) are significantly less mature than their global counterparts. A notable exception is the work of Martinez [25], which proved that there exists at most one local-nonglobal minimizer and gave conditions for its existence. Moreover, [25] presented a root-finding algorithm for the computation of the Lagrange Multiplier associated with the local-nonglobal minimizer that entailed the computation of the two smallest eigenvalues of the Hessian and the solution of a series of indefinite linear systems.

More recent work is based on a result from [1], which shows that each KKT point of the TRS, in the absence of linear equality constraints $Ax = b$ (i.e. with a norm constraint on x only), can be extracted from an eigenpair of

$$M := \begin{bmatrix} -P & qq^T/r^2 \\ I & -P \end{bmatrix}.$$

This result was used in [30] to calculate the local-nonglobal minimizer under the assumption that it exists. We extend these results by proving that both checking if the local-nonglobal minimizer exists and computing it, in the case where it exists, can be performed at the same time by simply calculating the two rightmost eigenpairs of M . Crucially, and similarly to the results of [1] for the global minimizer, this allows the use of efficient factorization-free methods such as the Arnoldi method for the detection and computation of the local-nonglobal minimizer. Furthermore, we show that this approach can efficiently handle equality constraints $Ax = b$ in the Trust Region Subproblem, without the need for variable reduction, by means of a projected Arnoldi method.

The paper is organised as follows. In Sect. 2 we present an introduction to the TRS and present novel results for the detection/computation of its local-nonglobal minimizer and for the incorporation of linear equality constraints. In Sect. 3.1 we introduce an active-set algorithm for solving (P) starting from the special case where $r_{\min} = r_{\max}$, and then generalizing to any r_{\min}, r_{\max} . Finally, in Sect. 4, we conclude with a series of numerical results.

Notation used: \mathbb{S}^n denotes the set of symmetric $n \times n$ real matrices. Given a vector (or matrix) x , x^T denotes its transpose and x^H its conjugate transpose.

2 The trust-region subproblem

This section concerns the computation of global as well as local-nonglobal minimizers for (TRS). Although this section only considers solutions on the boundary $\|x\| = r$, the results can be trivially extended to the interior $\|x\| \leq r$. Indeed, if an interior solution exists then the TRS is essentially convex (P is positive semidefinite in the nullspace of A [15, §1]), thus no local-nonglobal solution exists and the interior solution is obtained by solving an equality constrained quadratic problem. The presence of a (necessarily global) interior solution can be detected by checking the sign of the Lagrange Multiplier of the global boundary solution(s) [25, Lemma 2.2].

For clarity of exposition, we will first assume that there are no linear equality constraints in (TRS), resulting in the following problem:

$$\begin{aligned} &\text{minimize } \frac{1}{2}x^T Px + q^T x \\ &\text{subject to } \|x\|_2 = r, \end{aligned} \tag{T}$$

where $x \in \mathbb{R}^n$ is the decision variable, and $P \in \mathbb{S}^n, q \in \mathbb{R}^n, r \in \mathbb{R}_+$ are the problem data. For the rest of the section we will assume that $n > 1$, excluding the trivial one dimensional case where the feasible set consists of two points, and that $r > 0$. We will then show in Sect. 2.1 how to extend the results of this section to allow for the inclusion of linear equality constraints. Finally, in Sect. 2.3 we prove the main result of this section, which is presented in a separate subsection to facilitate the presentation.

In the sequel we will make frequent use of the eigendecomposition of $P := W\Lambda W^T$ defined by an orthonormal matrix $W := [w_1 \dots w_n]$ and $\Lambda := \text{diag}(\lambda_1, \dots, \lambda_n)$ where $\lambda_1 \leq \dots \leq \lambda_n$.

Every KKT point of (T) satisfies

$$Px + q + \mu x = 0 \Rightarrow x = -(P + \mu I)^{-1}q, \tag{1}$$

where μ is a Lagrange multiplier associated with the norm constraint². Equation (1) is well defined when $\mu \neq -\lambda_i, i = 1, \dots, n$. For the purposes of clarity of the subsequent introduction, which follows [27, §4], we will assume that this holds. However, the conclusions of this section are independent of this assumption.

Using the feasibility of x , we have

$$\|x\|_2^2 = r^2 \Leftrightarrow \|(P + \mu I)^{-1}q\|_2^2 - r^2 = 0. \tag{2}$$

Then, noting that

$$\begin{aligned} \|(P + \mu I)^{-1}q\|_2^2 &= \|(W(\Lambda + \mu I)W^T)^{-1}q\|_2^2 = \|W(\Lambda + \mu I)^{-1}W^Tq\|_2^2 \\ &= \|(\Lambda + \mu I)^{-1}W^Tq\|_2^2 = \sum_{i=1}^n \frac{(w_i^T q)^2}{(\lambda_i + \mu)^2}, \end{aligned}$$

² Technically, μ is the Lagrange multiplier of the equivalent constraint $\frac{1}{2}\|x\|_2^2 = \frac{1}{2}r^2$. We define (T) with $\|x\|_2 = r$ for simplicity and to match the notation of [1].




Fig. 1 A typical example of the secular equation $s(\mu)$. The dashed vertical lines mark the locations of the eigenvalues of $-P$, i.e. $-\lambda_1, \dots, -\lambda_n$. The Lagrange multiplier of the global solution of (T) is the rightmost root of $s(\mu)$ (marked by a cross). The TRS associated with this Figure also exhibits a local-nonglobal minimizer with a Lagrange multiplier that corresponds to the second-rightmost root of $s(\mu)$ (marked by a circle)

we can express the rightmost condition in (2) as

$$s(\mu) := \sum_{i=1}^n \frac{(w_i^T q)^2}{(\lambda_i + \mu)^2} - r^2 = 0. \tag{3}$$

Determining the KKT points of (T) is therefore equivalent to finding the roots of s , which is often referred to as the *secular equation* [27]. We depict s for a particular choice of P, q, r in Fig. 1. As might be expected from the tight connection between polynomial root-finding problems and eigenproblems, solving $s(\mu) = 0$ is equivalent to the following eigenproblem [1, Equation (22)]:

$$\underbrace{\begin{bmatrix} -P & qq^T/r^2 \\ I & -P \end{bmatrix}}_{:=M} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \mu \begin{bmatrix} z_1 \\ z_2 \end{bmatrix}. \tag{4}$$

This elegant result was first noted more than 50 years ago [9, Equation (2.21)], only to be later disregarded as inefficient by some of the same authors [10], and then rediscovered by [1] who highlighted its great applicability owing to the remarkable efficacy of modern eigensolvers [24].

The relation between the spectrum of M and the Lagrange multipliers of (T) is formally stated below:

Lemma 1 *The spectrum of M includes the Lagrange multiplier of every KKT point of (T).*

Proof This is a consequence of Proposition 2 of Sect. 2.3 and [9, Theorem 4.1]. \square

Note that the eigenvalues of M are complex in general since M is asymmetric. Furthermore, Lemma 1 and all the subsequent results allow for potential “degenerate” KKT points with $\mu = -\lambda_i$.

We will focus on the KKT points that are minimizers of (T) rather than maximizers or saddle points. The characterization of global minimizers is widely known in the literature. The following theorem shows that the Lagrange multiplier μ^g of the global minimizer corresponds to the rightmost eigenvalue of M in \mathbb{C} , which is always real:

Theorem 1 *A KKT point of (T) with Lagrange multiplier μ^g is a global minimizer if and only if μ^g is the rightmost eigenvalue of M . Furthermore, we necessarily have that $\mu^g \in [-\lambda_1, \infty)$.*

Proof See [1, Theorem 3.2]. \square

In general, (T) is guaranteed to possess a unique global optimizer, except perhaps in a special case that has deservedly been given a special name:

Definition (Hard case) Problem (T) belongs to the *hard case* when $\mu^g = -\lambda_1$.

We will see that in the hard case (T) is still guaranteed to have a global minimizer, but it is not necessarily unique. Furthermore, the computation of the minimizer(s) can be more challenging than in the “standard” case.

Moreover, due to the nonconvexity of (T) there can exist a local minimizer that is not global. We derive an analogous result for this minimizer that is called “local-nonglobal” [25]:

Theorem 2 *Problem (T) has a second-order sufficient local-nonglobal minimizer if and only if it does not belong to the hard case and the second-rightmost eigenvalue of M is real and simple. Furthermore, if such a minimizer exists, then its Lagrange multiplier μ^l is equal to this second-rightmost eigenvalue.*

Proof See Sect. 2.3. \square

The Lagrange multipliers for the global and local-nonglobal minimizers of (T) can therefore be identified by calculating the two rightmost eigenvalues³ of M , which can be done efficiently e.g. with the Arnoldi method.

Furthermore, for each of the respective multipliers a corresponding minimizer can be extracted immediately by the respective eigenvectors $[z_1^*; z_2^*]$ of M using as

$$x^* = -\text{sign}(q^T z_2^*) r \frac{z_1^*}{\|z_1^*\|_2}, \tag{5}$$

unless $z_1^* = 0$. Indeed, this follows from the next Proposition:

Proposition 1 *Every eigenpair $(\mu^*, [z_1^*; z_2^*])$ of M gives a KKT point (μ^*, x^*) for (T) where x^* is defined by (5), unless $z_1^* = 0$.*

³ Except perhaps when the local-nonglobal minimizer is not second-order sufficient.

Proof This proof is based on [1, §3.3] that shows the result only for the rightmost eigenpair of M .

We will first show the result for $y^* = \frac{-r^2}{q^T z_2^*} z_1^*$ and then show that $y^* = x^*$. Note that using x^* is preferred over y^* because it can avoid unnecessary numerical errors according to [1, §3.3]. We also emphasize that all of the proof assumes that z_1^* is nonzero.

For y^* , we have to show that it is well defined, feasible and stationary. Regarding the first two points, note that (4) (first row) gives

$$-Pz_1^* + qq^T z_2^*/r^2 = \mu^* z_1^* \Rightarrow (z_2^*)^T (P + \mu^* I)z_1^* = (z_2^*)^T q q^T z_2^*/r^2,$$

or since $(P + \mu^* I)z_2^* = z_1^*$ due to (4) (second row), we have $\|z_1^*\|_2^2 = (q^T z_2^*/r)^2$, which implies that $q^T z_2^* \neq 0$ because $z_1^* \neq 0$ by assumption. Thus:

$$\|z_1^*\|_2^2 = (q^T z_2^*/r)^2 \Rightarrow \left\| -\frac{r^2}{q^T z_2^*} z_1^* \right\|_2^2 = r^2 \Rightarrow \|y^*\|_2 = r.$$

To show stationarity, note that the first row of (4) also gives:

$$\begin{aligned} -Pz_1^* + qq^T z_2^*/r^2 = \mu^* z_1^* &\Rightarrow P \frac{-z_1^* r^2}{q^T z_2^*} + q + \mu^* \frac{-z_1^* r^2}{q^T z_2^*} = 0 \\ &\Rightarrow Py^* + q + \mu^* y^* = 0. \end{aligned}$$

Finally, we show that $y^* = x^*$:

$$y^* = y^* \frac{r}{\|y^*\|_2} = \frac{-r^2 z_1^*}{q^T z_2^*} \frac{|q^T z_2^*|}{\|-r^2 z_1^*\|_2} r = -\text{sign}(q^T z_2^*) r \frac{z_1^*}{\|z_1^*\|_2} = x^*.$$

□

Dealing with the hard case: Extracting a solution x^* via (5) is not possible when $z_1^* = 0$, since (5) is not well-defined in that case. In this case, (μ^*, z_2^*) is an eigenpair of $-P$ due to (4). Note that, according to [25, Lemma 3.3], this can never happen for the local-nonglobal minimizer; thus we can always extract the local-nonglobal minimizer with (5). However, it is possible for global minimizers in the hard case, i.e. the case where $\mu^* = -\lambda_1$. In the hard case, the necessary and sufficient conditions for global optimality are, due to the KKT conditions of (T) and Theorem (1), the following:

$$(P - \lambda_1 I)x + q = 0 \tag{6}$$

$$\|x\|_2 = r. \tag{7}$$

Note that $P - \lambda_1 I$ is singular, and the above system of equations represents the intersection of an affine subspace with a sphere. This intersection must be non-empty,

since (T) necessarily has a global minimizer as it arises from the minimization of a smooth function over a compact subset of \mathbb{R}^n . One can then solve (6)–(7), by computing the minimum length solution y_{\min} of $(P - \lambda_1 I)x + q = 0$ and return $x^* = y_{\min} + \alpha v$ where α is a scalar such that $\|x^*\|_2 = \|y_{\min} + \alpha v\|_2 = r$ and v is any null-vector of $P - \lambda_1 I$. Interestingly, when $z_1^* = 0$, z_2^* is a null vector of $P - \lambda_1 I$ due to (4), thus the only additional computation for extracting a solution when $z_1^* = 0$ is finding the minimum-length solution of a symmetric linear system. This can be achieved e.g. with MINRES-QLP [6], or with the Conjugate Gradient method [1, Theorem 4.3].

2.1 Equality-constrained trust-region subproblems

We will now extend the preceding results of this section to also account for the presence of linear equality constraints, $Ax = b$, in the TRS. We will show that, given an operator that projects an n -dimensional vector to the nullspace of A , we can calculate global as well as local-nonlocal minima of (TRS) by applying a projected Arnoldi method to M .

For simplicity, and without loss of generality⁴, we will assume that $b = 0$, thus solving the following problem:

$$\begin{aligned} &\text{minimize } \frac{1}{2}x^T Px + q^T x \\ &\text{subject to } \|x\|_2 = r \\ &\quad Ax = 0 \end{aligned} \tag{8}$$

where $x \in \mathbb{R}^n$ is the decision variable, $P \in \mathbb{S}^n$, i.e. an $n \times n$ symmetric matrix, A is a $m \times n$ matrix of full row rank and q is an n -dimensional vector. Similarly to the preceding results for (T), we assume that $n - m > 1$ and $r > 0$.

In principle, (8) can be solved with the “eigenproblem approach” described in this section if we reduce (8) into a smaller TRS subproblem via a nullspace elimination procedure [27, §15.3], obtaining

$$\begin{aligned} &\text{minimize } \frac{1}{2}y^T \tilde{P}y + \tilde{q}^T y \\ &\text{subject to } \|y\|_2 = r \end{aligned} \tag{9}$$

where we define Z to be a $n \times (n - m)$ orthonormal matrix with $\mathcal{R}(Z) = \mathcal{N}(A)$ and

$$\tilde{P} := Z^T P Z, \quad \tilde{q} := Z^T q.$$

According to the preceding results of this section, global and local-nonglobal minimizer(s) y^* of (9) can be computed by calculating the two rightmost eigenvectors of

$$\tilde{M} := \begin{bmatrix} -\tilde{P} & \tilde{q}\tilde{q}^T/r^2 \\ I & -\tilde{P} \end{bmatrix}.$$

⁴ Problems with linear constraints of the form $Ax = b$ can be transformed to (8) via a change of variables $\tilde{x} = x + x_0$ where $x_0 \perp \mathcal{N}(A)$ such that $Ax_0 = b$, resulting in a quadratic cost and constraints: $A\tilde{x} = 0$, $\|\tilde{x}\|_2 = \sqrt{r^2 - \|x_0\|_2^2}$.

The respective solution(s) x^* of (P) can then be recovered as

$$x^* = Zy^*. \tag{10}$$

Some disadvantages of this approach are that a nullspace basis is required and that the structure or sparsity of the problem might be destroyed.

We will present an alternative method that avoids these issues, mirroring standard approaches for solving equality constrained quadratic programs [14]. To this end, note that

$$\tilde{M} = \begin{bmatrix} -\tilde{P} & \tilde{q}\tilde{q}^T/r^2 \\ I & -\tilde{P} \end{bmatrix} = \begin{bmatrix} Z^T & 0 \\ 0 & Z^T \end{bmatrix} \begin{bmatrix} -P & qq^T/r^2 \\ I & -P \end{bmatrix} \begin{bmatrix} Z & 0 \\ 0 & Z \end{bmatrix}.$$

Since Z is an orthonormal matrix, every eigenvalue-eigenvector pair $\{\mu, [\tilde{z}_1; \tilde{z}_2]\}$ of \tilde{M} corresponds to an eigenvalue-eigenvector pair $\{\mu, [z_1; z_2]\}$ of the matrix

$$\begin{bmatrix} ZZ^T & 0 \\ 0 & ZZ^T \end{bmatrix} \underbrace{\begin{bmatrix} -P & qq^T/r^2 \\ I & -P \end{bmatrix}}_{=M} \begin{bmatrix} ZZ^T & 0 \\ 0 & ZZ^T \end{bmatrix} \tag{11}$$

with $Z[\tilde{z}_1; \tilde{z}_2] = [z_1; z_2]$. Note that (11) also has an extra eigenvalue at zero with multiplicity $2m$.

Recalling that Z is an orthonormal matrix spanning $\mathcal{N}(A)$, we note that multiplication with ZZ^T is equivalent to projection onto the nullspace of A . In practice, eigenvalues of (11) can therefore be calculated using the Arnoldi method on M where we project the starting vector and every requested matrix-vector product with M onto the nullspace of $\begin{bmatrix} A & 0 \\ 0 & A \end{bmatrix}$.

Note that the two rightmost eigenvalues of \tilde{M} correspond to the two rightmost eigenvalues of (11) if at least two eigenvalues of \tilde{M} have positive real parts. We will show that we can always transform (8) so that this holds. Indeed if we shift P to $P - \alpha I$ with an appropriate $\alpha \in \mathbb{R}$ so that \tilde{P} become negative definite⁵, then \tilde{M} has at least two eigenvalues with positive real part:

Lemma 2 *If $\tilde{P} \in \mathbb{S}^{n-m}$ is negative definite then \tilde{M} has at least two eigenvalues with positive real part, unless we trivially have $n - m = 1$.*

Proof To avoid redefining the notation used in the introduction of §2 we will show the equivalent result for (P, M) , i.e. that M has at least two eigenvalues with positive real part when P is negative definite and $n > 1$.

To do so, it suffices to show that s has at most one root with non-positive real part. This is because M has $2n \geq 4$ eigenvalues and the spectrum of M matches the roots of s on the left-hand complex plane (shown in Proposition 2 that we will prove on Sect. 2.3, given that P is negative definite).

⁵ Such an α can be obtained by applying e.g., the Lanczos method, or the Gershgorin circle Theorem on P .

To show this, note that s has at most one (simple) real root with nonpositive real part since

$$s'(\alpha) = -2 \sum_{j=1}^n \frac{(w_j^T q)^2}{(\lambda_j + \alpha)^3} \tag{12}$$

is positive for any $\alpha \leq 0$ since $\lambda_1, \dots, \lambda_n < 0$ by assumption. To conclude the proof, it suffices to show that all the roots of s with nonpositive real part are real. Indeed, note that for any $a, b \in \mathbb{R}$ we have

$$\text{Im}(s(a + bi)) = - \sum_{j=1}^n \frac{2(\lambda_j + a)b}{((\lambda_j + a)^2 + b^2)^2} (w_j^T q)^2,$$

where i is the imaginary unit. A detailed derivation of the above equation is provided in Lemma 4. For any $a \leq 0$ we have $-a \geq 0 > \lambda_1, \dots, \lambda_n \Rightarrow 2(\lambda_j + a) < 0, j = 1, \dots, n$ by assumption. Thus $\text{Im}(s(a + bi)) \neq 0 \Rightarrow s(a + bi) \neq 0$ for any $a \leq 0$ unless $b = 0$. □

Such a shift would not affect the (local/global) optimizers of (P) as

$$\frac{1}{2}x^T(P - \alpha I)x + q^T x = \frac{1}{2}x^T P x + q^T x - \frac{\alpha}{2}\|x\|_2^2$$

or, since $\|x\|_2 = r$,

$$\frac{1}{2}x^T(P - \alpha I)x + q^T x = \frac{1}{2}x^T P x + q^T x - \frac{\alpha}{2}r^2,$$

where $\alpha \in \mathbb{R}$. Thus, we can always transform (8) so that the two rightmost eigenvalues of \tilde{M} are equal to the two rightmost eigenvalues of (11).

Having computed the two rightmost eigenvalues/vectors with a projected Arnoldi method applied to M , global/local solutions of (P) can be obtained from the respective eigenvectors. Indeed, recall that according to (5) and (10), an appropriate (rightmost/second-rightmost) eigenvector $[\tilde{z}_1^*; \tilde{z}_2^*]$ of (11) gives a solution x^* to (P) as follows

$$x^* = Zy^* = -Z \text{sign}(\tilde{q}^T \tilde{z}_2^*) r \frac{\tilde{z}_1^*}{\|\tilde{z}_1^*\|_2} = -\text{sign}(q^T (Z\tilde{z}_2^*)) r \frac{Z\tilde{z}_1^*}{\|Z\tilde{z}_1^*\|_2} \tag{13}$$

$$= -\text{sign}(q^T z_2^*) r \frac{z_1^*}{\|z_1^*\|_2} \tag{14}$$

where $[z_1^*; z_2^*]$ is an eigenvector of M . Thus, x^* can be extracted solely by the eigenvector of M unless $z_1^* = 0$.

Dealing with the hard case: In the case where $z_1^* = 0$, which can appear in the hard case, one would have to calculate the minimum length solution of

$$(\tilde{P} + \mu^* I)y + \tilde{q} = 0 \tag{15}$$

and then construct a global solution x^* as $x^* = Z(y + \alpha \tilde{z}_2^*) = Zy + \alpha z_2^*$, where α is a scalar such that $\|x^*\|_2 = r$. However, the minimum length solution of (15) can be obtained by the minimum length solution of

$$(P + \mu^* I)x + q = 0 \tag{16}$$

$$Ax = 0 \tag{17}$$

which can be calculated via projected MINRES-QLP or a projected CG algorithm [1, Theorem 4.3] [14].

2.2 An algorithm for computing global and local solutions of the TRS

Using the ideas presented in this section, we now present Algorithm 1, a practical algorithm for computing global and second-order sufficient local-nonglobal solutions of equality constrained trust-region subproblems. In the remainder of this section, we establish the correctness of Algorithm 1.

Algorithm 1 uses the results of this section (in particular Theorems 1 and 2, Equation (13) and the paragraph “dealing with the hard case” of Sect. 2.1) to compute the minimizers of (8). The only novel point is the way we detect the presence of the local-nonglobal minimisers, which we proceed to show to be valid.

If the TRS is in the hard case, then the algorithm should only return global minimiser(s). Indeed, this is the case as

1. If $z_1^s = 0$, then the algorithm terminates at line 16, according to the remarks of the paragraph “Dealing with the hard case” of Sect. 2.1.
2. If $z_1^s \neq 0$, then there must exist another eigenvector associated with μ^s besides $[z_1^s; z_2^s]$. This is because $\mu^s = -\lambda_1$ thus according to Proposition 2 of Sect. 2.3 there exists a vector $u \in \mathcal{N}(P - \lambda_1 I)$ that is orthogonal to q which makes $(\mu^s, [0; u])$ an eigenpair of M due to (4).

Thus $\mu^s = \mu^l$, which implies that μ^l is not simple. As a result the algorithm will only return a global minimizer using (5) at line 11.

If the TRS is not in the hard case, then we must have $z_1^s \neq 0$ (see remarks of paragraph “Dealing with the hard case” of Sect. 2). Thus the algorithm will return either in lines 9 or 11, detecting correctly the presence of the local-nonglobal minimiser according to Theorem 2.

Finally, we emphasize that in the hard case the TRS can have an infinite number of solutions, but only one or two of them are returned by Algorithm 1.

Algorithm 1: Calculation of global and local-nonglobal minimizers of the equality constrained trust-region subproblem (8)

- 1 **Given** the problem data $P \in \mathbb{S}^n, q \in \mathbb{R}^n, r \in \mathbb{R}_+, A \in \mathbb{R}^{m \times n}$ where A is full row rank with $m < n - 1$ and $\Pi : \mathbb{C}^n \mapsto \mathbb{C}^n$, a projector into the nullspace of A ;
- 2 Shift P by a multiple of identity so that it becomes negative definite;
- 3 Run an Arnoldi method on

$$\Pi \left(\begin{bmatrix} -P & qq^T/r^2 \\ I & -P \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \right) = \mu \begin{bmatrix} z_1 \\ z_2 \end{bmatrix},$$

where $\Pi \left(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \right) := \begin{bmatrix} \Pi(x_1) \\ \Pi(x_2) \end{bmatrix}$ for any $x_1, x_2 \in \mathbb{C}^n$, starting from an initial point $\Pi(z) \in \mathbb{C}^{2n}$, to calculate a rightmost eigenvalue/vector pair (μ^g, z^g) ;

- 4 **if** $z_1^g \neq 0$ **then**
 - 5 $x^g \leftarrow -\text{sign}(q^T z_2^g) r \frac{z_1^g}{\|z_1^g\|_2}$;
 - 6 Resume the projected Arnoldi method to compute a second-rightmost eigenvalue/vector (μ^ℓ, z^ℓ) ;
 - 7 **if** μ^ℓ is real and simple **then**
 - 8 $x^\ell \leftarrow -\text{sign}(q^T z_2^\ell) r \frac{z_1^\ell}{\|z_1^\ell\|_2}$;
 - 9 **Return** x^g and x^ℓ as the unique global and local-nonglobal minimizers;
 - 10 **else**
 - 11 **Return** x^g as a global minimizer and state that no second-order sufficient local-nonglobal minimizer exists.;
 - 12 **end**
 - 13 **else**
 - 14 Compute the minimum length solution x_{\min} of

$$\begin{aligned} (P + \mu^g I)x + q &= 0 \\ Ax &= 0 \end{aligned}$$

via projected MINRES-QLP or projected CG [1, Theorem 4.3];

 - 15 Solve a quadratic equation to find α_1, α_2 such that $\|x_{\min} + \alpha_1/2 z_2^g\|_2 = r$;
 - 16 **Return** $x_{\min} + \alpha_1 z_2^g, x_{\min} + \alpha_2 z_2^g$ as global solutions, state that no local-nonglobal minimizer exists;
 - 17 **end**
-

2.3 Proof of Theorem 2

In this subsection we provide a proof of our main theoretical result, Theorem 2, which shows that (T) has a second-order sufficient local-nonglobal minimizer if and only if it does not belong to the hard case and the second-rightmost eigenvalue of the matrix

$$M := \begin{bmatrix} -P & qq^T/r^2 \\ I & -P \end{bmatrix},$$

with counted algebraic multiplicities, is real and simple. Furthermore, we show that if such a minimizer exists, its Lagrange multiplier μ^ℓ is equal to this second-rightmost eigenvalue.

The proof is based on the following Lemma, which follows from [25]:

Lemma 3 *A KKT point x^ℓ of (T) with Lagrange multiplier μ^ℓ is a second-order sufficient local-nonglobal minimizer iff $q \not\perp w_1$, $\mu^\ell \in (-\lambda_1, -\lambda_2)$ and $s'(\mu^\ell) > 0$.*

Proof See [25, Theorem 3.1]. □

Before we use Lemma 3 we will need the following two ancillary results:

Lemma 4 *For any complex number $a + bi \in \mathbb{C}$ where a, b are real numbers such that a is in the domain of s and $b \neq 0$, we have:*

$$\text{Re}(s(a + bi)) < s(a). \tag{18}$$

Proof Recall the definition of s :

$$s(a + bi) = \sum_{j=1}^n \frac{(w_j^T q)^2}{(\lambda_j + a + bi)^2} - r^2$$

and define, for convenience, $\kappa_j = w_j^T q$, $j = 1, \dots, n$. Thus

$$\begin{aligned} s(a + bi) &= -r^2 + \sum_{j=1}^n \frac{\kappa_j^2}{(\lambda_j + a + bi)^2} \\ &= -r^2 + \sum_{j=1}^n \kappa_j^2 \frac{(\lambda_j + a)^2 - b^2}{((\lambda_j + a)^2 + b^2)^2} - \kappa_j^2 \frac{2(\lambda_j + a)b}{((\lambda_j + a)^2 + b^2)^2} i, \end{aligned}$$

hence

$$\begin{aligned} \text{Re}(s(a + bi)) &= -r^2 + \sum_{j=1}^n \kappa_j^2 \frac{(\lambda_j + a)^2 - b^2}{((\lambda_j + a)^2 + b^2)^2} \\ &< -r^2 + \sum_{j=1}^n \kappa_j^2 \frac{(\lambda_j + a)^2}{((\lambda_j + a)^2)^2}, \end{aligned}$$

or, equivalently, $\text{Re}(s(a + bi)) < s(a)$. □


Lemma 5 *If $s(x) < 0$ for some real x in the domain of $s : \mathbb{C} \mapsto \mathbb{C}$ then s has an equal number of poles and zeros with real parts in $[x, \infty)$.*

Proof We will prove this result with the argument principle. We propose encircling the half-space

$$S(x) := \{c \in \mathbb{C} | c = a + bi, \text{ where } a \geq x, b \in \mathbb{R}\}$$

with the contour of Fig. 2 and show that it maps to a contour that lies strictly in the left half plane. The suggested closed contour C intersects the real axis at x and at ∞ .

Fig. 2 Suggested closed contour for encircling $S(x)$



It consists of two segments: $C_1 = \{\omega + \beta i \mid \beta \in \mathbb{R}\}$, i.e. a segment that is parallel to the imaginary axis, and C_2 , a half-circle with infinite radius.

By the argument principle [2] we have $\frac{1}{2\pi i} \int_C \frac{s'(z)}{s(z)} dz = Z - P$, where Z, P are the number of roots and poles of s in the domain enclosed by C . It follows that the number of poles in S is equal to the number of zeros in S if $\int_C \frac{s'(z)}{s(z)} dz = 0$, which holds if there are no poles or zeros of s on C and $s(C)$ does not enclose the origin. By examining (3) we conclude that s has only real, finite poles. Thus, the contour C does not intersect any pole of s . It remains to show that there exist no zeros of s on C and $s(C)$ does not enclose the origin. We show both by proving that every point in $s(C) = s(C_1) \cup s(C_2)$ has negative real part. For $s(C_2)$, it is easy to see that, since $s(C_2) = \{-r^2\}$ as $\lim_{|w| \rightarrow \infty} s(w) = -r^2$. Regarding $s(C_1)$ note that $s(x) < 0$, and thus Lemma 4 gives $\text{Re}(s(x + \beta i)) \leq s(x) < 0 \forall \beta \in \mathbb{R}$, i.e., any point in $s(C_1)$ has negative real part. □

We will now connect the second-rightmost root of the secular equation s with the results of Lemma 3. Note that, due to Lemma 4, if s has a real root μ then it cannot also have an unreal root with real part μ and vice versa. Thus, the concept of “the second-rightmost root of s ” is well defined as long as we assume (or are prepared to check) that this root is real.

Lemma 6 (i) Suppose that $q \not\prec w_1$ and $\mu \in (-\lambda_2, -\lambda_1)$ is a simple root of s with $s'(\mu) > 0$. Then, μ is the second-rightmost root of s in \mathbb{C} .

(ii) Suppose that $q \not\prec w_1, w_2, \lambda_1 \neq \lambda_2, \mu \in \mathbb{R}$ is the second-rightmost root of s in \mathbb{C} and μ is simple. Then, $\mu \geq -\lambda_2$.

Proof Note that in both cases $q \not\prec w_1$, thus (T) does not belong to the hard case [1]. Hence s has a real root in $(-\lambda_1, \infty)$ (Theorem 1) which we denote by μ^g . Moreover, μ^g is simple because $s'(\mu^g) = -2 \sum_{j=1}^n (w_j^T q)^2 / (\lambda_j + \mu^g)^3$ is negative since $\mu^g > -\lambda_1 > \dots > -\lambda_n$.

In the subsequent analysis, we will frequently refer to the poles of s which are

$$\{-\lambda_i \mid i = 1, \dots, n \text{ such that } q^T w_i \neq 0\},$$

and have double multiplicity, and the following parametric set

$$S(x) := \{c \in \mathbb{C} \mid c = a + bi, \text{ where } a \geq x, b \in \mathbb{R}\}$$

defined for any $x \in \mathbb{R}$.

We will begin with claim (i). Since $\mu < -\lambda_1 < \mu^s$ we conclude that μ is not the rightmost root of s . In order to show that it is the *second* rightmost, it suffices to show that there are at most two roots of s in $S(\mu)$.

Since $s(\mu) = 0$ and $s'(\mu) > 0$, we have $s(\mu - \epsilon) < 0$ for a sufficiently small $\epsilon > 0$. According to Lemma 5, s has equal number of zeros and poles in $S(\mu - \epsilon)$. By assumption $q \not\perp w_1$ and $\mu > -\lambda_2$, thus $-\lambda_1$ is the only (double) pole of s in $S(\mu - \epsilon)$. Hence, the number of zeros in $S(\mu - \epsilon) \supset S(\mu)$ must be two. This concludes the proof of claim (i).

We will now prove claim (ii). Assume the contrary, i.e. that $\mu < -\lambda_2$. Consider first the case where $s'(\mu) > 0$. Since $s(\mu) = 0$ we can choose a sufficiently small $\epsilon > 0$ such that $s(\mu - \epsilon) < 0$. Thus s has equal number of zeros and poles in $S(\mu - \epsilon)$ (Lemma 5). By assumption $\mu < -\lambda_2$ and $q \not\perp w_1, w_2$, and thus the double poles $-\lambda_1$ and $-\lambda_2$ are in $S(\mu - \epsilon)$. Thus s has (at least) four zeros in $S(\mu - \epsilon)$, and since ϵ can be arbitrarily small, the same holds for $S(\mu)$.

We conclude that there should be extra roots for s in $S(\mu)$, besides the simple roots μ^s and μ . Moreover, these extra roots cannot have real part equal to μ as according to Lemma 4 $\text{Re}(s(\mu + bi)) < s(\mu) = 0$ for any real $b \neq 0$. Thus their real parts must be in (μ, ∞) . This means that μ is not the second-rightmost root of s , which is a contradiction.

Finally, if $s'(\mu) < 0$, then, using the same arguments as before, we conclude that there exist four roots in $S(\mu + \epsilon)$ for any sufficiently small $\epsilon > 0$. Thus, there exist extra roots (besides the simple root μ^s) with real part in (μ, ∞) , which, again, contradicts the assumption that μ is the second-rightmost root of s . □

Lemma 6, in combination with Lemma 3, provides an “almost” iff relationship between the second-rightmost root of s and the local-nonglobal minimizer of (T). However, the quantity of primary interest is the second-rightmost eigenvalue of M instead of the second-rightmost root of s . The following result characterizes the spectrum of M and its tight connection with the roots of s :

Proposition 2 *The spectrum of M consists of:*

- the roots of s with matched algebraic multiplicity; and
- all of the eigenvalues $-\lambda_i$ of $-P$ that are non-simple for $-P$ or have w_i (the i -th eigenvector of P) orthogonal to q .

Furthermore, every $-\lambda_i$ that is in the spectrum of M is non-simple for M .

Proof We will prove the result by deriving a closed form expression for the characteristic polynomial of M . Since $-I$ and $P + \mu I$ commute, we have [31, Theorem 3]:

$$\det(\mu I - M) = \det \left(\begin{bmatrix} P + \mu I & -qq^T/r^2 \\ -I & P + \mu I \end{bmatrix} \right) = \det((P + \mu I)^2 - qq^T/r^2)$$

or, using the matrix determinant lemma,

$$\det(\mu I - M) = -r^{-2} \left(q^T (P + \mu I)^{-2} q - r^2 \right) \det((P + \mu I)^2)$$

for all $\mu \neq -\lambda_1, \dots, -\lambda_n$. Thus, recalling (2)–(3) we have

$$\det(\mu I - M) = -r^{-2} \left(\sum_{i=1}^n \frac{(w_i^T q)^2}{(\lambda_i + \mu)^2} - r^2 \right) \prod_{i=1}^n (\lambda_i + \mu)^2 \tag{19}$$

$$= -r^{-2} \sum_{i=1}^n (w_i^T q)^2 \prod_{j \neq i} (\lambda_j + \mu)^2 + \prod_{i=1}^n (\lambda_i + \mu)^2, \tag{20}$$

for all $\mu \neq -\lambda_1, \dots, -\lambda_n$. Thus, (20) and the characteristic polynomial of M coincide in all \mathbb{C} except perhaps on $2n$ points. It follows from continuity arguments that (20) is in fact the characteristic polynomial of M .

By examining (19) (2nd leftmost) and (20), we conclude that the eigenvalues of M include all the roots of s with matched algebraic multiplicity. The secular equation s has $2n$ roots whenever the eigenvalues of P are distinct and $w_i^T q \neq 0$ for all $i = 1, \dots, n$. Otherwise, i.e. for every i with λ_i non-simple or $w_i^T q = 0$, $M \in \mathbb{R}^{2n \times 2n}$ has extra non-simple eigenvalues at these $-\lambda_i$, which constitute the only differences between the spectrum of M and the roots of s . □

The following result follows directly from the above Proposition and will be useful for the rest of this section.

Corollary 1 *No simple eigenvalue of M is in the spectrum of $-P$.*

We are now ready to prove the main result of this section, Theorem 2. The proof will occupy the rest of this subsection. Note, again, that, due to Lemma 4 and Proposition 2, if M has a real eigenvalue μ then it cannot also have an unreal eigenvalue with real part μ and vice versa. Thus, the concept of “the second-rightmost eigenvalue of M ” is well defined as long as we assume (or are prepared to check) that this eigenvalue is real.

“only-if” In this case we know that μ^ℓ is the Lagrange multiplier of a second-order sufficient local-nonglobal minimizer and we have to show that

- (i) (T) is not in the hard case.
- (ii) μ^ℓ is a simple eigenvalue of M .
- (iii) μ^ℓ is the second-rightmost eigenvalue of M .

Points (i), and (ii) follow directly from Lemma 3 and Proposition 2 (note that $q \not\perp w_1$ implies that (T) is not in the hard case [1]). In order to prove (iii) we will first show that the spectrum (counted with algebraic multiplicities) of M with real part larger than $-\lambda_2$ coincides with the roots of s in the same region. This follows from Proposition 2, because λ_1 is simple (due to $\mu^\ell \in (-\lambda_2, -\lambda_1)$) and $q^T w_1 \neq 0$. Point (iii) then follows from Lemma 6(i) which shows that μ^ℓ is the second-rightmost root of s , and thus the second-rightmost eigenvalue of M .

“if” In this case, we know that (T) is not in the hard case and that μ^ℓ is the second-rightmost eigenvalue of M which is real and simple, and we want to show that μ^ℓ is the Lagrange multiplier of the second-order sufficient local-nonglobal minimizer of (T).

Note that, by assumption, μ^ℓ is in the spectrum of M but, due to Corollary 1, not in that of $-P$. Thus μ^ℓ is a root of s due to Proposition 2. As a result, it suffices to show that μ^ℓ is in $(-\lambda_2, -\lambda_1)$ and $s'(\mu^\ell) > 0$ (Lemma 3).

We will first show that $\mu^\ell \in (-\lambda_2, -\lambda_1)$. Since we are not in the hard case, M has a real eigenvalue in $(-\lambda_1, \infty)$ (Theorem 1) which is simple and unique in $(-\lambda_1, \infty)$ due to Proposition 2 and the fact that $s'(\mu) = -2 \sum_{j=1}^n (w_j^T q)^2 / (\lambda_j + \mu)^3$ is negative in that region. Thus, $-\lambda_1$ is not in the spectrum of M as otherwise it would be its second-rightmost eigenvalue (which by assumption is real) and it would imply that μ^ℓ is in the spectrum of $-P$, which we have already excluded. Thus, $\mu^\ell \in (-\infty, -\lambda_1)$, and $\mu^\ell \neq -\lambda_2$.

It remains to show that $\mu^\ell \geq -\lambda_2$. If $q^T w_2 = 0$ then, $(-\lambda_2, [0; w_2])$ is an eigenpair of M , thus we must have $\mu^\ell > -\lambda_2$ so as to avoid μ^ℓ (the second-rightmost eigenvalue of M) being in the spectrum of $-P$. Consider now the case where $q^T w_2 \neq 0$. Note that $\lambda_1 \neq \lambda_2$ and $q \not\perp w_1$ as otherwise $-\lambda_1$ is in the spectrum of M (Proposition 2), which, according to the paragraph above is a contradiction. Thus, $\lambda_1 \neq \lambda_2$ and $q \not\perp w_1, w_2$ resulting in $\mu^\ell \geq -\lambda_2$ due to Lemma 6(ii).

Finally we show that $s'(\mu^\ell) > 0$, thereby concluding the proof. Note that $s'(\mu^\ell) \neq 0$ since by assumption μ^ℓ is a simple root of s . Assume the contrary, i.e. $s'(\mu^\ell) < 0$. Since s is convex in $(-\lambda_2, -\lambda_1)$ [25, (3.12)] and $\lim_{\mu \rightarrow -\lambda_1} s(\mu) = \infty$ we conclude that there must exist another root of s in $(\mu^\ell, -\lambda_1)$. This contradicts the assumption that μ^ℓ is the second-rightmost root of s .

3 An active-set algorithm for (P)

We are now in a position to present the main contribution of this paper, an active-set algorithm for solving (P). We will first present an algorithm for the special that $r_{\min} = r_{\max} := r$ and then describe how to generalize for any r_{\min}, r_{\max} .

3.1 Solving (P) when $r_{\min} = r_{\max} := r$

In this section we introduce a primal active-set approach for the optimization of (P) when $r_{\min} = r_{\max} := r$. It will be useful in the subsequent analysis to recall the Karush-Kuhn-Tucker (KKT) conditions of (P), which are

$$\nabla f(x) + A^T \kappa + \mu x = 0 \tag{21}$$

$$\kappa \geq 0 \tag{22}$$

$$\kappa_i (a_i^T x - b_i) = 0, \quad i = 1, \dots, m \tag{23}$$

$$Ax \leq b, \quad x^T x = r^2. \tag{24}$$

As is typical from a primal active-set approach, our algorithm starts from a given feasible point of (P) and generates iterates that remain feasible for (P) and have non-increasing objective values. At every iteration we treat a subset of the inequality constraints $Ax \leq b$ as equalities. We refer to this subset as the *working set* \mathcal{W}_k and define

$$\bar{A}_k = [a_i]_{i \in \mathcal{W}_k}, \quad \bar{b}_k = [b_i]_{i \in \mathcal{W}_k},$$

where $[\cdot]$ denotes vertical (row-wise) concatenation. To simplify the analysis, we will assume that one of the simplest and most common constraint qualification holds for every iterate of the algorithm:

Assumption 1 Linear Independence Constraint Qualification [LICQ]

The LICQ holds for every iterate of the suggested Algorithm, i.e. $\begin{bmatrix} \bar{A}_k \\ \bar{x}_k^T \end{bmatrix}$ is full row rank.

We will first give a brief, schematic description of our algorithm. At every iteration k of our active-set based minimization procedure, we use Algorithm 1 to compute minimizers of the subproblem

$$\begin{aligned} &\text{minimize } \frac{1}{2}x^T Px + q^T x \\ &\text{subject to } \|x\|_2 = r \\ &\quad \bar{A}_k x = \bar{b}_k. \end{aligned} \tag{SP}$$

We then attempt to move towards those minimizers in that hope that we either (i) ‘hit’ a new constraint or (ii) set x_{k+1} to a (potentially local) minimizer of (SP). Unfortunately due to the nonconvexity of the problem, we might achieve neither (i) nor (ii) by simply moving towards the minimizer of (SP), as the objective of (P) is not guaranteed to be non-increasing along these moves. In that case, which we will show can happen only when (SP) has a unique minimizer, we perform a projected gradient descent followed (if necessary) by a move towards the global minimizer of (SP) that is guaranteed to achieve either (i) or (ii)⁶. Finally, at the end of each iteration, we check for termination and update the working set.

Our overall active-set procedure is shown in Algorithm 2. We now proceed in the following subsections that describe in detail each of the Algorithm’s steps. For the rest of this section, we will say that a point is ‘feasible’ when it is in the feasible region of (P). Also, given a point x we will call the projection of $\nabla f(x)$ to the nullspace of constraint normals of (SP), $\begin{bmatrix} \bar{A}_k \\ x^T \end{bmatrix}$, as ‘projected gradient of $f(\cdot)$ at x ’.

3.1.1 Move towards the minimizers of (SP)

Notice that Algorithm 1 always returns a global minimizer x^g of (SP) and possibly a second minimizer x^* .

⁶ Technically, in this case x_{k+1} will be a second order necessary stationary point of (SP).

Algorithm 2: Active Set method for solving (P) when $r_{\min} = r_{\max} := r$

```

1  Assume a feasible starting point  $x_0$ 
2  Set  $W_0$  to be a subset of the active constraints in  $x_0$ 
3  for  $k = 0, 1, 2, \dots$  do
4  | Solve (SP) with Algorithm 1;
5  | if Algorithm 1 returned two distinct minimizers then
6  | |  $x_{k+1} \leftarrow$  solution of the 2D subproblem defined by  $x_k$  and the two returned minimizers.
7  | else
8  | |  $x_{k+1} \leftarrow$  solution of the 2D subproblem defined by  $x_k$ , the returned (global) minimizer and
9  | | the projected gradient of  $f$  at  $x_k$ .
10 | end
11 | if  $x_{k+1}$  is not a minimizer of (SP) and no new constraint was "hit" then
12 | | Run projected gradient descent starting from  $x_{k+1}$  until a new constraint is "hit" or until
13 | | convergence, and store the resulting point in  $x_{k+1}$ ;
14 | | if PGD converged to a feasible  $x_{k+1}$  of indefinite projected Hessian then
15 | | | Compute a suitable limiting direction  $d$  along which  $x_{k+1}$  is a local maximum;
16 | | |  $x_{k+1} \leftarrow$  solution of the 2D subproblem on the plane defined by  $x_{k+1}$ , a global
17 | | | minimizer of (SP) and  $d$ ;
18 | | if a new constraint was "hit" then
19 | | | Obtain  $W_{k+1}$  by adding one of the blocking constraints to  $W_k$ ;
20 | | else
21 | | |  $W_{k+1} \leftarrow$  check_multipliers( $x_{k+1}, W_k$ );
22 | end
23 Function check_multipliers( $x, W$ )
24 | Compute the Lagrange multipliers ( $\kappa, \mu$ ) at  $x$  that satisfy (21)-(24) with  $A = [a_i]_{i \in W}$ ;
25 | if  $\kappa \geq 0$  then
26 | | Terminate Algorithm with  $x$  as the returned solution
27 | else
28 | |  $j \leftarrow \operatorname{argmin}_{j \in W} \kappa$ ;
29 | |  $W \leftarrow W \setminus \{j\}$ ;
30 | end

```

Let us consider first the case where two distinct minimizers are returned which is treated in lines 5-6 of Algorithm 2. In this case it is always possible to either set x_{k+1} to a minimizer of (SP) or "hit" a new constraint as follows. Consider the circle defined by the intersection of the sphere $\|x\|_2 = r$ and the plane defined by x_k and the two returned minimizers. When (SP) is constrained in this circle, then it can be reduced to a two-dimensional TRS. As such, we will show that it has at most 2 maximizers and 2 minimizers, except perhaps when x_k is a global minimizer of (SP) in which case we can trivially set $x_{k+1} = x_k$.

Proposition 3 *A two dimensional TRS has at most 2 minimizers and 2 maximizers except when its objective is constant across its feasible region.*

Proof It suffices to show the result only for the minimizers, as a negation of the TRS's objective would show the same result for its maximizers. Suppose that the TRS does not belong to the hard case. Then it has a unique global minimizer and at most one local-nonglobal minimizer [25]. If it belongs to the hard case, then only global minimizers can exist. These are intersections of the affine subspace (6) and the sphere (7). Note

that this intersection is either a distinct point, two points or a circle. In the latter case, every feasible point of (T) is a global minimizer. This completes the proof. \square

We can identify the two minimizers of this two-dimensional TRS as x^g and x^* . It follows that at least in one of the circular arcs connecting x_k with x^g and x^* the objective value is always less than $f(x_k)$. Thus, by moving into that circular arc one will either end up with x^g or x^* or “hit” a new constraint.

On the other hand, if a single global minimizer x^g was returned, which corresponds to lines 7-8 of Algorithm 2, then we consider the circle defined by the intersection of the sphere $\|x\|_2 = r$ and the plane defined by x_k, x^g and the projected gradient of $f(\cdot)$ at x_k . Obviously, if x^g is feasible for (P) then we choose $x_{k+1} = x^g$, but otherwise we are not guaranteed to “hit” a new constraint. This is because the associated two-dimensional TRS might possess a second minimizer, which is not necessarily a stationary point of (SP), but might be the best feasible solution on this circle.

3.1.2 Perform projected gradient descent (if necessary)

This subsection concerns the case where the procedure described in the previous subsection could neither set x_{k+1} to a minimizer of SP nor “hit” a new constraint. In this case (lines 10-14 of Algorithm 2), we proceed by performing projected gradient descent (PGD) on (SP) starting from the current iterate of the active-set algorithm. This is guaranteed to converge to a KKT point x^s with $f(x^s) \leq f(x_k)$ [4, Theorem 4.5]. Suppose that x^s is not feasible for (P). Since the feasible region is a closed set, PGD exits the feasible region in finitely many steps. By stopping the projected gradient descent just before it exits the feasible region, we can find a point x_{k+1} with a newly “hit” constraint.

Consider now the case where x^s is feasible for (P). If the minimum eigenvalue λ_{\min} of the projected Hessian of $f(\cdot)$ at x^s is nonnegative (lines 12-14 of Algorithm 2), then we set $x_{k+1} = x^s$ and we proceed to the steps outlined in the next subsection. Otherwise $\lambda_{\min} < 0$, and there exists a feasible sequence $\{z_k\}$ (w.r.t. to (SP)) converging to x^s with an appropriate limiting direction d

$$d := \lim_{k \rightarrow \infty} \frac{z_k - x^s}{\|z_k - x^s\|_2}, \tag{25}$$

such that

$$f(z_k) = f(x^s) + \frac{1}{2} \lambda_{\min} \|z_k - x^s\|^2 + \mathcal{O}(\|z_k - x^s\|^2). \tag{26}$$

In practice, we can choose d equal to some projected eigenvector associated with $\lambda_{\min} < 0$. Consider now the circle defined by the intersection of the sphere $\|x\|_2 = r$ and the plane defined by x^s, x^g and the limiting direction d . As before, this can be reduced to a two-dimensional TRS that possesses at most 2 minimizers and 2 maximizers (Proposition 3). We can identify two of them: x^g which is a global minimum and x^s which, according to (26), is a local maximum. It follows that in at least one of the two circular arcs connecting x^s with x^g the objective value is always less than $f(x^s)$.

Thus, by moving into that circular arc, and since x^s is feasible for (P) and x^g infeasible, we can identify a suitable feasible point x_{k+1} such that $f(x^g) \leq f(x_{k+1}) \leq f(x^s)$ for which a new constraint, not in the current working set, becomes active.

3.1.3 Update the working set and check for termination

If a new constraint was identified in the steps above, then we update the working set and proceed to the next iteration (lines 15-16 of Algorithm 2). Otherwise, x_{k+1} was set to a second-order necessary stationary point of (SP) and we proceed as follows (lines 17-18 of Algorithm 2). Stationarity of x_{k+1} gives

$$\nabla f(x_{k+1}) + \mu x_{k+1} + \sum_{i \in \mathcal{W}_k} a_i \kappa_i = 0$$

for some Lagrange multipliers μ and $\kappa_i \forall i \in \mathcal{W}_k$. Thus, (x_{k+1}, μ, κ) satisfy the first KKT condition (21) if we define $\kappa_i = 0 \forall i \notin \mathcal{W}_k$. It follows from feasibility of x_{k+1} and the definition of κ that the third and fourth KKT conditions (23)-(24) also hold. If we also have that $\kappa_i \geq 0 \forall i \in \mathcal{W}_k$, then (x_{k+1}, μ, κ) is a KKT pair for (P) and we terminate the algorithm (lines 23-24 of Algorithm 2). If, on the other hand, $\kappa_i < 0$ for some $i \in \mathcal{W}_k$ (lines 25-27 of Algorithm 2), then (22) is not satisfied and x_{k+1} is not a local minimizer for (P). In fact, the objective $f(\cdot)$ may be decreased by dropping one of the constraints corresponding to a negative Lagrange multiplier [27, Section 12.3].

3.1.4 Finite termination

We will show that Algorithm 2 terminates in finitely many outer iterations, assuming that $x_{k+1} \neq x_k$ for every $k = 1, 2, \dots$. Indeed at every iteration the algorithm either

- Moves to a stationary point of the current TRS subproblem; or
- Activates a new constraint.

Since there can be at most m constraints in the working set it follows that x_k visits a stationary point of the k -th TRS subproblem periodically (at least every m iterations). Furthermore, note that every TRS subproblem, resulting from a particular working set, has at most $2n$ sets of stationary points of the same objective value [9, Theorem 4.1 and subsequent comments]. Since there are a finite number of different working set configurations, it follows that there exist finitely many such sets. On the other hand, every constraint that is deleted from the working set has an associated negative Lagrange multiplier, thus every iterate x_k is not a stationary point of the next iteration. Furthermore, note that if x_k is not a stationary point for the $k + 1$ subproblem then our algorithm generates x_{k+1} with $f(x_{k+1}) < f(x_k)$ unless $x_k = x_{k+1}$ which is excluded by assumption. This means that once the algorithm visits one of these sets of stationary points with equal objective value, it can never visit it again. Hence, the algorithm terminates after finitely many iterations.

We believe that it is not possible to bound the number of iterations by a polynomial in m, n . Indeed, even finding an initial feasible point is NP-complete, as we show in Sect. 3.2.1.

3.1.5 Further remarks

Similarly to active-set algorithms for quadratic programs, we can always update the working set such that \bar{A}_k is full row rank. However, LICQ might still fail to be satisfied when the gradient of the spherical constraint lies on $\mathcal{R}(\bar{A}_k^T)$. In these cases we terminate the algorithm without guarantees about local optimality.

3.2 Solving (P) for any r_{\min}, r_{\max}

We are now ready to present an active-set algorithm that solves (P) for any r_{\min} and r_{\max} . At every iteration of the suggested active-set algorithm, the spherical inequality constraint will either be in the working set or not. If it is, then we iterate as described in Algorithm 2; otherwise we proceed with a generic (nonconvex) quadratic programming active-set algorithm [27, §16.8] [12,18]. We switch between the two algorithms when an iterate of the QP algorithm hits the spherical boundary or when the Lagrange Multiplier μ of the norm constraint in Algorithm 2 is negative and less than any other Lagrange Multiplier κ_i .

In the special the case where $r_{\min} = 0$, the norm constraint reduces to $\|x\|_2 \leq r_{\max}$. In this case, when a switch from Algorithm 2 to the QP Algorithm happens, the projected Hessian of the Lagrangian of $f(\cdot)$ at x_k is positive semidefinite, i.e.

$$Q^T(P + \mu I)Q \geq 0 \tag{27}$$

where $Q := [Q_1 \ q_2]$ is defined by the “thin” QR decomposition of $[\bar{A}_k^T \ x_k]$ and q_2 is an appropriate vector. Recall that for a switch to happen, we need $\mu < 0$; hence $Q^T P Q$ is positive definite. Thus, $Q_1^T P Q_1$, i.e. the projected Hessian of the next iteration that is to be handled by the generic QP algorithm, has at most one nonpositive eigenvalue [29, Corollary 4.1]. This means that even the popular, but less flexible, class of “inertia controlling” QP algorithms can be used as part of the suggested active-set algorithm for solving (P) when $r_{\min} = 0$.

3.2.1 Finding an initial point

Algorithm 2 is a primal active-set algorithm and, as such, it requires an initial feasible point. Unfortunately, finding such a point is, in general, NP-complete as we formally establish at the end of this subsection. Nevertheless, we can find a feasible point for (P) with standard tools, albeit in exponential worst-case complexity, as we proceed to show. Indeed, define the following problems:

$$\begin{array}{ll} \text{maximize } x^T x & \text{minimize } x^T x \\ \text{subject to } Ax \leq b & \text{subject to } Ax \leq b \\ \|x\|_\infty \leq r_{\min} & \end{array}$$

i.e. a convex minimization and a convex maximization problem. Denote with x_{\min}^* and x_{\max}^* the solutions to the convex minimization and maximization problems respec-

tively. Then, (P) is feasible iff $\|x_{\min}^*\|_2 \leq r_{\max}$ and $\|x_{\max}^*\|_2 \geq r_{\min}$, in which case a feasible point for (P) can be found by interpolating between x_{\min}^* and x_{\max}^* .

The convex minimization problem above can be solved in polynomial time with e.g. interior point or first order [32] methods. On the other hand, the convex maximization problem can have exponential worst-case complexity, but it can be solved to local or even global optimality with standard commercial solvers e.g. with IBM CPLEX.

Finally, we formally show that determining if (P) is feasible is NP-complete:

Proposition 4 *Determining if there is a solution to*

$$\begin{aligned} & \text{find } x \\ & \text{subject to } Ax \leq b \\ & r_{\min} \leq \|x\|_2 \leq r_{\max} \end{aligned} \tag{F}$$

is NP-complete.

Proof Determining if (F) has a solution is NP since we can easily check whether a candidate point x is feasible for (F). Furthermore, it can be decomposed into the following two independent problems: (i) “is there a point in the polytope $Ax \leq b$ such that $\|x\|_2 \leq r_{\max}$?” and (ii) “is there a point in the polytope $Ax \leq b$ such that $\|x\|_2 \geq r_{\min}$?” Problem (i) can be answered in polynomial time with e.g. interior point methods. Thus determining if (F) has a solution is reducible to problem (ii) which is known to be NP-complete [26, Problem 3]. \square

4 Applications and experiments

In this section we present numerical results of the suggested active-set algorithms on a range of numerical optimization problems. A Julia implementation of the algorithm, along with code for the generation of the results of this section is available at:

<https://github.com/oxfordcontrol/QPnorm.jl>

As described in the previous sections, the algorithm is based on a TRS solver, and a general (nonconvex) Quadratic Programming solver. We implemented these dependencies as separate packages that we also release publicly as listed below.

- TRS.jl: A Julia package for the computation of global and local-nonglobal minimizers of

$$\begin{aligned} & \text{minimize } \frac{1}{2}x^T Px + q^T x \\ & \text{subject to } \|x\| = r \text{ or } \|x\| \leq r \\ & Ax = b, \end{aligned} \tag{P}$$

essentially implementing in Julia the theoretical results of Sect. 2 and [1]. Available at:

<https://github.com/oxfordcontrol/TRS.jl>

- GeneralQP.jl: A Julia implementation of [13], i.e. an “inertia controlling” active-set solver for nonconvex, dense quadratic problems, with efficient

and numerically stable routines for updating QR decompositions of the working set and LDLt factorizations of the projected Hessian. Available at: <https://github.com/oxfordcontrol/GeneralQP.jl>
This solver is useful as a part of the suggested algorithm for solving (P) according to the remarks of Sect. 3.2.

For simplicity our implementations of the active-set algorithms are based on dense linear algebra that uses a QR factorization to compute/update a nullspace basis for every working set. Thus the presented results are limited to dense problems, except for Sect. 4.3 where the special structure of the constraints of (32) results in trivial, sparse QR factorizations of the nullspace bases.

Before presenting the results we will discuss some practical considerations regarding the suggested active-set algorithm. The eigenproblems (4) associated with the trust-region subproblems of each working set are solved with ARPACK [24]. In the occasional cases where ARPACK fails, a direct eigensolver is used. Finally, before solving the subproblem (SP) of every iteration, we perform a few projected gradient steps with the hope to quickly activate a new constraint.

We now proceed with the results, starting with the simplest case of dense random problems.

4.1 Random dense constant norm QPs

We compared the performance of our algorithm against the state-of-the-art non-linear solver `Ipopt` [33] with its default parameters. We use the `Julia` interface of `Ipopt`, `Ipopt.jl`, which exhibits negligible interface overhead times due to the excellent interfacing capabilities of `Julia` with `C++`.

We consider a set of these randomly generated problems with varying number of variables n . A feasible point for each problem instance is calculated for our Algorithm as described in Sect. 3.2.1. The time required to compute the initial feasible points is included in the subsequent results.

Figure 3 (left) shows the execution times. We observe that our algorithm is significantly faster than `Ipopt` by a factor of up to 50. Both of the solvers achieve practically identical objectives (relative difference less than 10^{-9}) in all of the problem instances, except in two cases where there is a considerable difference due to the fact that the solvers ended up converging in different local minimizers. Finally, Fig. 3 (right) shows the infeasibility of the returned solution, where we observe that our solver returns solutions of significantly smaller (i.e. better) infeasibility.

4.2 Computing search directions for sequential quadratic programming

Sequential Quadratic Programming (SQP) is a powerful and popular algorithm that aims to solve the problem

$$\begin{aligned} & \text{minimize } g(x) \\ & \text{subject to } h(x) \leq 0, \end{aligned} \tag{28}$$

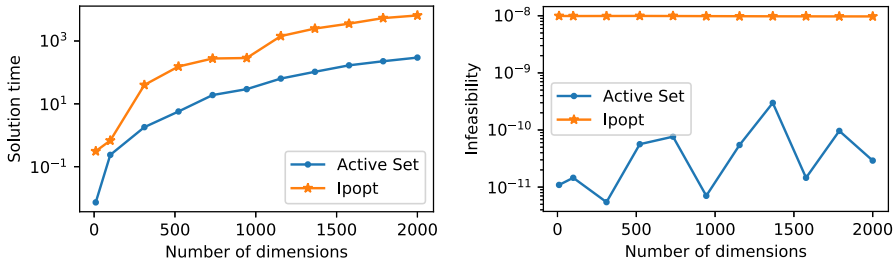


Fig. 3 Timing results (left) and maximum feasibility violation (right) of the solution of random problems with varying size n generated with the following parameters $P = \text{Symmetric}(\text{randn}(n, n))$, $q = \text{randn}(n)$, $A = \text{randn}(1.5n, n)$, $b = \text{randn}(n)$ and $r = 100$. Hardware used: Intel E5-2640v3, 64GB memory. The maximum feasibility violation is defined as $\max((Ax^* - b)_1, \dots, (Ax^* - b)_m, \frac{1}{2} \|x^*\|_2^2 - r^2, 0)$ for a solution x^*

where $g : \mathbb{R}^n \mapsto \mathbb{R}$ and $h : \mathbb{R}^n \mapsto \mathbb{R}^m$ are general (nonlinear) functions. At every iteration, SQP minimizes the nonlinear objective over some *search directions*. These search directions can, for example, be obtained by the optimization of a quadratic approximation of the original objective subject to some linear approximation of the original constraints. It is common to introduce a norm constraint $\|\delta x\|_2 \leq r$ that ensures that the solution will be inside a confidence or *trust* region of this approximation thus resulting in a problem of the form (P). This is particularly useful at points where the Hessian of the nonlinear objective is singular or indefinite as the search directions computed by the respective QPs might be unrealistically large or unbounded.

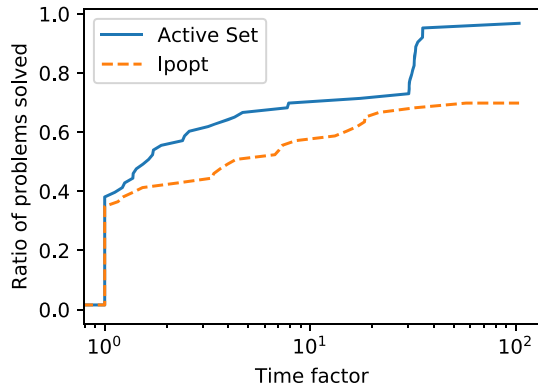
Typically, SQP algorithms solve problems simpler than (P) by e.g. not including the inequality constraints (P) in explicit form for the calculation of the search directions [27, §18.2]. We will show, however, that our solver is capable of computing search directions from (P) directly, without the need for these simplifications.

We demonstrate this on all the problems of the CUTEst collection [17] that have linear constraints and number of variables less than 2000. For each of these problems we consider the quadratic approximation

$$\begin{aligned}
 &\text{minimize } f(\delta x) := \frac{1}{2} \delta x^T \nabla^2 g(x_0) \delta x + \delta x^T \nabla g(x_0) + c \\
 &\text{subject to } \|\delta x\|_2^2 \leq 1 \\
 &\qquad \qquad \nabla h(x_0) \delta x + h(x_0) \leq 0
 \end{aligned} \tag{29}$$

***where x_0 is the closest feasible point (in the 2-norm) to the initial point suggested by CUTEst, which we compute with GUROBI [19], and $\delta x := x - x_0$. We discard problems where GUROBI fails to calculate an initial feasible point. Furthermore, we do not consider problems for which (29) is convex, since these problems can be solved to global optimality in polynomial time with standard solvers such as MOSEK or COSMO.jl [11]. Finally, some problems in the CUTEst collection are parametric; if the default parameters lead to a problem with more than 2000 variables then we choose a parameter that, if possible, leads to a number of variables close to, but not exceeding, 2000. Table 4 list all the 63 problems considered, along with any non-default parameters.

Fig. 4 Performance graph of the suggested algorithm against I_{popt} on calculating SQP search directions on all 63 problems of the CUTEst library listed in Table 4. Matrices are passed to I_{popt} in a sparse format. Hardware used: Intel E5-2640v3, 64GB memory



We compare the performance of our algorithm against I_{popt} with its default options. For some problems, I_{popt} terminates without indication of failure but returns a low quality solution. We consider a problem as “solved” using the same criteria as [8]. That is, we require that the overall error of the KKT conditions is less than 10^{-4} , defined as

$$\text{error} := \max(\epsilon_{\text{primal inf}}, \epsilon_{\text{dual inf}}, \epsilon_{\text{stationarity}}, \epsilon_{\text{compl}}) \tag{30}$$

with

$$\begin{aligned} \epsilon_{\text{primal inf}} &:= \max\left(0, \delta x^T \nabla h(x_0) + h(x_0), \|\delta x\|_2 - 1\right) \\ \epsilon_{\text{dual inf}} &:= -\min(0, \kappa_1, \dots, \kappa_m, \mu) \\ \epsilon_{\text{stationarity}} &:= \left\| \nabla^2 g(x_0) \delta x + \nabla g(x_0) + \kappa^T \nabla h(x_0) + 2\mu \delta x \right\|_{\infty} \\ \epsilon_{\text{compl}} &:= \max(\{\min(\kappa_i, \|\nabla h(x_0) \delta x + h(x_0)\|_i)\}, \min(\mu, \|\delta x\|_2^2 - 1)), \end{aligned}$$

and κ, μ are the Lagrange multipliers corresponding to the linear inequality and norm constraints of (29). Figure 4 shows the performance profile for the problems considered. The performance profile suggests that our algorithm significantly outperforms I_{popt} on this set of problems especially in reliability, in the sense that we consistently obtain high quality solutions. Note that unlike the dense implementation of our solver, I_{popt} can exploit the sparsity of the problems efficiently. Further speedups can be brought to our algorithm with a sparse implementation which might allow its use in large scale sparse problems. Moreover, as an active-set solver, our algorithm can efficiently exploit warm starting, which might be highly desirable in SQP where repeated solution of problems of the form (29) is required as part of the SQP procedure for minimizing (28). This is in contrast to I_{popt} whose Interior Point nature makes warm starting very difficult to exploit.

4.3 Sparse principal component analysis

Principal Component Analysis (PCA) is a standard, widely used dimensionality reduction algorithm. It has applications in numerous fields, including statistics, machine learning, bioinformatics, genetics, meteorology and others. Given a $k \times n$ data matrix D that consists of k points of n variables, PCA suggests a few linear combinations of these variables, which we call *principal vectors*, that explain as much variance of the data as possible. Standard PCA amounts to the following problem

$$\begin{aligned} &\text{maximize } x^T \Sigma x \\ &\text{subject to } \|x\|_2 \leq 1 \end{aligned}$$

where Σ is the covariance matrix of the data. We will assume for the rest of this section that the data matrix D is “centered”, i.e. that it has column-wise zero mean. We can then consider Σ to be the empirical covariance matrix $\Sigma := \frac{D^T D}{k-1} \in \mathbb{S}_+^n$. The above problem is essentially an eigenvalue problem on Σ or a singular value problem on D and, as such, can be solved with standard linear algebra tools.

In general, each principal vector is a linear combination of all the variables, i.e. typically all components of the principal vectors are non-zero. This can pose an issue of interpretability of the reduced dimensions, as it is often desirable to express the principal vector as a combination of a few variables, especially when the variables are associated with a user interpretable meaning. To alleviate this problem, sparsity has to be enforced in the original PCA problem, resulting in a new optimization problem that aims to identify a small set of variables the linear combination of which will hopefully still be able to explain a significant proportion of the variance of the data.

Unfortunately, enforcing sparsity in PCA results in a combinatorial optimization problem. Various remarkably efficient and scalable heuristics have been suggested to avoid this intractability [23,34]. We will focus on one of the most popular heuristics that is based on a lasso type constraint originally introduced by [22], resulting in the following optimization problem:

$$\begin{aligned} &\text{minimize } f(x) := -x^T \Sigma x \\ &\text{subject to } \|x\|_2 \leq 1 \\ &\qquad\qquad \|x\|_1 \leq \gamma, \end{aligned} \tag{31}$$

where $x \in \mathbb{R}^n$ is the decision variable and $\gamma \geq 1$ is a parameter that controls the sparsity of the solution.

We will now show how Problem (31) can be solved with our algorithm. Note that (31) is not in the standard norm-constrained form (P) that we address. However, we show in the “Appendix” that it is equivalent to

$$\begin{aligned} &\text{minimize } g(w) := -(w_1 - w_2)^T \Sigma (w_1 - w_2) \\ &\text{subject to } \|w\|_2^2 \leq 1 \\ &\qquad\qquad \mathbf{1}^T w \leq \gamma \\ &\qquad\qquad w \geq 0, \end{aligned} \tag{32}$$

where $[w_1^T \ w_2^T] := w^T \in \mathbb{R}^{2n}$, in the sense that (31) and (32) have the same optimal value and every minimizer of (32) \bar{w} defines a minimizer $\bar{x} = \bar{w}_1 - \bar{w}_2$ for (31). Also, the optimal \bar{w} has the same cardinality as the respective optimal $\bar{x} = \bar{w}_1 - \bar{w}_2$.

The reader might think that solving (32) with an active-set method requires a solve time that is polynomial in the number of variables, thus limiting its scalability as compared to simpler gradient based methods that are scalable but might have inferior accuracy and, potentially, weaker convergence guarantees. This is not true however, since our algorithm takes advantage of the sparsity of the iterates. Indeed, excluding the lasso constraint, the working set \mathcal{W}_k of Algorithm 2 applied in Problem (32) can be interpreted as the set of variables w_i that are fixed to zero. Thus every subproblem can be trivially reduced to k dimensions, where k is the number of variables w_i that are not included in the current working set. This property is particularly attractive when the user is interested in a very sparse solution and has been recognized in the literature on ℓ_1 -penalized active-set algorithms [3, §6].

We will demonstrate the scalability and flexibility of our algorithm by applying it in a large “bag of words” dataset from articles of the New York Times newspaper obtained by the University of California, Irvine (UCI) Machine Learning Repository⁷. In this dataset, every row of the associated data matrix D corresponds to a word used in any of the documents and every column to a document. The entries of the matrix are the number of occurrences of a word in a document. The dataset contains $m \approx 300,000$ articles that consist of $n \approx 102,660$ words (or keywords) with ≈ 70 million non-zeros in the dataset matrix D . We seek to calculate a few sparse principal vectors, each of which will hopefully correspond to a user-interpretable category, like politics, economics, the arts etc.

The usefulness of Sparse PCA on this dataset has been already demonstrated in the literature. Both [28,35] have generated a set of five sparse principal vectors, each of which is a linear combination of five words. The results of [28] are listed in Table 1. The five resulting principal vectors have distinct, user interpretable associated meanings that correspond to sports, economics, politics, education, and US foreign policy. The resulting words can be used to create user interpretable categories of the documents [35]. However, this interpretability is in general lost when we increase the cardinality of the sparse principal vectors, i.e. when we require more representative words for each category. For example, the result of applying TPower [34] (a state-of-the-art algorithm for Sparse PCA) in the NYtimes dataset to generate a sparse principal vector that consists of 30 words, listed in Table 2, does not have a distinct associated meaning. For reference, TPower runs in 4.91 s on a standard Laptop with an Intel i7-5557U CPU and 8 GB of memory and has a resulting variance of 58.79. Another popular algorithm, ℓ_1 -GPower, runs in 355.86 s and results in a variance of 56.66. The results of our active-set algorithm are better than ℓ_1 -GPower⁸ and comparable (slightly inferior) to TPower: our algorithm runs in 5.22 s and has a resulting variance equal to 58.23. However, our algorithm is more general, in the sense that it can solve

⁷ Available at <http://archive.ics.uci.edu/ml/datasets/Bag+of+Words>.

⁸ Note that ℓ_1 penalized algorithms (such as our algorithm and ℓ_1 -GPower) do not allow for the specification of the desired sparsity directly. Instead, their sparsity is controlled by a penalty parameter. We find a penalty parameter that results in 30 nonzeros via means of a binary search, while exploiting warm starting of the solvers.

Table 1 Results of [28] on applying sparse PCA on the NYtimes dataset

Vector #1	Vector #2	Vector #3	Vector #4	Vector #5
game	companies	campaign	children	attack
play	company	president	program	government
player	million	zzz_al_gore	school	official
season	percent	zzz_bush	student	zzz_u_s
team	stock	zzz_george_bush	teacher	zzz_united_states

Each column lists words that correspond to nonzero entries in the respective principal vectors. The resulting principal vectors have clear associated meanings that correspond to sports, economics, elections, education, and the US

Table 2 Results of applying sparse PCA on the NYtimes dataset with TPower, allowing weights with mixed signs

Words	Weights	Words (continued)	Weights (continued)
teamed	-0.45197	hitch	-0.083649
gameday	-0.3698	playland	-0.081834
seasonable	-0.34698	leaguer	-0.07607
playful	-0.28105	nightcap	-0.072675
playa	-0.24374	percentage	0.27263
gamesmanship	-0.20957	companywide	0.23443
coachable	-0.13991	companion	0.14581
pointe	-0.13315	stockade	0.12727
runaround	-0.11266	marketability	0.12368
win95	-0.1104	millionaire	0.11005
yardage	-0.10503	governmental	0.092267
guzzle	-0.093052	billionaire	0.090433
player	-0.089249	businesses	0.076009
wonder	-0.086247	zzz_bush_administration	0.072731
ballad	-0.085514	fundacion	0.070655

Each column lists words that correspond to nonzero entries in the generated principal vector. Note that increasing the cardinality of the sparse principal vector has severely affected the interpretability of the principal vector as compared to the results of Table 1 that have clear associated meaning

any problem of the form (P) which, as we proceed to show, allows for variants of (31) that can generate principal vectors with more meaningful interpretation.

Indeed, improving interpretability of the results of Sparse PCA can be achieved by the incorporation of additional constraints in the underlying optimization problem. Unfortunately, most of the algorithms that are specific for sparse PCA do not allow for the presence of extra constraints. An exception is perhaps GPower of [23], which is a generic framework for minimizing a concave function f over a compact subset \mathcal{F} . The iterates of GPower are generated as

$$\begin{aligned}
 x_{k+1} = \operatorname{argmin} \quad & x^T f'(x_k) \\
 \text{subject to } & x \in \mathcal{F}
 \end{aligned}
 \tag{33}$$

Table 3 Results of applying non-negative sparse PCA on the NYtimes dataset with the suggested active-set algorithm

Vector #1	Vector #2	Vector #3	Vector #4	Vector #5
team	percent	zzz_al_gore	official	school
game	company	zzz_bush	government	student
season	million	zzz_george_bush	attack	children
player	companies	campaign	zzz_u_s	program
play	stock	president	zzz_united_states	teacher
games	market	election	zzz_bush	parent
point	billion	political	military	high
coach	business	zzz_white_house	zzz_american	percent
run	fund	republican	palestinian	public
win	analyst	voter	zzz_taliban	education
yard	money	vote	zzz_afghanistan	kid
guy	zzz_enron	presidential	terrorist	district
won	firm	democratic	war	test
played	investor	tax	leader	family
hit	sales	zzz_republican	zzz_israel	college
ball	industry	administration	group	home
playing	plan	zzz_clinton	bin	child
league	cost	ballot	laden	class
left	investment	votes	country	group
home	quarter	plan	administration	job
fan	deal	democrat	security	boy
shot	financial	zzz_washington	forces	money
field	customer	zzz_congress	zzz_israeli	help
playoff	economy	poll	zzz_pakistan	mother
night	chief	support	american	system
goal	price	candidate	zzz_washington	friend
final	executive	zzz_florida	nation	girl
start	growth	candidates	terrorism	private
quarterback	earning	governor	weapon	voucher
football	share	vice	foreign	grade

Each column lists words that correspond to nonzero entries in the respective principal vectors. The words are sorted in decreasing magnitude of the respective weights. Given a set of computed principal vectors $\{x_i\}$, the next principal vector is computed on the deflated data matrix $D - \sum (Dx_i)x_i^T$. The algorithm is initialized with the 30 most positive, or most negative, entries (depending on which one leads to higher explained variance) of the first singular vector of the respective matrix. Furthermore we perform a standard post-processing polishing step at each resulting principal vector, as described in [23, §4.2]. Note that, unlike Table 2 our algorithm results in principal vectors that have clear, user interpretable associated meanings that corresponds to sports, economics, politics, education, and US foreign policy. These words could be used for document classification [35]

where $f'(\cdot)$ is a subgradient of $f(\cdot)$. [23] showed that for the case of Sparse PCA of (31) the above optimization problem reduces to two matrix multiplications with D and D^T . However, in the presence of general linear inequality constraints, solving (33) can be at least as hard as solving a linear program (LP), which is obviously considerably more computationally demanding than two matrix multiplications.

In contrast, due to the generality of our approach, imposing any linear constraint is straightforward for our algorithm. To demonstrate this, we impose a non-negativity constraint on the elements of the principal vector to avoid cases where the principal vectors consist of elements that have opposing meanings which happens, for example, in the results of TPower listed in Table 2. Table 3 lists five sparse principal vectors obtained with this approach which take 24.37 s to compute on a standard Laptop with an Intel i7-5557U CPU and 8GB of memory. Note that, unlike the results of Table 2, the principal vectors of Table 3 have a clear associated meaning that can be labeled as sports, economics, politics, US foreign policy and education. These words could for example be used for organizing the documents in a user interpretable way [35].

5 Conclusion

This paper introduced an active-set algorithm for the solution of quadratic functions subject to linear constraints and a single norm constraint. The suggested algorithm is based on repeated solutions of the TRS, for which we derived novel theoretical results regarding its local-nonglobal minimizer. The usefulness of the suggested algorithm was demonstrated in a range of real world experiments.

Appendix

In this Section we show that the suggested active-set algorithm of Sect. 3.2 can perform ℓ_1 sparse PCA, that is, to solve the problem

$$\begin{aligned} &\text{minimize } f(x) := -x^T \Sigma x \\ &\text{subject to } \|x\|_2 \leq 1 \\ &\qquad\qquad \|x\|_1 \leq \gamma. \end{aligned}$$

Theorem 3 *Problems (31) and (32) have the same optimal value and every (local) minimizer \bar{w} of (32) defines a (local) minimizer $\bar{x} = \bar{w}_1 - \bar{w}_2$ for (31).*

Proof First note that both Problems (31) and (32) have a minimizer as their objective is smooth and their feasible set compact. For any feasible x of (31), the choice

$$w_1 = x_+ \quad \text{and} \quad w_2 = -x_-,$$

where x_+ and x_- are defined element-wise as follows

$$x_+ := \max(x, 0) \quad \text{and} \quad x_- := \min(x, 0),$$

is feasible for (32) since $\|w_1\|_2^2 + \|w_2\|_2^2 = \|x\|_2^2 = 1$ and $\mathbf{1}^T w_1 + \mathbf{1}^T w_2 = \|x\|_1 \leq \gamma$, while $f(x) = g([w_1; w_2])$ since

$$-(w_1 - w_2)^T \Sigma (w_1 - w_2) = -x^T \Sigma x.$$

Thus the optimal value of (32) is less than that of (31).

Moreover, for any minimizer (\bar{w}_1, \bar{w}_2) of (32), choosing $\bar{x} = \bar{w}_1 - \bar{w}_2$ gives $f(\bar{x}) = g([\bar{w}_1; \bar{w}_2])$ with \bar{x} feasible for (31) as we proceed to show. Indeed, note that $\bar{w}_1^T \bar{w}_2 \geq 0$ and $\|\bar{w}_1\|_2^2 + \|\bar{w}_2\|_2^2 = 1$ thus

$$\|\bar{x}\|_2^2 = \|\bar{w}_1\|_2^2 + \|\bar{w}_2\|_2^2 - 2\bar{w}_1^T \bar{w}_2 \leq \|\bar{w}_1\|_2^2 + \|\bar{w}_2\|_2^2 = 1 \tag{34}$$

and

$$\|\bar{x}\|_1 = \|\bar{w}_1 - \bar{w}_2\|_1 \leq \|\bar{w}_1\|_1 + \|\bar{w}_2\|_1 \leq \gamma. \tag{35}$$

Assume that \bar{x} is not a minimizer for (31). Then, there exists an \tilde{x} with $\|\tilde{x} - \bar{x}\|_\infty$ arbitrarily small such that $f(\tilde{x}) > f(\bar{x})$. Thus $(\tilde{w}_1, \tilde{w}_2) := (\tilde{x}_+, -\tilde{x}_-)$ is feasible for (32) with $g([\tilde{w}_1; \tilde{w}_2]) = f(\tilde{x}) < f(\bar{w}) = g([\bar{w}_1; \bar{w}_2])$ and

$$\left\| \begin{bmatrix} \tilde{w}_1 - \bar{w}_1 \\ \tilde{w}_2 - \bar{w}_2 \end{bmatrix} \right\|_\infty = \left\| \begin{bmatrix} \tilde{x}_+ - \bar{x}_+ \\ -(\tilde{x}_- - \bar{x}_-) \end{bmatrix} \right\|_\infty \leq \|\tilde{x} - \bar{x}\|_\infty \tag{36}$$

i.e. $(\tilde{w}_1, \tilde{w}_2)$ is arbitrarily close to (\bar{w}_1, \bar{w}_2) . This contradicts the assumption that (\bar{w}_1, \bar{w}_2) is a local minimizer for (32).

Finally, since the optimal value of (32) is less than the one of (31) and the global optimum (\bar{w}_1, \bar{w}_2) of (32) defines a feasible \bar{x} for (31) with $g([\bar{w}_1; \bar{w}_2]) = f(\bar{x})$ we conclude that the optimal values of (31) and (32) coincide. \square

Furthermore, the following complementarity condition holds which implies that the sparsity of a solution \bar{w} of (32) is equal to the sparsity of the solution $\bar{x} = \bar{w}_1 - \bar{w}_2$ for Problem (31).

Lemma 7 Every (local) minimizer \bar{w} of (32) has $\bar{w}_1^T \bar{w}_2 = 0$.

Proof Assume the contrary, i.e. that $\bar{w}_1^T \bar{w}_2 \neq 0$. Due to feasibility of $\bar{w} = [\bar{w}_1; \bar{w}_2]$ we have $\bar{w} \geq 0$. Thus, there exists an $i \in \{1, \dots, n\}$ such that $\bar{w}_{(i)}, \bar{w}_{(i+n)} > 0$. Define \tilde{w} element-wise as

$$\tilde{w}_{(k)} = \begin{cases} \bar{w}_{(k)} - \epsilon, & k = i, i + n \\ \bar{w}_{(k)}, & \text{otherwise} \end{cases} \tag{37}$$

for $\epsilon > 0$ sufficiently small so that $\tilde{w} > 0$. We will show that for a sufficiently small, positive α the point $(1 + \alpha)\tilde{w}$ is feasible and of smaller objective value than the local minimizer \bar{w} while $\|\bar{w} - (1 + \alpha)\tilde{w}\|$ can become arbitrarily small for an appropriate choice of α and ϵ . This will be a contradiction of (local) optimality of \bar{w} .

Table 4 Comparison of the suggested active-set algorithm against t_{ipopt} in calculating SQP search directions on 63 nonconvex problems of the CUTEst library

Problem	Parameter	n	m	t	$t_{\text{ipopt}}^{\text{dense}}$	$t_{\text{ipopt}}^{\text{sparse}}$	δf^*	$\delta f_{\text{ipopt}}^*$	Error	Error $_{\text{ipopt}}$
AVION2		49	113	4.444538E-02	4.670914E-02	1.723294E-02	-3.645564E+02	-3.645789E+02	1.479577E-10	5.426607E-04
BLOCKQP1	$N = 100$	210	521	7.184425E-03	1.308548E+00	5.072931E-02	-6.745688E+00	-6.745688E+00	9.992007E-16	9.472623E-09
BLOCKQP2	$N = 100$	210	521	1.333553E-02	1.234218E+00	5.287435E-02	-6.745688E+00	-6.745688E+00	6.661338E-16	9.374904E-09
BLOCKQP3	$N = 100$	210	521	1.551006E-02	7.109685E-01	3.531169E-02	-4.018526E+00	-4.018526E+00	2.164935E-15	9.224197E-09
BLOCKQP4	$N = 100$	210	521	9.190950E-03	7.812505E-01	4.003907E-02	-4.018526E+00	-4.018526E+00	1.015854E-14	8.997483E-09
BLOCKQP5	$N = 100$	210	521	9.125178E-03	7.250970E-01	3.328339E-02	-4.018526E+00	-4.018526E+00	1.160183E-14	9.325761E-09
BLOWEYA	$N = 100$	202	304	2.057082E-01	6.153741E-01	4.804298E-02	-2.902540E-03	-2.903560E-03	1.103075E-12	4.457334E-05
BLOWEYB	$N = 100$	202	304	3.518280E-01	5.780202E-01	4.473963E-02	-2.574329E-03	-2.575152E-03	1.452606E-12	4.834189E-05
BLOWEYC	$N = 100$	202	304	3.704172E-02	5.043310E-01	4.545440E-02	-2.714783E-03	-2.715942E-03	3.039952E-15	3.156699E-05
EQC		9	21	2.374110E-04	3.362073E-02	3.607733E-02	-1.653931E+00	-4.903961E+03	0.000000E+00	5.729870E+00
EXPFITA		5	22	9.989510E-04	5.855086E-03	7.309660E-03	-1.529569E+01	-1.529569E+01	7.813195E-12	1.159203E-07
EXPFITB		5	102	8.978320E-04	1.018311E-02	1.349212E-02	-7.615627E+01	-7.615627E+01	1.547296E-11	2.935301E-06
EXPFITC		5	502	2.641202E-03	4.953891E-02	6.193263E-02	-3.810092E+02	-3.810092E+02	8.583134E-13	1.278743E-04
FERRISDC	$n = 100$	1100	2110	1.632799E+01	4.871856E+02	2.109886E+02	-2.314228E-01	-2.314213E-01	9.975523E-09	4.939870E-05
GOULDQP1		32	81	4.174306E-03	2.652730E-02	1.418419E-02	-8.790888E+02	-8.790888E+02	9.783285E-13	7.851076E-05
HIMMELBJ		45	61	4.420100E+00	3.329924E-02	2.115145E-02	-7.390332E+01	-7.391036E+01	3.586903E+00	1.144493E-05
HS105		8	17	1.351519E-03	4.024936E-03	4.477112E-03	-8.081024E+01	-8.081024E+01	1.452172E-12	2.669379E-07
HS24		2	5	3.546430E-04	6.277818E-03	6.520859E-03	-1.311008E-01	-1.311008E-01	2.220446E-16	1.623253E-08
HS36		3	7	2.946660E-04	5.016933E-03	5.453507E-03	-1.832051E+02	-1.832051E+02	7.105427E-14	9.974060E-09
HS37		3	8	2.950220E-04	5.691497E-03	4.989736E-03	-1.832051E+02	-1.832051E+02	9.947598E-14	9.974059E-09
HS41		4	9	1.630200E-04	4.818620E-03	5.279881E-03	-1.562500E-02	-1.562500E-02	5.551115E-17	1.329137E-08

Table 4 continued

Problem	Parameter	n	m	t	$t_{\text{ipopt}}^{\text{dense}}$	$t_{\text{ipopt}}^{\text{sparse}}$	δf^*	$\delta f_{\text{ipopt}}^*$	Error	Error _{ipopt}
HS44		4	10	4.904870E-04	8.343742E-03	6.473175E-03	-1.304760E+00	-1.304760E+00	2.442491E-15	3.577725E-08
HS44NEW		4	10	6.557700E-04	4.320227E-03	4.416368E-03	-1.847553E+00	-1.847553E+00	2.220446E-15	8.037391E-09
HS55		6	14	2.517480E-04	5.743211E-03	5.443675E-03	-1.000000E+00	-2.236068E+00	4.000000E+00	5.616085E-05
LEUVEN7		360	1605	5.521594E-01	6.015317E+00	3.374051E-01	-4.858191E+02	-4.858192E+02	2.160050E-12	5.712026E-05
LINCONT		1257	2435	6.700989E+01	4.709799E+02	6.443721E-01	-1.113381E-18	-1.558091E-05	2.899090E-10	3.035288E-07
MPC10		1530	3571	2.982143E+01	5.949118E+02	9.300218E-01	-3.155428E+03	-3.155438E+03	1.636501E-11	9.900918E-05
MPC11		1530	3571	3.138670E+01	6.484969E+02	9.605708E-01	-3.155339E+03	-3.155349E+03	2.258621E-11	1.353352E-04
MPC12		1530	3571	3.015642E+01	6.625879E+02	9.935781E-01	-3.155068E+03	-3.155078E+03	1.451272E-11	9.204114E-05
MPC13		1530	3571	3.139020E+01	6.049481E+02	8.909481E-01	-3.154808E+03	-3.154819E+03	1.516229E-11	1.890977E-04
MPC14		1530	3571	3.036349E+01	6.181861E+02	9.324590E-01	-3.154810E+03	-3.154821E+03	2.793515E-11	1.047976E-04
MPC15		1530	3571	3.230341E+01	5.987057E+02	9.142717E-01	-3.153846E+03	-3.153859E+03	2.210518E-11	1.069349E-04
MPC16		1530	3571	3.241830E+01	6.266897E+02	9.186500E-01	-3.153841E+03	-3.153853E+03	1.616848E-11	1.070100E-04
MPC2		1530	3571	2.824377E+01	5.600366E+02	8.767303E-01	-3.157144E+03	-3.157151E+03	1.161907E-11	1.056836E-04
MPC3		1530	3571	2.783808E+01	6.103022E+02	9.163411E-01	-3.157321E+03	-3.157329E+03	1.224183E-11	1.047022E-04
MPC4		1530	3571	3.018797E+01	6.347866E+02	9.388815E-01	-3.156964E+03	-3.156972E+03	1.548417E-11	1.033535E-04
MPC5		1530	3571	2.943601E+01	6.344659E+02	9.331457E-01	-3.157075E+03	-3.157083E+03	1.574789E-11	1.044170E-04
MPC6		1530	3571	3.035928E+01	5.720201E+02	9.081594E-01	-3.151828E+03	-3.151837E+03	1.558678E-11	1.905436E-04
MPC7		1530	3571	2.897083E+01	6.455158E+02	9.593955E-01	-3.155357E+03	-3.155366E+03	1.857497E-11	1.770294E-04
MPC8		1530	3571	2.761454E+01	5.973750E+02	9.020797E-01	-3.155408E+03	-3.155418E+03	1.359729E-11	1.503306E-04
MPC9		1530	3571	2.949440E+01	6.456766E+02	9.470979E-01	-3.155392E+03	-3.155402E+03	1.521752E-11	1.044509E-04
NASH		72	138	2.023284E-02	9.083129E-02	2.347952E-02	-1.335068E-16	-2.562575E-05	1.101341E-13	2.652532E-07

Table 4 continued

Problem	Parameter	n	m	t	$f_{\text{ipopt}}^{\text{dense}}$	$f_{\text{ipopt}}^{\text{sparse}}$	δf^*	$\delta f_{\text{ipopt}}^*$	Error	Error _{ipopt}
NCVXQP1	$N = 1000$	1000	2500	3.468481E-01	5.037756E+01	2.530511E-01	-7.824962E+04	-7.824962E+04	9.094947E-12	3.340764E-07
NCVXQP2	$N = 1000$	1000	2500	3.615064E-01	5.018531E+01	2.535988E-01	-5.736680E+04	-5.736680E+04	1.273293E-11	2.509192E-07
NCVXQP3	$N = 1000$	1000	2500	3.681914E-01	4.789657E+01	2.396625E-01	-4.998258E+04	-4.998258E+04	1.273293E-11	2.732257E-07
NCVXQP4	$N = 1000$	1000	2250	3.495667E-01	3.453693E+01	2.044327E-01	-7.824962E+04	-7.824962E+04	5.184120E-11	3.340765E-07
NCVXQP5	$N = 1000$	1000	2250	3.553980E-01	3.450249E+01	2.057896E-01	-5.736680E+04	-5.736680E+04	1.714398E-09	2.509195E-07
NCVXQP6	$N = 1000$	1000	2250	3.642895E-01	3.295662E+01	1.933023E-01	-4.998258E+04	-4.998258E+04	5.820766E-11	2.732277E-07
NCVXQP7	$N = 1000$	1000	2750	3.584502E-01	7.042072E+01	2.941084E-01	-7.824962E+04	-7.824962E+04	5.129550E-10	3.340764E-07
NCVXQP8	$N = 1000$	1000	2750	3.669888E-01	7.371150E+01	2.933839E-01	-5.736680E+04	-5.736680E+04	8.305506E-09	2.509189E-07
NCVXQP9	$N = 1000$	1000	2750	3.835907E-01	6.705865E+01	2.787119E-01	-4.998258E+04	-4.998258E+04	2.119577E-08	2.732247E-07
PENTAGON		6	15	1.087885E-03	9.631484E-03	9.260053E-03	-2.699815E-02	-2.66403E-02	8.267692E-15	2.092105E-06
QC		9	22	2.120020E-04	1.348843E-02	1.216467E-02	-1.114234E+02	-1.114236E+02	0.000000E+00	1.875249E-05
QCNEW		9	21	2.276460E-04	1.657338E-02	1.513761E-02	-1.653931E+00	-4.112503E+01	0.000000E+00	1.990100E-04
QPNBLEND		83	157	1.501935E-02	6.937614E-02	2.064464E-02	-9.282172E-03	-9.283588E-03	2.504663E-13	4.679297E-05
QPNBOEII		384	891	6.116935E-01	1.115413E+01	1.653756E-01	-4.366064E+01	-4.366092E+01	7.786588E-12	8.287324E-05
QPNBOEIJ		143	363	4.073949E-02	1.011505E+00	6.225512E-02	-1.003864E+03	-1.004448E+03	1.224766E-09	8.420519E-05
QPNSTAIR		467	905	7.413311E-01	1.500511E+01	2.324631E-01	-1.150055E+02	-1.150130E+02	4.352719E-11	2.310892E-04
SOSQP1	$N = 1000$	2000	5001	1.025841E+00	3.387020E+02	2.187173E-01	-5.000000E-01	-5.000000E-01	3.108624E-14	4.572343E-09
SOSQP2	$N = 1000$	2000	5001	1.189768E+00	2.171972E+03	1.061519E+00	-5.000000E-01	-5.000000E-01	4.940492E-15	4.988599E-09
STATIC3		434	240	9.991469E-01	7.633038E-01	5.708679E-02	-1.473959E+04	-1.473959E+04	4.599613E-06	1.455835E-04
STNQPI	$P = 10$	1025	2560	3.909290E-01	5.022181E+01	1.634204E-01	-1.629919E+03	-1.629919E+03	2.557954E-13	1.049186E-08
STNQP2	$P = 10$	1025	2560	3.919327E-01	4.356643E+01	1.601439E-01	-1.629919E+03	-1.629919E+03	4.547474E-13	1.047931E-08

These include all the problems with linear inequality constraints, at most $n = 2000$ variables and an indefinite Hessian at the initial point. The initial point is computed with GURQBI as the nearest feasible point (in the two norm) to the initial point suggested by CURTEST. The column m denotes the number of inequality constraints, while columns t , δf^* , error denote the solution time (in seconds), the difference in objective value $f(\delta x^*) - f(0)$ achieved and the residual error defined by (30) of the suggested active-set solver. Similarly, $\delta f_{\text{ipopt}}^*$, error_{ipopt} denote respectively identical metrics for Ipopt. $f_{\text{ipopt}}^{\text{sparse}}$ ($f_{\text{ipopt}}^{\text{dense}}$) denotes Ipopt's timings when the matrices are passed in sparse (dense) format. Hardware used: Intel E5-2640v3, 64GB memory

Indeed, note that $\mathbf{1}^T \tilde{w} < \mathbf{1}^T \bar{w} \leq \gamma$ and $\|\tilde{w}\|_2 < \|\bar{w}\|_2 \leq 1$, thus $(1+\alpha)\tilde{w}$ is feasible for any $\alpha > 0$ sufficiently small. Moreover, $g(\tilde{w}) = g(\bar{w})$ giving $g((1+\alpha)\tilde{w}) = (1+\alpha)^2 g(\bar{w}) < g(\bar{w})$, since $g(\bar{w}) < 0$ as otherwise \bar{w} would be a global maximizer of (32). \square

Finally, we present detailed comparison results of our algorithm against `Ipropt` on the problems described in Sect. 4.2.

Acknowledgements We are grateful to the detailed feedbacks from an anonymous reviewer that significantly improved the quality of this document. Furthermore, we thank Nick Trefethen for providing valuable comments. This work was supported by the EPSRC AIMS CDT Grant EP/L015987/1 and Schlumberger. We would also like to acknowledge the support of the NII International Internship Program, which funded the first author's visit to NII, during which this collaboration was launched. We thank Ron Estrin for valuable discussions. We used the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work. <https://doi.org/10.5281/zenodo.22558>.

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. Adachi, S., Iwata, S., Nakatsukasa, Y., Takeda, A.: Solving the trust-region subproblem by a generalized eigenvalue problem. *SIAM J. Optim.* **27**(1), 269–291 (2017)
2. Ahlfors, L.: *Complex Analysis: An Introduction to the Theory of Analytic Functions of One Complex Variable*, vol. 2. McGraw-Hill, New York (1966)
3. Bach, F., Jenatton, R., Mairal, J., Obozinski, G.: Optimization with sparsity-inducing penalties. *Found. Trends Mach. Learn.* **4**(1), 1–106 (2012)
4. Beck, A., Vaisbourd, Y.: Globally solving the trust region subproblem using simple first-order methods. *SIAM J. Optim.* **28**(3), 1951–1967 (2018)
5. Boufounos, P.T., Baraniuk, R.G.: 1-Bit compressive sensing. In: 2008 42nd Annual Conference on Information Sciences and Systems, pp. 16–21 (2008)
6. Choi, S., Paige, C., Saunders, M.: MINRES-QLP: a Krylov subspace method for indefinite or singular symmetric systems. *SIAM J. Sci. Comput.* **33**(4), 1810–1836 (2011)
7. Conn, A.R., Gould, N.M., Toint, P.: *Trust-Region Methods*. Society for Industrial and Applied Mathematics, Philadelphia (2000)
8. Ferreau, H.J., Kirches, C., Potschka, A., Bock, H.G., Diehl, M.: qpOASES: a parametric active-set algorithm for quadratic programming. *Math. Program. Comput.* **6**(4), 327–363 (2014)
9. Forsythe, G., Golub, G.H.: On the stationary values of a second-degree polynomial on the unit sphere. *J. Soc. Ind. Appl. Math.* **13**(4), 1050–1068 (1965)
10. Gander, W., Golub, G.H., von Matt, U.: A constrained eigenvalue problem. *Linear Algebra Appl.* **114–115**, 815–839 (1989). Special Issue Dedicated to Alan J. Hoffman
11. Garstka, M., Cannon, M., Goulart, P.J.: COSMO: a conic operator splitting method for convex conic problems. [arXiv:1901.10887](https://arxiv.org/abs/1901.10887) (2019)
12. Gill, P., Murray, W., Saunders, M., Wright, M.: Inertia-controlling methods for general quadratic programming. *SIAM Rev.* **33**(1), 1–36 (1991)
13. Gill, P.E., Murray, W.: Numerically stable methods for quadratic programming. *Math. Program.* **14**(1), 349–372 (1978)

14. Gould, N., Hribar, M., Nocedal, J.: On the solution of equality constrained quadratic programming problems arising in optimization. *SIAM J. Sci. Comput.* **23**(4), 1376–1395 (2001)
15. Gould, N., Lucidi, S., Roma, M., Toint, P.: Solving the trust-region subproblem using the Lanczos method. *SIAM J. Optim.* **9**(2), 504–525 (1999)
16. Gould, N., Orban, D., Toint, P.: Numerical methods for large-scale nonlinear optimization. *Acta Numerica* **14**, 299–361 (2005)
17. Gould, N., Orban, D., Toint, P.: CUTEst: a constrained and unconstrained testing environment with safe threads for mathematical optimization. *Comput. Optim. Appl.* **60**(3), 545–557 (2015)
18. Gould, N., Toint, P.: Numerical methods for large-scale non-convex quadratic programming. In: *Applied Optimization*, pp. 149–179. Springer US (2002)
19. Gurobi Optimization LLC: Gurobi Optimizer Reference Manual (2018). <http://www.gurobi.com>
20. Jeyakumar, V., Li, G.Y.: Trust-Region problems with linear inequality constraints: exact SDP relaxation, global optimality and robust optimization. *Math. Program.* **147**(1), 171–206 (2014)
21. Jolliffe, I.: *Principal Component Analysis*, pp. 1094–1096. Springer, Berlin (2011)
22. Jolliffe, I.T., Trendafilov, N.T., Uddin, M.: A modified principal component technique based on the LASSO. *J. Comput. Graph. Stat.* **12**(3), 531–547 (2003)
23. Journée, M., Nesterov, Y., Richtárik, P., Sepulchre, R.: Generalized power method for sparse principal component analysis. *J. Mach. Learn. Res.* **11**, 517–553 (2010)
24. Lehoucq, R., Sorensen, D., Yang, C.: *ARPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia (1998)
25. Martínez, J.: Local minimizers of quadratic functions on Euclidean balls and spheres. *SIAM J. Optim.* **4**(1), 159–176 (1994)
26. Murty, K.G., Kabadi, S.N.: Some NP-complete problems in quadratic and nonlinear programming. *Math. Program.* **39**(2), 117–129 (1987)
27. Nocedal, J., Wright, S.J.: *Numerical Optimization*, 2nd edn. Springer, New York (2006)
28. Richtárik, P., Takáč, M., Damla Ahipaşaoğlu, S.: Alternating Maximization: Unifying Framework for 8 Sparse PCA Formulations and Efficient Parallel Codes. arXiv e-prints [arXiv:1212.4137](https://arxiv.org/abs/1212.4137) (2012)
29. Saad, Y.: *Numerical Methods for Large Eigenvalue Problems*. Society for Industrial and Applied Mathematics, Philadelphia (2011)
30. Salahi, M., Taati, A., Wolkowicz, H.: Local nonglobal minima for solving large-scale extended Trust-Region Subproblems. *Comput. Optim. Appl.* **66**(2), 223–244 (2017)
31. Silvester, J.R.: Determinants of block matrices. *Math. Gazette* **84**(501), 460–467 (2000)
32. Stellato, B., Banjac, G., Goulart, P., Bemporad, A., Boyd, S.: OSQP: An Operator Splitting Solver for Quadratic Programs. ArXiv e-prints (2017)
33. Wächter, A., Biegler, L.T.: On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Program.* **106**(1), 25–57 (2006)
34. Yuan, X., Zhang, T.: Truncated power method for sparse eigenvalue problems. *J. Mach. Learn. Res.* **14**(1), 899–925 (2013)
35. Zhang, Y., El Ghaoui, L.: Large-scale sparse principal component analysis with application to text data. In: Shawe-Taylor, J., Zemel, R.S., Bartlett, P.L., Pereira, F., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*, vol. 24, pp. 532–539. Curran Associates, Inc, New York (2011)