SHAPE-CHANGING L-SR1 TRUST-REGION METHODS

JOHANNES BRUST, OLEG BURDAKOV, JENNIFER B. ERWAY, ROUMMEL F. MARCIA, AND YA-XIANG YUAN

Abstract. In this article, we propose a method for solving the trust-region subproblem when a limited-memory symmetric rank-one matrix is used in place of the true Hessian matrix. The method takes advantage of two shape-changing norms to decompose the trust-region subproblem into two separate problems, one of which has a closed-form solution and the other one is easy to solve. Sufficient conditions for global solutions to both subproblems are given. The proposed solver makes use of the structure of limited-memory symmetric rank-one matrices to find solutions that satisfy these optimality conditions. Solutions to the trust-region subproblem are computed to high-accuracy even in the so-called "hard case".

1. Introduction

In this article, we describe a method for minimizing a quadratic function defined by a limited-memory symmetric rank-one (L-SR1) matrix subject to a norm constraint; i.e., for a given $\mathbf{x}_k$,

$$\underset{\mathbf{p}\in\mathbb{R}^n}{\text{minimize}} \;\; \mathcal{Q}(\mathbf{p}) \triangleq \mathbf{g}^T\mathbf{p} + \frac{1}{2}\mathbf{p}^T\mathbf{B}\mathbf{p} \quad \text{subject to} \quad \|\mathbf{p}\| \leq \delta, \tag{1}$$

where $\mathbf{g} \triangleq \nabla f(\mathbf{x}_k)$, $\mathbf{B}$ is an L-SR1 approximation to $\nabla^2 f(\mathbf{x}_k)$, δ is a positive constant, and $\|\cdot\|$ is a given norm. At each iteration of a trust-region method, the *trust-region subproblem* (1) must be solved to obtain a step direction. The norm used in (1) not only defines the trust region shape but also determines the difficulty of solving each subproblem. In large-scale optimization, solving (1) represents the bulk of the computational effort in trust-region methods; thus, the choice of norm has significant consequences for the overall trust-region method.

The most widely-used norm chosen to define the trust-region subproblem is the two-norm. One reason for this choice of norm is that the necessary and sufficient conditions for a global solution to the subproblem defined by the two-norm is well-known [11, 18, 22]; many methods exploit these conditions to compute high-accuracy solutions to the trust-region subproblem (see e.g., [7, 8, 9, 13, 1, 18]). The infinity-norm is sometimes used to define the subproblem; however, when $\mathbf{B}$ is indefinite, as can be the case when $\mathbf{B}$ is a L-SR1 matrix, the subproblem is an NP hard [19, 25]. For more discussion on norms other than the infinity-norm we refer the reader to [5].

Date: July 14, 2016.

Key words and phrases. Large-scale unconstrained optimization, nonconvex optimization, trust-region methods, limited-memory quasi-Newton methods, symmetric rank-one update, shape-changing norm.

J. B. Erway is supported in part by National Science Foundation grant CMMI-1334042.

R. F. Marcia is supported in part by National Science Foundation grant CMMI-1333326.

In this article, we consider the trust-region subproblem defined by a *shape-changing* norm originally proposed in [3]. Generally speaking, shape-changing norms are norms that depend on $\mathbf{B}$; thus, in the quasi-Newton setting where the quasi-Newton matrix $\mathbf{B}$ is updated each iteration, the shape of the trust region changes each iteration. One of the earliest references to shape-changing norms is found in [12] where a norm is implicitly defined by the product of a permutation matrix and a unit lower triangular matrix that arise from a symmetric indefinite factorization of $\mathbf{B}$. Perhaps the most widely-used shape-changing norm is the so-called "elliptic norm" given by $\|\mathbf{x}\|_B \triangleq \mathbf{x}^{\mathbf{T}}\mathbf{B}\mathbf{x}$, where $\mathbf{B}$ is a positive-definite matrix (see, e.g., [6]). A well-known use of this norm is found in the Steihaug method [23], and, more generally, truncated preconditioned conjugate-gradients (CG) [6]; these methods reformulate a two-norm trust-region subproblem using an elliptic norm to maintain the property that the iterates from preconditioned CG are increasing in norm. Other examples of shape-changing norms include those defined by vectors in the span of $\mathbf{B}$ (see, e.g., [6]).

1.1. **Overview of the proposed method.** The two shape-changing norms used in this article were originally proposed in [3] (developed in [2]) to decompose the (1) in such a way that global solutions can be computed efficiently. Specifically, the shape-changing norms decouple the trust-region subproblem into two subproblems, one of which has a closed-form solution while the other can be solved very efficiently using techniques borrowed from [2, 1]. This work can be viewed as an extension of [2] in the case when L-SR1 matrices are used to define the trust-region subproblem, allowing high-accuracy subproblem solutions to be computed by exploiting the structure of L-SR1 matrices.

This paper is organized as follows: In Section 2, we review L-SR1 matrices, including the compact representation for these matrices and a method to efficiently compute their eigenvalues and a partial eigenbasis. In Section 3, we demonstrate how the shape-changing norms decouple the original trust-region subproblem into two problems and describe the proposed solver for each subproblem. Finally, we show how to construct a global solution to (1) from the solutions of the two decoupled subproblems. Optimality conditions are presented for each of these decoupled subproblems in Section 4. In Section 5, we demonstrate the accuracy of the proposed solver, and concluding remarks can be found in Section 6.

1.2. **Notation.** In this article, the identity matrix of dimension d is denoted by $\mathbf{I}_d = [\mathbf{e}_1 | \cdots | \mathbf{e}_d]$, and depending on the context the subscript d may be suppressed. Finally, we assume that all L-SR1 updates are computed so that the L-SR1 matrix is well defined.

2. L-SR1 matrices

Suppose $f : \mathbb{R}^n \to \mathbb{R}$ is a smooth objective function and $\{x_i\}$, $i = 0, \ldots k$, is a sequence of iterates, then the symmetric rank-one (SR1) matrix is defined using pairs (s_i, y_i) where

$$\mathbf{s}_i \triangleq \mathbf{x}_{i+1} - \mathbf{x}_i \quad \text{and} \quad \mathbf{y}_i \triangleq \nabla f(\mathbf{x}_{i+1}) - \nabla f(\mathbf{x}_i),$$

and ∇f denotes the gradient of f. Specifically, given an initial matrix $\mathbf{B}_0$, $\mathbf{B}_{k+1}$ is defined recursively as

$$\mathbf{B}_{k+1} \triangleq \mathbf{B}_k + \frac{(\mathbf{y}_k - \mathbf{B}_k\mathbf{s}_k)(\mathbf{y}_k - \mathbf{B}_k\mathbf{s}_k)^T}{(\mathbf{y}_k - \mathbf{B}_k\mathbf{s}_k)^T\mathbf{s}_k}, \tag{2}$$

provided $(\mathbf{y}_k - \mathbf{B}_k\mathbf{s}_k)^T\mathbf{s}_k \neq 0$. In practice, $\mathbf{B}_0$ is often taken to be scalar multiple of the identity matrix; for the duration of this article we assume that $\mathbf{B}_0 = \gamma_k\mathbf{I}$, $\gamma_k \in \mathbb{R}$. *Limited-memory* symmetric rank-one matrices (L-SR1) store and make use of only the m most-recently computed pairs $\{(\mathbf{s}_i, \mathbf{y}_i)\}$, where $m \ll n$ (for example, Byrd et al. [4] suggest $m \in [3, 7]$).

The SR1 update is a member of the Broyden class of updates (see, e.g., [21]). Unlike widely-used updates such as the Broyden-Fletcher-Goldfarb-Shanno (BFGS) and the Davidon-Fletcher-Powell (DFP) updates, this update can yield indefinite matrices; that is, SR1 matrices can incorporate negative curvature information. In fact, SR1 matrices have convergence properties superior to other widely-used positive-definite quasi-Newton matrices such as BFGS [5]. (For more background on the SR1 update formula, see, e.g., [14, 15, 16, 21, 24, 26].)

2.1. **Compact representation.** The compact representation of SR1 matrices can be used to compute the eigenvalues and a partial eigenbasis of these matrices. In this section, we review the compact formulation of SR1 matrices.

To begin, we define the following matrices:

$$\begin{aligned} \mathbf{S}_k &\triangleq \begin{bmatrix} \mathbf{s}_0 & \mathbf{s}_1 & \mathbf{s}_2 & \cdots & \mathbf{s}_k \end{bmatrix} \in \mathbb{R}^{n\times(k+1)}, \\ \mathbf{Y}_k &\triangleq \begin{bmatrix} \mathbf{y}_0 & \mathbf{y}_1 & \mathbf{y}_2 & \cdots & \mathbf{y}_k \end{bmatrix} \in \mathbb{R}^{n\times(k+1)}. \end{aligned}$$

The matrix $\mathbf{S}_k^T\mathbf{Y}_k \in \mathbb{R}^{(k+1)\times(k+1)}$ can be written as the sum of the following three matrices:

$$\mathbf{S}_k^T\mathbf{Y}_k = \mathbf{L}_k + \mathbf{D}_k + \mathbf{R}_k,$$

where $\mathbf{L}_k$ is strictly lower triangular, $\mathbf{D}_k$ is diagonal, and $\mathbf{R}_k$ is strictly upper triangular. Then, $\mathbf{B}_{k+1}$ can be written as

$$\mathbf{B}_{k+1} = \gamma_k\mathbf{I} + \mathbf{\Psi}_k\mathbf{M}_k\mathbf{\Psi}_k^T, \tag{3}$$

where $\mathbf{\Psi}_k \in \mathbb{R}^{n\times(k+1)}$ and $\mathbf{M}_k \in \mathbb{R}^{(k+1)\times(k+1)}$. In particular, $\mathbf{\Psi}_k$ and $\mathbf{M}_k$ are given by

$$\mathbf{\Psi}_k = \mathbf{Y}_k - \gamma_k\mathbf{S}_k \quad \text{and} \quad \mathbf{M}_k = (\mathbf{D}_k + \mathbf{L}_k + \mathbf{L}_k^T - \gamma_k\mathbf{S}_k^T\mathbf{S}_k)^{-1}.$$

The right side of equation (3) is the *compact representation* of $\mathbf{B}_{k+1}$; this representation is due to Byrd et al. [4, Theorem 5.1]. For the duration of this paper, we assume that updates are only accepted when both the next SR1 matrix $\mathbf{B_{k+1}}$ is well-defined and $\mathbf{M}_k$ exists [4, Theorem 5.1]. For notational simplicity, we assume $\mathbf{\Psi}_k$ has full column rank; when $\mathbf{\Psi}_k$ does not have full column rank, we refer to reader to [2] for the modifications needed for computing the eigenvalues reviewed in Section 2.2. Notice that the computation of $\mathbf{M}_k$ is computationally admissible since it is a very small square matrix.

2.2. **Eigenvalues.** In this subsection, we demonstrate how the eigenvalues and a partial eigenbasis can be computed for SR1 matrices. In general, this derivation can be done for any limited-memory quasi-Newton matrix that admits a compact representation; in particular, it can be done for any member of the Broyden convex class [10]. This discussion is based on [2].

Consider the problem of computing the eigenvalues of $\mathbf{B}_{k+1}$, which is assumed to be an L-SR1 matrix, obtained from performing $(k+1)$ rank-one updates to $\mathbf{B}_0 = \gamma\mathbf{I}$. For notational simplicity, we drop subscripts and consider the compact representation of $\mathbf{B}$:

$$\mathbf{B} = \gamma\mathbf{I} + \mathbf{\Psi}\mathbf{M}\mathbf{\Psi}^T, \tag{4}$$

The "thin" QR factorization can be written as $\mathbf{\Psi} = \mathbf{Q}\mathbf{R}$ where $\mathbf{Q} \in \mathbb{R}^{n\times(k+1)}$ and $\mathbf{R} \in \mathbb{R}^{(k+1)\times(k+1)}$ is invertible because, as it was assumed above, $\mathbf{\Psi}$ has full column rank. Then,

$$\mathbf{B} = \gamma\mathbf{I} + \mathbf{Q}\mathbf{R}\mathbf{M}\mathbf{R}^T\mathbf{Q}^T. \tag{5}$$

The matrix $\mathbf{R}\mathbf{M}\mathbf{R}^{\mathbf{T}} \in \mathbb{R}^{(k+1)\times(k+1)}$ is of a relatively small size, and thus, it is computationally inexpensive to compute its spectral decomposition, i.e., $\mathbf{R}\mathbf{M}\mathbf{R}^T = \mathbf{U}\hat{\Lambda}\mathbf{U}^T$, where $\mathbf{U} \in \mathbb{R}^{(k+1)\times(k+1)}$ is orthogonal and $\hat{\Lambda} = \mathrm{diag}(\hat{\lambda}_1, \ldots, \hat{\lambda}_{k+1})$.

Thus,

$$\mathbf{B} = \gamma\mathbf{I} + \mathbf{Q}\mathbf{U}\hat{\Lambda}\mathbf{U}^T\mathbf{Q}^T.$$

Since both $\mathbf{Q}$ and $\mathbf{U}$ have orthonormal columns, $\mathbf{P}_\parallel \triangleq \mathbf{Q}\mathbf{U} \in \mathbb{R}^{n\times(k+1)}$ also has orthonormal columns. Let $P_\perp$ denote the matrix whose columns form an orthonormal basis for $\left(P_\parallel\right)^\perp$. Thus, the spectral decomposition of $\mathbf{B}$ is given by $\mathbf{B} = \mathbf{P}\Lambda_\gamma\mathbf{P}^T$, where

$$\mathbf{P} \triangleq \begin{bmatrix}\mathbf{P}_\parallel & \mathbf{P}_\perp\end{bmatrix} \quad \text{and} \quad \Lambda_\gamma \triangleq \begin{bmatrix}\Lambda & 0 \\ 0 & \gamma\mathbf{I}_{n-(k+1)}\end{bmatrix}, \tag{6}$$

with $\Lambda_\gamma = \mathrm{diag}(\lambda_1, \ldots, \lambda_n)$ and $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_{k+1}) = \hat{\Lambda} + \gamma\mathbf{I} \in \mathbb{R}^{(k+1)\times(k+1)}$.

We emphasize three important properties of the eigendecomposition. First, all eigenvalues of $\mathbf{B}$ are explicitly obtained and represented by Λ_γ. Second, only the first $(k+1)$ eigenvectors of $\mathbf{B}$ can be explicitly computed, if needed; they are represented by $\mathbf{P}_\parallel$. In particular, since $\mathbf{\Psi} = \mathbf{Q}\mathbf{R}$, then

$$\mathbf{P}_\parallel = \mathbf{Q}\mathbf{U} = \mathbf{\Psi}\mathbf{R}^{-1}\mathbf{U}. \tag{7}$$

If $\mathbf{P}_\parallel$ needs only be available to compute matrix-vector products then one can avoid explicitly forming $\mathbf{P}_\parallel$ by storing $\mathbf{\Psi}$, $\mathbf{R}$, and $\mathbf{U}$. Third, the eigenvalues given by the parameter γ can be interpreted as an estimate of the curvature of f in the space spanned by the columns of $\mathbf{P}_\perp$. While there is no reason to assume the function f has negative curvature throughout the entire subspace $\mathbf{P}_\perp$, in this paper, we consider the case $\gamma \leq 0$ for the sake of completeness.

An alternative approach to computing the eigenvalues of $\mathbf{B}$ is presented in [17]. This method replaces the QR factorization of $\mathbf{\Psi}$ with the SVD and an eigendecomposition of a $(k+1)\times(k+1)$ matrix and $t\times t$ matrix, respectively, where $t \leq (k+1)$. For more details, see [17].

For the duration of this article, we assume the first $(k+1)$ eigenvalues in Λ_γ are ordered in increasing values, i.e., $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_{k+1})$ where $\lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_{k+1}$ and that r is the multiplicity of λ_1, i.e., $\lambda_1 = \lambda_2 = \cdots = \lambda_r < \lambda_{r+1}$. For details on updating this partial spectral decomposition when a new quasi-Newton pair is computed, see [10].

3. Proposed method

The proposed method is able to solve the L-SR1 trust-region subproblem to high accuracy, even when $\mathbf{B}$ is indefinite. The method makes use of the eigenvalues of $\mathbf{B}$ and the factors of $\mathbf{P}_\parallel$. To describe the method, we first transform the trust-region subproblem (1) so that the quadratic objective function becomes separable. Then, we describe the shape-changing norms proposed in [3, 2] that decouples the separable problem into two minimization problems, one of which has a closed-form solution while the other can be solved very efficiently. Finally, we show how these solutions can be used to construct a solution to the original trust-region subproblem.

3.1. **Transforming the Trust-Region Subproblem.** Let $\mathbf{B} = \mathbf{P}\Lambda_\gamma\mathbf{P}^T$ be the eigendecomposition of $\mathbf{B}$ described in Section 2.2. Letting $\mathbf{v} = \mathbf{P}^T\mathbf{p}$ and $\mathbf{g}_\mathbf{P} = \mathbf{P}^T\mathbf{g}$, the objective function $\mathcal{Q}(\mathbf{p})$ in (1) can be written as a function of v:

$$\mathcal{Q}(\mathbf{p}) = \mathbf{g}^T\mathbf{p} + \frac{1}{2}\mathbf{p}^T\mathbf{B}\mathbf{p} = \mathbf{g}_\mathbf{P}^T\mathbf{v} + \frac{1}{2}\mathbf{v}^T\Lambda_\gamma\mathbf{v} \triangleq q(\mathbf{v}).$$

With $\mathbf{P} = \begin{bmatrix}\mathbf{P}_\parallel & \mathbf{P}_\perp\end{bmatrix}$, we partition $\mathbf{v}$ and $\mathbf{g}_\mathbf{P}$ as follows:

$$\mathbf{v} = \mathbf{P}^T\mathbf{p} = \begin{bmatrix}\mathbf{P}_\parallel^T\mathbf{p} \\ \mathbf{P}_\perp^T\mathbf{p}\end{bmatrix} = \begin{bmatrix}\mathbf{v}_\parallel \\ \mathbf{v}_\perp\end{bmatrix} \quad \text{and} \quad \mathbf{g}_\mathbf{P} = \begin{bmatrix}\mathbf{P}_\parallel^T\mathbf{g} \\ \mathbf{P}_\perp^T\mathbf{g}\end{bmatrix} = \begin{bmatrix}\mathbf{g}_\parallel \\ \mathbf{g}_\perp\end{bmatrix},$$

where $\mathbf{v}_\parallel, \mathbf{g}_\parallel \in \mathbb{R}^{(k+1)}$ and $\mathbf{v}_\perp, \mathbf{g}_\perp \in \mathbb{R}^{n-(k+1)}$. Then,

$$\begin{aligned}
q(\mathbf{v}) &= \begin{bmatrix}\mathbf{g}_\parallel^T & \mathbf{g}_\perp^T\end{bmatrix}\begin{bmatrix}\mathbf{v}_\parallel \\ \mathbf{v}_\perp\end{bmatrix} + \frac{1}{2}\begin{bmatrix}\mathbf{v}_\parallel^T & \mathbf{v}_\perp^T\end{bmatrix}\begin{bmatrix}\Lambda & \\ & \gamma\mathbf{I}_{n-(k+1)}\end{bmatrix}\begin{bmatrix}\mathbf{v}_\parallel \\ \mathbf{v}_\perp\end{bmatrix} \\
&= \mathbf{g}_\parallel^T\mathbf{v}_\parallel + \mathbf{g}_\perp^T\mathbf{v}_\perp + \frac{1}{2}\left(\mathbf{v}_\parallel^T\Lambda\mathbf{v}_\parallel + \gamma\left\|\mathbf{v}_\perp\right\|^2\right) \\
&= q_\parallel\left(\mathbf{v}_\parallel\right) + q_\perp\left(\mathbf{v}_\perp\right), \qquad (8)
\end{aligned}$$

where

$$q_\parallel\left(\mathbf{v}_\parallel\right) \triangleq \mathbf{g}_\parallel^T\mathbf{v}_\parallel + \frac{1}{2}\mathbf{v}_\parallel^T\Lambda\mathbf{v}_\parallel \quad \text{and} \quad q_\perp\left(\mathbf{v}_\perp\right) \triangleq \mathbf{g}_\perp^T\mathbf{v}_\perp + \frac{\gamma}{2}\left\|\mathbf{v}_\perp\right\|^2.$$

Thus, the trust-region subproblem (1) can be expressed as

$$\underset{\|\mathbf{P}\mathbf{v}\|\le\delta}{\text{minimize}} \; q(\mathbf{v}) = q_\parallel\left(\mathbf{v}_\parallel\right) + q_\perp\left(\mathbf{v}_\perp\right). \qquad (9)$$

Note that the function $q(\mathbf{v})$ is now separable in $\mathbf{v}_\parallel$ and $\mathbf{v}_\perp$. To completely decouple (9) into two minimization problems, we use a shape-changing norm so that the norm constraint $\|\mathbf{P}\mathbf{v}\| \le \delta$ decouples into separate constraints, one involving $\mathbf{v}_\parallel$ and the other involving $\mathbf{v}_\perp$.

3.2. **Shape-Changing Norms.** Consider the following shape-changing norms proposed in [3, 2]:

$$\begin{aligned}
\|\mathbf{p}\|_{\mathbf{P},2} &\triangleq \max\left(\|\mathbf{P}_\parallel^T\mathbf{p}\|_2, \|\mathbf{P}_\perp^T\mathbf{p}\|_2\right) = \max\left(\|\mathbf{v}_\parallel\|_2, \|\mathbf{v}_\perp\|_2\right), \qquad (10) \\
\|\mathbf{p}\|_{\mathbf{P},\infty} &\triangleq \max\left(\|\mathbf{P}_\parallel^T\mathbf{p}\|_\infty, \|\mathbf{P}_\perp^T\mathbf{p}\|_2\right) = \max\left(\|\mathbf{v}_\parallel\|_\infty, \|\mathbf{v}_\perp\|_2\right). \qquad (11)
\end{aligned}$$

We refer to them as the $(\mathbf{P}, 2)$ and the $(\mathbf{P}, \infty)$ norms, respectively. Since $\mathbf{p} = \mathbf{P}\mathbf{v}$, the trust-region constraint in (9) can be expressed in these norms as

$$\begin{aligned}
\|\mathbf{P}\mathbf{v}\|_{\mathbf{P},2} \le \delta \quad &\text{if and only if} \quad \|\mathbf{v}_\parallel\|_2 \le \delta \text{ and } \|\mathbf{v}_\perp\|_2 \le \delta, \\
\|\mathbf{P}\mathbf{v}\|_{\mathbf{P},\infty} \le \delta \quad &\text{if and only if} \quad \|\mathbf{v}_\parallel\|_\infty \le \delta \text{ and } \|\mathbf{v}_\perp\|_2 \le \delta.
\end{aligned}$$

Thus, from (9), the trust-region subproblem is given for the $(\mathbf{P},2)$ norm by

$$\underset{\|\mathbf{Pv}\|_{\mathbf{P},2}\le\delta}{\text{minimize}}\ q(\mathbf{v}) = \underset{\|\mathbf{v}_\parallel\|_2\le\delta}{\text{minimize}}\ q_\parallel\left(\mathbf{v}_\parallel\right) + \underset{\|\mathbf{v}_\perp\|_2\le\delta}{\text{minimize}}\ q_\perp\left(\mathbf{v}_\perp\right), \tag{12}$$

and using the $(\mathbf{P},\infty)$ norm it is given by

$$\underset{\|\mathbf{Pv}\|_{\mathbf{P},\infty}\le\delta}{\text{minimize}}\ q(\mathbf{v}) = \underset{\|\mathbf{v}_\parallel\|_\infty\le\delta}{\text{minimize}}\ q_\parallel\left(\mathbf{v}_\parallel\right) + \underset{\|\mathbf{v}_\perp\|_2\le\delta}{\text{minimize}}\ q_\perp\left(\mathbf{v}_\perp\right). \tag{13}$$

As shown in [2], these norms are equivalent to the two-norm, i.e.,

$$\begin{aligned}\frac{1}{\sqrt{2}}\|\mathbf{p}\|_2 \ &\le\ \ \|\mathbf{p}\|_{\mathbf{P},2} \ \ \le \ \|\mathbf{p}\|_2\\ \frac{1}{\sqrt{k+1}}\|\mathbf{p}\|_2 \ &\le\ \ \|\mathbf{p}\|_{\mathbf{P},\infty} \ \ \le \ \|\mathbf{p}\|_2.\end{aligned}$$

Note that the equivalence factors depend on the number of stored quasi-Newton pairs $(k+1)$ and not on the number of variables (n).

We now show how to solve the decoupled subproblems.

3.3. **Solving for the optimal $\mathbf{v}_\perp^*$.** The subproblem

$$\underset{\|\mathbf{v}_\perp\|_2\le\delta}{\text{minimize}}\quad q_\perp\left(\mathbf{v}_\perp\right) \equiv \mathbf{g}_\perp^T\mathbf{v}_\perp + \frac{\gamma}{2}\|\mathbf{v}_\perp\|_2^2 \tag{14}$$

appears in both (12) and (13); its optimal solution can be computed by formula. For the quadratic subproblem (14) the solution $\mathbf{v}_\perp^*$ must satisfy the following optimality conditions found in [11, 18, 22] associated with (14): For some $\sigma_\perp^* \in \mathbb{R}^+$,

$$\left(\gamma + \sigma_\perp^*\right)\mathbf{v}_\perp^* = -\mathbf{g}_\perp, \tag{15a}$$

$$\sigma_\perp^*\left(\|\mathbf{v}_\perp^*\|_2 - \delta\right) = 0, \tag{15b}$$

$$\|\mathbf{v}_\perp^*\|_2 \le \delta, \tag{15c}$$

$$\gamma + \sigma_\perp^* \ge 0. \tag{15d}$$

Note that the optimality conditions are satisfied by $(\mathbf{v}_\perp^*, \sigma_\perp^*)$ given by

$$\mathbf{v}_\perp^* = \begin{cases} -\frac{1}{\gamma}\mathbf{g}_\perp & \text{if } \gamma > 0 \text{ and } \left\|\mathbf{g}_\perp\right\|_2 \le \delta|\gamma|,\\ \delta\mathbf{u} & \text{if } \gamma \le 0 \text{ and } \|\mathbf{g}_\perp\|_2 = 0,\\ -\frac{\delta}{\|\mathbf{g}_\perp\|_2}\mathbf{g}_\perp & \text{otherwise,}\end{cases} \tag{16}$$

and

$$\sigma_\perp^* = \begin{cases} 0 & \text{if } \gamma \ge 0 \text{ and } \left\|\mathbf{g}_\perp\right\|_2 \le \delta|\gamma|,\\ \frac{\|\mathbf{g}_\perp\|_2}{\delta} - \gamma & \text{otherwise,}\end{cases} \tag{17}$$

where $\mathbf{u} \in \mathbb{R}^{n-(k+1)}$ is any unit vector with respect to the two-norm.

3.4. **Solving for the optimal $\mathbf{v}_\parallel^*$.** In this section, we detail how to solve for the optimal $\mathbf{v}_\parallel^*$ when either the $(\mathbf{P},\infty)$-norm or the $(\mathbf{P},2)$-norm is used to define the trust-region subproblem.

$(\mathbf{P},\infty)$-norm solution. If the shape-changing $(\mathbf{P},\infty)$-norm is used in (9), then the subproblem in $\mathbf{v}_\parallel$ is

$$\underset{\left\|\mathbf{v}_\parallel\right\|_\infty\le\delta}{\text{minimize}}\quad q_\parallel\left(\mathbf{v}_\parallel\right) = \mathbf{g}_\parallel^T\mathbf{v}_\parallel + \frac{1}{2}\mathbf{v}_\parallel^T\Lambda\mathbf{v}_\parallel. \tag{18}$$

The solution to this problem is computed by separately minimizing $(k+1)$ scalar quadratic problems of the form

$$\underset{|[\mathbf{v}_\parallel]_i|\le\delta}{\text{minimize}} \quad q_{\parallel,i}([\mathbf{v}_\parallel]_i) = \left[\mathbf{g}_\parallel\right]_i \left[\mathbf{v}_\parallel\right]_i + \frac{\lambda_i}{2}\left(\left[\mathbf{v}_\parallel\right]_i\right)^2, \qquad 1 \le i \le (k+1). \tag{19}$$

The minimizer depends on the convexity of $q_{\parallel,i}$, i.e., the sign of λ_i. The solution to (19) is given as follows:

$$[\mathbf{v}_\parallel^*]_i = \begin{cases} -\frac{\left[\mathbf{g}_\parallel\right]_i}{\lambda_i} & \text{if } \left|\frac{\left[\mathbf{g}_\parallel\right]_i}{\lambda_i}\right| \le \delta \text{ and } \lambda_i > 0, \\ c & \text{if } \left[\mathbf{g}_\parallel\right]_i = 0,\ \lambda_i = 0, \\ -\text{sgn}(\left[\mathbf{g}_\parallel\right]_i)\delta & \text{if } \left[\mathbf{g}_\parallel\right]_i \neq 0,\ \lambda_i = 0, \\ \pm\delta & \text{if } \left[\mathbf{g}_\parallel\right]_i = 0,\ \lambda_i < 0, \\ -\frac{\delta}{\left|\left[\mathbf{g}_\parallel\right]_i\right|}\left[\mathbf{g}_\parallel\right]_i & \text{otherwise,} \end{cases} \tag{20}$$

where c is any real number in $[-\delta, \delta]$ and "sgn" denotes the signum function (see [2] for details).

$(\mathbf{P}, 2)$-norm solution: If the shape-changing $(\mathbf{P}, 2)$-norm is used in (9), then the subproblem in $\mathbf{v}_\parallel$ is

$$\underset{\|\mathbf{v}_\parallel\|_2\le\delta}{\text{minimize}} \quad q_\parallel\left(\mathbf{v}_\parallel\right) = \mathbf{g}_\parallel^T\mathbf{v}_\parallel + \frac{1}{2}\mathbf{v}_\parallel^T\Lambda\mathbf{v}_\parallel. \tag{21}$$

The solution $\mathbf{v}_\parallel^*$ must satisfy the following optimality conditions [11, 18, 22] associated with (21): For some $\sigma_\parallel^* \in \mathbb{R}^+$,

$$(\Lambda + \sigma_\parallel^*\mathbf{I})\mathbf{v}_\parallel^* = -\mathbf{g}_\parallel, \tag{22a}$$

$$\sigma_\parallel^*\left(\|\mathbf{v}_\parallel^*\|_2 - \delta\right) = 0, \tag{22b}$$

$$\|\mathbf{v}_\parallel^*\|_2 \le \delta, \tag{22c}$$

$$\lambda_i + \sigma_\parallel^* \ge 0 \quad \text{for } 1 \le i \le (k+1). \tag{22d}$$

A solution to the optimality conditions (22a)-(22d) can be computed using the method found in [1]. For completeness, we outline the method here; this method depends on the sign of λ_1. Throughout these cases, we make use of the expression of $\mathbf{v}_\parallel$ as a function of $\sigma_\parallel$. That is, from the first optimality condition (22a), we write

$$\mathbf{v}_\parallel\left(\sigma_\parallel\right) = -\left(\Lambda + \sigma_\parallel\mathbf{I}\right)^{-1}\mathbf{g}_\parallel, \tag{23}$$

with $\sigma_\parallel \neq -\lambda_i$ for $1 \le i \le (k+1)$.

Case 1 ($\lambda_1 > 0$). When $\lambda_1 > 0$, the unconstrained minimizer is computed (setting $\sigma_\parallel^* = 0$):

$$\mathbf{v}_\parallel(0) = -\Lambda^{-1}\mathbf{g}_\parallel. \tag{24}$$

If $\mathbf{v}_\parallel(0)$ is feasible, i.e., $\|\mathbf{v}_\parallel(0)\|_2 \le \delta$ then $\mathbf{v}_\parallel^* = \mathbf{v}_\parallel(0)$ is the global minimizer; otherwise, $\sigma_\parallel^*$ is the solution to the secular equation (28) (discussed below). The minimizer to the problem (21) is then given by

$$\mathbf{v}_\parallel^* = -\left(\Lambda + \sigma_\parallel^*\mathbf{I}\right)^{-1}\mathbf{g}_\parallel. \tag{25}$$

Case 2 ($\lambda_1 = 0$). If $\mathbf{g}_\parallel$ is in the range of Λ, i.e., $[g_\parallel]_i = 0$ for $1 \le i \le r$, then set $\sigma_\parallel = 0$ and let

$$\mathbf{v}_\parallel(0) = -\Lambda^\dagger \mathbf{g}_\parallel,$$

where $\dagger$ denotes the pseudo-inverse. If $\|\mathbf{v}_\parallel(0)\|_2 \le \delta$, then

$$\mathbf{v}_\parallel^* = \mathbf{v}_\parallel(0) = -\Lambda^\dagger \mathbf{g}_\parallel$$

satisfies all optimality conditions (with $\sigma_\parallel^* = 0$). Otherwise, i.e., if either $[\mathbf{g}_\parallel]_i \neq 0$ for some $1 \le i \le r$ or $\|\Lambda^\dagger \mathbf{g}_\parallel\|_2 > \delta$, then $\mathbf{v}_\parallel^*$ is computed using (25), where $\sigma_\parallel^*$ solves the secular equation in (28) (discussed below).

Case 3 ($\lambda_1 < 0$): If $\mathbf{g}_\parallel$ is in the range of $\Lambda - \lambda_1 \mathbf{I}$, i.e., $[g_\parallel]_i = 0$ for $1 \le i \le r$, then we set $\sigma_\parallel = -\lambda_1$ and

$$\mathbf{v}_\parallel(-\lambda_1) = -\left(\Lambda - \lambda_1 \mathbf{I}\right)^\dagger \mathbf{g}_\parallel.$$

If $\|\mathbf{v}_\parallel(-\lambda_1)\|_2 \le \delta$, then the solution is given by

$$\mathbf{v}_\parallel^* = \mathbf{v}_\parallel(-\lambda_1) + \alpha \mathbf{e}_1, \tag{26}$$

where $\alpha = \sqrt{\delta^2 - \left\|\mathbf{v}_\parallel(-\lambda_1)\right\|_2^2}$. (This case is referred to as the "hard case" [6, 18].) Note that $v_\parallel^*$ satisfies the first optimality condition (22a):

$$\left(\Lambda - \lambda_1 \mathbf{I}\right)\mathbf{v}_\parallel^* = \left(\Lambda - \lambda_1 \mathbf{I}\right)\left(\mathbf{v}_\parallel(-\lambda_1) + \alpha \mathbf{e}_1\right) = -\mathbf{g}_\parallel.$$

The second optimality condition (22b) is satisfied by observing that

$$\|\mathbf{v}_\parallel^*\|_2^2 = \|\mathbf{v}_\parallel(-\lambda_1)\|_2^2 + \alpha^2 = \delta^2.$$

Finally, since $\sigma_\parallel^* = -\lambda_1 > 0$ the other optimality conditions are also satisfied.

On the other hand, if $[\mathbf{g}_\parallel]_i \neq 0$ for some $1 \le i \le r$ or $\|(\Lambda - \lambda_1 \mathbf{I})^\dagger \mathbf{g}_\parallel\|_2 > \delta$, then $\mathbf{v}_\parallel^*$ is computed using (25), where $\sigma_\parallel^*$ solves the secular equation (28).

The secular equation. We now summarize how to find a solution of the so-called *secular equation*. Note that from (23),

$$\|\mathbf{v}_\parallel(\sigma_\parallel)\|_2^2 = \sum_{i=1}^{k+1} \frac{(\mathbf{g}_\parallel)_i^2}{(\lambda_i + \sigma_\parallel)^2}.$$

If we combine the terms above that correspond to the same eigenvalues and remove the terms with zero numerators, then for $\sigma_\parallel \neq -\lambda_i$, we have

$$\|\mathbf{v}_\parallel(\sigma_\parallel)\|_2^2 = \sum_{i=1}^{\ell} \frac{\bar{a}_i^2}{(\bar{\lambda}_i + \sigma_\parallel)^2},$$

where $\bar{a}_i \neq 0$ for $i = 1, \ldots, \ell$ and $\bar{\lambda}_i$ are *distinct* eigenvalues of $\mathbf{B}$ with $\bar{\lambda}_1 < \bar{\lambda}_2 < \cdots < \bar{\lambda}_\ell$. Next, we define the function

$$\phi_\parallel(\sigma_\parallel) = \begin{cases} \dfrac{1}{\sqrt{\displaystyle\sum_{i=1}^{\ell} \frac{\bar{a}_i^2}{(\bar{\lambda}_i + \sigma_\parallel)^2}}} - \dfrac{1}{\delta} & \text{if } \sigma_\parallel \neq -\bar{\lambda}_i \text{ where } 1 \le i \le \ell \\ -\dfrac{1}{\delta} & \text{otherwise.} \end{cases} \tag{27}$$

From the optimality conditions (22b) and (22d), if $\sigma_\|^* \neq 0$, then $\sigma_\|^*$ solves the *secular equation*

$$\phi_\| \left(\sigma_\|\right) = 0, \tag{28}$$

with $\sigma_\| \geq \max\{0, -\lambda_1\}$. Note that $\phi_\|$ is monotonically increasing and concave down on the interval $[-\lambda_1, \infty)$; thus, Newton's method can be used to efficiently compute $\sigma_\|^*$ in (28).

More details on the solution method for subproblem (21) are given in [1].

3.5. **Computing $\mathbf{p}^*$.** Given $\mathbf{v}^* = [\mathbf{v}_\|^* \ \mathbf{v}_\perp^*]^T$, the solution to the trust-region subproblem (1) using either the $(\mathbf{P}, 2)$ or the $(\mathbf{P}, \infty)$ norms is

$$\mathbf{p}^* = \mathbf{P}\mathbf{v}^* = \mathbf{P}_\|\mathbf{v}_\|^* + \mathbf{P}_\perp\mathbf{v}_\perp^*. \tag{29}$$

(Recall that using either of the two norms generates the same $v_\perp^*$ but different $v_\|^*$.) It remains to show how to form $\mathbf{p}^*$ in (29). Matrix-vector products involving $\mathbf{P}_\|$ are possible using (7), and thus, $\mathbf{P}_\|\mathbf{v}_\|^*$ can be computed; however, an implicit formula to compute products $\mathbf{P}_\perp$ is not available. To compute the second term, $\mathbf{P}_\perp\mathbf{v}_\perp^*$, we observe that $\mathbf{v}_\perp^*$, as given in (16), is a multiple of either $\mathbf{g}_\perp = \mathbf{P}_\perp^T\mathbf{g}$ or a vector $\mathbf{u}$ with unit length. In particular, define $\mathbf{u} = \frac{\mathbf{P}_\perp^T\mathbf{e}_i}{\|\mathbf{P}_\perp^T\mathbf{e}_i\|_2}$, where $i \in \{1, 2, \ldots, k+2\}$ is the first index such that $\left\|\mathbf{P}_\perp^T\mathbf{e}_i\right\|_2 \neq 0$. (Such an $\mathbf{e}_i$ exists since $\text{rank}(\mathbf{P}_\perp) = n-(k+1)$.) Thus, we obtain

$$\mathbf{p}^* = \mathbf{P}_\|\mathbf{v}_\|^* + \left(\mathbf{I} - \mathbf{P}_\|\mathbf{P}_\|^T\right)\mathbf{w}^*, \tag{30}$$

where

$$\mathbf{w}^* = \begin{cases} -\frac{1}{\gamma}\mathbf{g} & \text{if } \gamma > 0 \text{ and } \|\mathbf{g}_\perp\|_2 \leq \delta|\gamma|, \\ \frac{\delta}{\left\|\mathbf{P}_\perp^T\mathbf{e}_i\right\|_2}\mathbf{e}_i & \text{if } \gamma \leq 0 \text{ and } \|\mathbf{g}_\perp\|_2 = 0, \\ -\frac{\delta}{\|\mathbf{g}_\perp\|_2}\mathbf{g} & \text{otherwise.} \end{cases} \tag{31}$$

The quantities $\|\mathbf{g}_\perp\|_2$ and $\left\|\mathbf{P}_\perp^T\mathbf{e}_i\right\|_2$ are computed using the orthogonality of P, which implies

$$\left\|\mathbf{g}_\|\right\|_2^2 + \|\mathbf{g}_\perp\|_2^2 = \|\mathbf{g}\|_2^2, \quad \text{and} \quad \|\mathbf{P}_\|^T\mathbf{e}_i\|_2^2 + \|\mathbf{P}_\perp^T\mathbf{e}_i\|_2^2 = 1. \tag{32}$$

Then $\|\mathbf{g}_\perp\|_2 = \sqrt{\|\mathbf{g}\|_2^2 - \|\mathbf{g}_\|\|_2^2}$ and $\|\mathbf{P}_\perp^T\mathbf{e}_i\|_2 = \sqrt{1 - \|\mathbf{P}_\|^T\mathbf{e}_i\|_2^2}$. Note that $\mathbf{v}_\perp^*$ is never explicitly computed.

3.6. **Computational Complexity.** We estimate the cost of one iteration using the proposed method to solve the trust-region subproblem defined by shape-changing norms (10) and (11). We make the practical assumption that $\gamma > 0$. Assuming we have already obtained the Cholesky factorization of $\mathbf{\Psi}^T\mathbf{\Psi}$ associated with the previously-stored limited-memory pairs, it is possible to update the Cholesky factorization of the new $\mathbf{\Psi}^T\mathbf{\Psi}$ at a cost of $O(k^2)$. (Note that if γ is constant from one iteration to the next, then $\mathbf{\Psi}$ is updated with $O(2n)$ operations.) To form $\mathbf{\Psi}^T\mathbf{\Psi}$, we do not store $\mathbf{\Psi}$. Instead, we store and update the small $(k+1)\times(k+1)$ matrices $\mathbf{Y^TY}$, $\mathbf{S^TY}$, and $\mathbf{S^TS}$. Similarly, we compute matrix-vector products with $\mathbf{\Psi} = \mathbf{Y} - \gamma\mathbf{S}$ by computing matrix-vector products with $\mathbf{Y}$ and $\mathbf{S}$.

The eigendecomposition $\mathbf{RMR}^T = \mathbf{U}\hat{\mathbf{\Lambda}}\mathbf{U}^T$ costs $O(k^3) = \left(\frac{k^2}{n}\right)O(kn)$, where $k \ll n$. To compute $\mathbf{p}^*$ in (30), one needs to compute $\mathbf{v}^*$ from Section 3.4 and $\mathbf{w}^*$ from (31). The dominant cost for computing $\mathbf{v}^*$ and $\mathbf{w}^*$ is forming $\mathbf{\Psi}^T\mathbf{g}$, which requires $4kn$ operations. (In practice, this quantity is computed while solving the previous trust-region subproblem and can be stored to avoid recomputing when solving the current subproblem–see [2] for details.) Note that given $\mathbf{P}_{\parallel}^T\mathbf{g}$, the computation of $\mathbf{p}^*$ in (30) costs $O(4kn)$. Finally, the dominant cost to update $\mathbf{\Psi}^T\mathbf{\Psi}$ is $4kn$. Thus, dominant term in the total number of floating point operations is $4kn$. This is the same cost as for L-BFGS [20].

4. Numerical experiments

In this section, we report on numerical experiments with the proposed Shape-Changing SR1 (SC-SR1) algorithm implemented in MATLAB to solve limited-memory SR1 trust-region subproblems. The SC-SR1 algorithm was tested on randomly-generated problems of size $n = 10^3$ to $n = 10^7$, organized as five experiments when there is no closed-form solution to the shape-changing trust-region subproblem and one experiment designed to test the SC-SR1 method in the so-called "hard case". These six cases only occur using the $(\mathbf{P}, 2)$-norm trust region. (In the case of the $(\mathbf{P}, \infty)$ norm, $\mathbf{v}_{\parallel}^*$ has the closed-form solution given by (20).) The six experiments are outlined as follows:

(E1) $\mathbf{B}$ is positive definite with $\|\mathbf{v}_{\parallel}(0)\|_2 \geq \delta$.
(E2) $\mathbf{B}$ is positive semidefinite and singular with $[\mathbf{g}_{\parallel}]_i \neq 0$ for some $1 \leq i \leq r$.
(E3) $\mathbf{B}$ is positive semidefinite and singular with $[\mathbf{g}_{\parallel}]_i = 0$ for $1 \leq i \leq r$ and $\|\Lambda^{\dagger}\mathbf{g}_{\parallel}\|_2 > \delta$.
(E4) $\mathbf{B}$ is indefinite and $[\mathbf{g}_{\parallel}]_i = 0$ for $1 \leq i \leq r$ with $\|(\Lambda - \lambda_1 \mathbf{I})^{\dagger}\mathbf{g}_{\parallel}\|_2 > \delta$.
(E5) $\mathbf{B}$ is indefinite and $[\mathbf{g}_{\parallel}]_i \neq 0$ for some $1 \leq i \leq r$.
(E6) $\mathbf{B}$ is indefinite and $[\mathbf{g}_{\parallel}]_i = 0$ for $1 \leq i \leq r$ with $\|\mathbf{v}_{\parallel}(-\lambda_1)\|_2 \leq \delta$ (the "hard case").

For these experiments, $\mathbf{S}$, $\mathbf{Y}$, and $\mathbf{g}$ were randomly generated and then altered to satisfy the requirements described above by each experiment. All randomly-generated vectors and matrices were formed using the MATLAB `randn` command, which draws from the standard normal distribution. The initial SR1 matrix was set to $\mathbf{B}_0 = \gamma \mathbf{I}$, where $\gamma = |10 * \texttt{randn(1)}|$. Finally, the number of limited-memory updates $(k+1)$ was set to 5, and r was set to 2. In the five cases when there is no closed-form solution, SC-SR1 uses Newton's method to find a root of $\phi_{\parallel}$. We use the same procedure as in [1, Algorithm 2] to initialize Newton's method since it guarantees monotonic and quadratic convergence to σ^*. The Newton iteration was terminated when the ith iterate satisfied $\|\phi_{\parallel}(\sigma^i)\| \leq \texttt{eps} \cdot \|\phi_{\parallel}(\sigma^0)\| + \sqrt{\texttt{eps}}$, where σ^0 denotes the initial iterate for Newton's method and `eps` is machine precision. This stopping criteria is both a relative and absolute criteria, and it is the only stopping criteria used by SC-SR1.

In order to report on the accuracy of the subproblem solves, we make use of the following theorem, which is based on the optimality conditions for a global solution to (1) defined by the two-norm [11, 18]. This theorem characterizes global solutions for the trust-region subproblem defined by the $(\mathbf{P}, 2)$-norm.

Theorem 1. *A vector $\mathbf{p}^* \in \mathbb{R}^n$ such that $\left\|\mathbf{P}_{\parallel}^T\mathbf{p}^*\right\|_2 \leq \delta$ and $\left\|\mathbf{P}_{\perp}^T\mathbf{p}^*\right\|_2 \leq \delta$, is a global solution of (1) defined by the $(\mathbf{P}, 2)$-norm if and only if there exists unique $\sigma_{\parallel}^* \geq 0$ and $\sigma_{\perp}^* \geq 0$ such that*

$$\left(\mathbf{B} + \mathbf{C}_{\parallel}\right)\mathbf{p}^* + \mathbf{g} = 0, \quad \sigma_{\parallel}^*\left(\left\|\mathbf{P}_{\parallel}^T\mathbf{p}^*\right\|_2 - \delta\right) = 0, \quad \sigma_{\perp}^*\left(\left\|\mathbf{P}_{\perp}^T\mathbf{p}^*\right\|_2 - \delta\right) = 0,$$

where $\mathbf{C}_\parallel \triangleq \sigma_\perp^* \mathbf{I} + \left(\sigma_\parallel^* - \sigma_\perp^*\right) \mathbf{P}_\parallel \mathbf{P}_\parallel^T$*, the matrix* $\mathbf{B} + \mathbf{C}_\parallel$ *is positive semi-definite, and* $\mathbf{P} = [\mathbf{P}_\parallel \ \mathbf{P}_\perp]$ *and* $\Lambda = diag(\lambda_1, \ldots, \lambda_{k+1}) = \hat{\Lambda} + \gamma \mathbf{I}$ *are as in (6).*

Thus, for each experiment, we report the following: (i) the norm of the residual of the first optimality condition, `opt 1` $\triangleq \|(\mathbf{B} + \mathbf{C}_\parallel)\mathbf{p}^* + \mathbf{g}\|_2$, (ii) the combined complementarity condition, `opt 2` $\triangleq |\sigma_\parallel^*(\|\mathbf{P}_\parallel^T \mathbf{p}^*\|_2 - \delta)| + |\sigma_\perp^*(\|\mathbf{P}_\perp^T \mathbf{p}^*\|_2 - \delta)|$, (iii) $\sigma_\parallel^* + \lambda_1$, (iv) $\sigma_\perp^* + \gamma$, (v) $\sigma_\parallel^*$, (vi) $\sigma_\perp^*$, (vii) the number of Newton iterations ("itns"), and (viii) time. The quantities (iii) and (iv) are reported since the optimality condition that $\mathbf{B} + \mathbf{C}_\parallel$ is a positive semidefinite matrix is equivalent to $\gamma + \sigma_\perp^* \geq 0$ and $\lambda_i + \sigma_\parallel^* \geq 0$, for $1 \leq i \leq (k+1)$.

TABLE I. Experiment 1: $\mathbf{B}$ is positive definite with $\|\mathbf{v}_\parallel(0)\|_2 \geq \delta$.

n	opt 1	opt 2	$\sigma_\parallel^* + \lambda_1$	$\sigma_\perp^* + \gamma$	$\sigma_\parallel^*$	$\sigma_\perp^*$	itns	time
1.0e+03	1.80e-14	2.76e-14	1.69e+01	1.70e+02	4.23e+00	1.64e+02	2	6.74e-04
1.0e+04	1.26e-13	4.98e-14	4.04e+00	2.24e+02	1.03e+00	2.23e+02	2	1.27e-03
1.0e+05	1.39e-12	1.04e-12	3.77e+01	7.13e+03	9.43e+00	7.11e+03	2	1.29e-02
1.0e+06	2.09e-11	5.83e-12	3.83e+00	2.39e+03	9.60e-01	2.39e+03	2	1.39e-01
1.0e+07	6.13e-11	3.42e-11	4.77e+01	8.12e+04	1.19e+01	8.12e+04	1	1.48e+00

TABLE II. Experiment 2: $\mathbf{B}$ is positive semidefinite and singular and $[\mathbf{g}_\parallel]_i \neq 0$ for some $1 \leq i \leq r$.

n	opt 1	opt 2	$\sigma_\parallel^* + \lambda_1$	$\sigma_\perp^* + \gamma$	$\sigma_\parallel^*$	$\sigma_\perp^*$	itns	time
1.0e+03	1.81e-14	1.55e-13	6.27e+00	1.03e+02	6.27e+00	9.61e+01	3	7.11e-04
1.0e+04	2.12e-13	2.25e-13	6.07e+01	3.23e+03	6.07e+01	3.19e+03	5	1.40e-03
1.0e+05	6.50e-13	4.62e-13	2.29e+00	3.25e+02	2.29e+00	3.18e+02	3	1.22e-02
1.0e+06	1.10e-11	1.11e-11	2.80e+00	3.28e+03	2.80e+00	3.27e+03	3	1.50e-01
1.0e+07	1.16e-10	7.28e-11	4.39e+00	1.12e+04	4.39e+00	1.12e+04	3	1.49e+00

TABLE III. Experiment 3: $\mathbf{B}$ is positive semidefinite and singular with $[\mathbf{g}_\parallel]_i = 0$ for $1 \leq i \leq r$ and $\|\Lambda^\dagger \mathbf{g}_\parallel\|_2 > \delta$.

n	opt 1	opt 2	$\sigma_\parallel^* + \lambda_1$	$\sigma_\perp^* + \gamma$	$\sigma_\parallel^*$	$\sigma_\perp^*$	itns	time
1.0e+03	1.62e-14	6.12e-17	4.41e+00	4.56e+02	4.41e+00	4.49e+02	2	8.49e-04
1.0e+04	1.48e-13	1.07e-13	6.74e+00	2.20e+03	6.74e+00	2.19e+03	2	1.24e-03
1.0e+05	1.02e-12	8.36e-13	7.93e+00	1.15e+04	7.93e+00	1.15e+04	2	1.34e-02
1.0e+06	9.06e-12	2.26e-12	3.06e+00	7.42e+03	3.06e+00	7.41e+03	2	1.45e-01
1.0e+07	1.50e-10	1.09e-10	1.95e+00	9.48e+03	1.95e+00	9.48e+03	2	1.49e+00

Tables I-VI show the results of the experiments. In all tables, the residual of the two optimality conditions `opt 1` and `opt 2` are on the order of $1e^{-10}$ or smaller. Columns 4 and 5 in all the tables show that $\sigma_\parallel^* + \lambda_1$ and $\sigma_\perp^* + \gamma$ are nonnegative with $\sigma_\parallel \geq 0$ and $\sigma_\perp \geq 0$ (Columns 6 and 7, respectively). Thus, the solutions obtained by SC-SR1 for these experiments satisfy the optimality conditions to high accuracy.

Also reported in each table are the number of Newton iterations. In the first five experiments no more than five Newton iterations were required to obtain $\sigma_\parallel$ to high

TABLE IV. Experiment 4: $\mathbf{B}$ is indefinite and $[\mathbf{g}_\parallel]_i = 0$ for $1 \le i \le r$ with $\|(\Lambda - \lambda_1 \mathbf{I})^\dagger \mathbf{g}_\parallel\|_2 > \delta$.

n	opt 1	opt 2	$\sigma_\parallel^* + \lambda_1$	$\sigma_\perp^* + \gamma$	$\sigma_\parallel^*$	$\sigma_\perp^*$	itns	time
1.0e+03	2.92e-14	1.75e-14	3.51e+00	3.15e+02	3.64e+00	3.10e+02	2	8.01e-04
1.0e+04	9.76e-14	1.17e-13	3.91e+00	1.41e+03	4.40e+00	1.40e+03	2	1.24e-03
1.0e+05	1.37e-12	9.57e-13	1.18e+00	7.48e+02	1.88e+00	7.47e+02	2	1.43e-02
1.0e+06	9.19e-12	8.00e-12	7.16e+00	1.48e+04	7.59e+00	1.48e+04	2	1.40e-01
1.0e+07	1.26e-10	2.60e-11	3.71e+00	1.23e+05	4.71e+00	1.23e+05	2	1.48e+00

accuracy (Column 8). In the hard case, no Newton iterations are required since $\sigma_\parallel^* = -\lambda_1$. This is reflected in Table VI, where Column 4 shows that $\sigma_\parallel^* = -\lambda_1$ and Column 8 reports no Newton iterations.)

The final column reports the time required by SC-SR1 to solve each subproblem. Consistent with the best limited-memory methods, the time required to solve each subproblem appears to grow linearly with n.

Additional experiments were run with $\mathbf{g}_\parallel \to 0$. In particular, the experiments were rerun with g scaled by factors of 10^{-2}, 10^{-4}, 10^{-6}, 10^{-8}, and 10^{-10}. All experiments resulted in tables similar to those in Tables I-VI: the optimality conditions were satisfied to high accuracy, no more than three Newton iterations were required in any experiment to find $\sigma_\parallel^*$, and the CPU times are similar to those found in the tables.

TABLE V. Experiment 5: $\mathbf{B}$ is indefinite and $[\mathbf{g}_\parallel]_i \neq 0$ for some $1 \le i \le r$.

n	opt 1	opt 2	$\sigma_\parallel^* + \lambda_1$	$\sigma_\perp^* + \gamma$	$\sigma_\parallel^*$	$\sigma_\perp^*$	itns	time
1.0e+03	2.28e-14	6.05e-15	8.09e-01	5.70e+01	1.65e+00	5.12e+01	3	8.01e-04
1.0e+04	1.06e-13	3.18e-14	1.88e+00	1.74e+02	2.22e+00	1.68e+02	3	1.64e-03
1.0e+05	4.17e-13	5.96e-13	2.02e+00	4.16e+02	2.06e+00	4.12e+02	3	1.25e-02
1.0e+06	1.51e-11	6.98e-12	1.19e+00	1.38e+03	2.14e+00	1.36e+03	3	1.41e-01
1.0e+07	1.52e-10	4.36e-12	1.90e+00	4.90e+03	2.57e+00	4.90e+03	5	1.48e+00

TABLE VI. Experiment 6: $\mathbf{B}$ is indefinite and $[\mathbf{g}_\parallel]_i = 0$ for $1 \le i \le r$ with $\|\mathbf{v}_\parallel(-\lambda_1)\|_2 \le \delta$ (the "hard case").

n	opt 1	opt 2	$\sigma_\parallel^* + \lambda_1$	$\sigma_\perp^* + \gamma$	$\sigma_\parallel^*$	$\sigma_\perp^*$	itns	time
1.0e+03	2.60e-14	4.60e-15	0.00e+00	3.36e+02	9.78e-01	3.30e+02	0	7.63e-04
1.0e+04	1.57e-13	9.69e-14	0.00e+00	4.01e+03	6.84e-01	3.99e+03	0	1.44e-03
1.0e+05	1.39e-12	7.45e-13	0.00e+00	7.46e+02	1.06e-01	7.45e+02	0	1.54e-02
1.0e+06	1.14e-11	9.04e-13	0.00e+00	1.63e+03	5.81e-01	1.62e+03	0	1.57e-01
1.0e+07	9.00e-11	3.73e-11	0.00e+00	4.80e+04	3.40e-01	4.80e+04	0	1.68e+00

5. CONCLUDING REMARKS

In this paper, we consider minimizing function f approximating a Hessian using limited-memory SR1 matrix. Indefinite Hessian approxiamtions cannot be used in a linesearch context.

We presented a high-accuracy trust-region subproblem solver for when the Hessian is approximated by L-SR1 matrices. The method makes use of special shape-changing norms that decouple the original subproblem into two separate subproblems, one of which has a closed-form solution. Numerical experiments verify that solutions are computed to high accuracy in cases when there are no closed-form solutions and also in the so-called "hard case".

References

[1] J. J. Brust, J. B. Erway, and R. F. Marcia. On solving L-SR1 trust-region subproblems. Technical Report 2015-6, Wake Forest University, 2015.

[2] O. Burdakov, L. Gong, S. Zikrin, and Y.-X. Yuan. On efficiently combining limited-memory and trust-region techniques. *Mathematical Programming Computation*, pages 1–34, 2016.

[3] O. Burdakov and Y.-X. Yuan. On limited-memory methods with shape changing trust region. In *Proceedings of the First International Conference on Optimization Methods and Software*, page p. 21, 2002.

[4] R. H. Byrd, J. Nocedal, and R. B. Schnabel. Representations of quasi-Newton matrices and their use in limited-memory methods. *Math. Program.*, 63:129–156, 1994.

[5] A. R. Conn, N. I. M. Gould, and P. L. Toint. Convergence of quasi-Newton matrices generated by the symmetric rank one update. *Math. Programming*, 50(2, (Ser. A)):177–195, 1991.

[6] A. R. Conn, N. I. M. Gould, and P. L. Toint. *Trust-Region Methods*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2000.

[7] J. B. Erway and P. E. Gill. A subspace minimization method for the trust-region step. *SIAM Journal on Optimization*, 20(3):1439–1461, 2010.

[8] J. B. Erway, P. E. Gill, and J. D. Griffin. Iterative methods for finding a trust-region step. *SIAM Journal on Optimization*, 20(2):1110–1131, 2009.

[9] J. B. Erway and R. F. Marcia. Algorithm 943: MSS: MATLAB software for L-BFGS trust-region subproblems for large-scale optimization. *ACM Transactions on Mathematical Software*, 40(4):28:1–28:12, June 2014.

[10] J. B. Erway and R. F. Marcia. On efficiently computing the eigenvalues of limited-memory quasi-newton matrices. *SIAM Journal on Matrix Analysis and Applications*, 36(3):1338–1359, 2015.

[11] D. M. Gay. Computing optimal locally constrained steps. *SIAM J. Sci. Statist. Comput.*, 2(2):186–197, 1981.

[12] D. Goldfarb. The use of negative curvature in minimization algorithms. Technical Report 80-412, Cornell University, 1980.

[13] N. I. M. Gould, D. P. Robinson, and H. S. Thorne. On solving trust-region and other regularised subproblems in optimization. *Mathematical Programming Computation*, 2(1):21–57, 2010.

[14] I. Griva, S. G. Nash, and A. Sofer. *Linear and nonlinear programming*. Society for Industrial and Applied Mathematics, Philadelphia, 2009.

[15] C. T. Kelley and E. W. Sachs. Local convergence of the symmetric rank-one iteration. *Computational Optimization and Applications*, 9(1):43–63, 1998.

[16] H. F. Khalfan, R. H. Byrd, and R. B. Schnabel. A theoretical and experimental study of the symmetric rank-one update. *SIAM Journal on Optimization*, 3(1):1–24, 1993.

[17] X. Lu. *A study of the limited memory SR1 method in practice*. PhD thesis, Department of Computer Science, University of Colorado at Boulder, 1996.

[18] J. J. Moré and D. C. Sorensen. Computing a trust region step. *SIAM J. Sci. and Statist. Comput.*, 4:553–572, 1983.

[19] K. G. Murty and S. N. Kabadi. Some np-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39(2):117–129, 1987.

[20] J. Nocedal. Updating quasi-Newton matrices with limited storage. *Math. Comput.*, 35:773–782, 1980.

[21] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, New York, 2nd edition, 2006.

[22] D. C. Sorensen. Newton's method with a model trust region modification. *SIAM J. Numer. Anal.*, 19(2):409–426, 1982.

[23] T. Steihaug. The conjugate gradient method and trust regions in large scale optimization. *SIAM J. Numer. Anal.*, 20:626–637, 1983.

[24] W. Sun and Y.-x. Yuan. *Optimization theory and methods*, volume 1 of *Springer Optimization and Its Applications*. Springer, New York, 2006. Nonlinear programming.

[25] B. Vavasis. *Nonlinear Optimization: Complexity Issues*. International Series of Monographs on Computer Science. Oxford University Press, Oxford, England, 1992.

[26] H. Wolkowicz. Measures for symmetric rank-one updates. *Mathematics of Operations Research*, 19(4):815–830, 1994.

E-mail address: jbrust@ucmerced.edu

Applied Mathematics, University of California, Merced, Merced, CA 95343

E-mail address: oleg.burdakov@liu.se

Department of Mathematics, Linköping University, SE-581 83 Linköping, Sweden

E-mail address: erwayjb@wfu.edu

Department of Mathematics, Wake Forest University, Winston-Salem, NC 27109

E-mail address: rmarcia@ucmerced.edu

Applied Mathematics, University of California, Merced, Merced, CA 95343

E-mail address: yyx@lsec.cc.ac.cn

State Key Laboratory of Scientific and Engineering, Institute of Computational Mathematics and Scientific/Engineering Computing, AMSS, Chinese Academy of Sciences, Beijing, China.