The ω-Condition Number: Applications to Preconditioning and Low Rank Generalized Jacobian Updating [*][†]

Woosuk L. Jung[‡]   David Torregrosa-Belén[§]   Henry Wolkowicz[‡]

Revision January 31, 2025

Key words and phrases: $\kappa, \omega, \omega^{-2}$-condition numbers, preconditioning, generalized Jacobian, iterative methods, clustering of eigenvalues

AMS subject classifications: 15A12, 65F35, 65F08, 65G50, 49J52, 49K10, 90C32

Contents

---

[*]Emails resp.: w2jung@uwaterloo.ca, david.torregrosa@ua.es, hwolkowicz@uwaterloo.ca
[†]This report is available at URL: www.math.uwaterloo.ca/~hwolkowi/henry/reports/ABSTRACTS.html
[‡]Department of Combinatorics and Optimization, University of Waterloo, 200 University Avenue West, Waterloo, ON N2L 3G1, Canada
[§]Centro de Modelamiento Matemático, Centro de Modelamiento Matemático (CNRS IRL2807), Universidad de Chile, Beauchef 851, Santiago, Chile

List of Tables

List of Figures

Abstract

Preconditioning is essential in iterative methods for solving linear systems. It is also the implicit objective in updating approximations of Jacobians in optimization methods, e.g., in quasi-Newton methods. We study a nonclassic matrix condition number, the ω-*condition number*, ω for short. ω is the ratio of: the arithmetic and geometric means of the singular values, rather than the largest and smallest for the classical κ-condition number. The simple functions in ω allow one to exploit first order optimality conditions. We use this fact to derive *explicit* formulae for (i) ω-optimal low rank updating of generalized Jacobians arising in the context of nonsmooth Newton methods; and (ii) ω-optimal preconditioners of special structure for iterative methods for linear systems. In the latter context, we analyze the benefits of ω for (a) improving the clustering of eigenvalues; (b) reducing the number of iterations; and (c) estimating the actual condition of a linear system. Moreover we show strong theoretical connections between the ω-optimal preconditioners and incomplete Cholesky factorizations, and highlight the misleading effects arising from the inverse invariance of κ. Our results confirm the efficacy of using the ω-condition number compared to the κ-condition number.

1 Introduction

Preconditioning is essential in iterative and direct solutions of linear systems e.g., [4]. It is also the implicit objective in low rank updating of approximate Jacobians in optimization, e.g., in quasi-Newton methods [11]. In this paper we study the ω-condition number (abbreviated as ω), a nonclassic matrix condition number that, for a positive definite matrix, is the ratio of the arithmetic and geometric means of the eigenvalues, rather than the largest and smallest eigenvalues of the classical κ-condition number (abbreviated as κ). We emphasize that ω provides a more average indication rather than a worst case measure of conditioning. Moreover, unlike κ, ω is *not* invariant under inversion. This important property allows one to recall and exploit that it is the conditioning of the inverse that is important.

In particular, our original motivation is to find well conditioned, ω-optimal, low rank updates of the positive definite generalized Jacobian that arises in nonsmooth Newton methods e.g., [3]. We use optimality conditions to find *explicit formulae* for these low rank updates. As well our work includes explicit formulae for ω-optimal diagonal and sparse upper triangular preconditioners. We see that these latter relate to a sparse incomplete Cholesky

3

factorization, i.e., we modify the iterates of the Cholesky factorization by including the entire diagonal at each iteration. We then illustrate both the efficiency and effectiveness of using ω compared to κ when addressing positive definite linear systems. In particular, our empirics show that ω is effective in promoting the important property of clustering of eigenvalues and the related reduction in iteration counts in preconditioned conjugate gradient methods. linelabelline:clustering

In addition, we show that ω can be evaluated exactly following a Cholesky or LU factorization; and that it is a better indication of the conditioning of a problem when compared to κ. [1]

1.1 Background and Preliminaries

In numerical analysis, a condition number of a matrix A is the main tool in the study of error propagation in the problem of solving the linear equation $Ax = b$. The linear system $Ax = b$ is said to be well-conditioned when A has a low condition number. In particular, in the literature $\kappa(A)$ is used as a (worst case) measure of the conditioning of a linear system $Ax = b$, i.e., how much a solution x, the output, will change with respect to changes in the right-hand side b, the input:

$$\operatorname{cond}(A) := \frac{\|\Delta x\|/\|x\|}{\|\Delta b\|/\|b\|}, \tag{1.1}$$

e.g., [35, Sect. 1.3]. In general, iterative algorithms that solve $Ax = b$ require a large number of iterations to achieve a solution with sufficient accuracy if the problem is not well-conditioned, i.e., is ill-conditioned. For simplicity, in this paper, we restrict ourselves to A positive definite and so $\kappa(A) = \lambda_1(A)/\lambda_n(A) \left(= \kappa(A^{-1}) \right)$.

In order to improve the conditioning of a problem, preconditioners are employed for obtaining equivalent systems with better condition number. For example, in [7] a preconditioner that minimizes κ is obtained in the Broyden family of rank-two updates. Also, for applications to inexact Newton methods see [1, 2], where it is emphasized that the goal is to improve the *clustering of eigenvalues* around 1. The ω-condition number in particular uses *all* the eigenvalues, rather than just the largest and smallest as in the classical κ. A recent survey on preconditioning is given in [31]. We emphasize that though many heuristics are given, the main measure of conditioning, e.g. [31], is κ.[2] However, in our view, κ has the misleading property that it is *inverse invariant*.[3]

From the discussions in [9, 23], the (*worst case*) condition number for the linear system

---

[1]Note that MATLAB/CVX has a new command DET_ROOTN that efficiently evaluates the geometric mean in the denominator of ω.

[2]Links: Who Invented the Matrix Condition Number? and What is the Condition Number of a Matrix?

[3]In What is the Condition Number of a Matrix?, the derivation for κ for $Ax = b$ with $\sigma_{\max}, \sigma_{\min}$ largest and smallest singular values for A, respectively, is that $\|b\| \leqslant \sigma_{\max}\|x\|, \|\Delta b\| \geqslant \sigma_{\min}\|\Delta x\|$. However, one can equally argue using $\|\Delta x\| \leqslant \sigma_{\max}(A^{-1})\|\Delta b\|, \|x\| \geqslant \sigma_{\min}(A^{-1})\|b\|$ to get $\frac{\|\Delta x\|\|b\|}{\|x\|\|\Delta b\|} \leqslant \kappa(A^{-1})$.

$Ax = b$ is

$$\text{cond}(A, b) := \lim_{\epsilon \downarrow 0} \sup_{\substack{\|\Delta A\| \leqslant \epsilon \|A\| \\ \|\Delta b\| \leqslant \epsilon \|b\|}} \frac{\|(A + \Delta A)^{-1}(b + \Delta b) - A^{-1}b\|}{\epsilon \|A^{-1}b\|}$$
$$= \kappa(A) + \frac{\|A^{-1}\|\|b\|}{\|A^{-1}b\|} \quad \left(= \kappa(A^{-1}) + \frac{\|A^{-1}\|\|b\|}{\|A^{-1}b\|} \right), \tag{1.2}$$

where we have added the equivalence for the inverse invariance for emphasis. The nonstandard condition number ω was proposed in [11]. Interestingly enough, the authors show that the inverse-sized BFGS and sized DFP [30] are obtained as optimal quasi-Newton updates with respect to this measure. In contrast to the worst case condition number κ, we argue that ω is a more average type condition number and provides a better measure for improving conditioning. Moreover, it distinguishes between the conditioning of A and A^{-1}. We illustrate that ω presents advantages with respect to the classic κ. Both are pseudoconvex over the open convex cone of positive definite matrices, $\mathbb{S}^n_{++}$; thus a local minimum is a global minimum. But, κ is differentiable if, and only if, both largest and smallest eigenvalues are singletons, while ω is differentiable on all of $\mathbb{S}^n_{++}$. This facilitates obtaining explicit formulae for optimal preconditioners and avoids expensive calculations, see e.g., [11] and Section 2.1, below.[4] Moreover, it is expensive to evaluate the classic condition number [22] as it uses both $\|A\|, \|A^{-1}\|$. For large scale, one often uses the ℓ_1 approximation in [22]. We show that we can find the exact value of the ω-condition number when a Cholesky or LU factorization is done. Finally, we show that ω, and particularly $\sqrt{\omega(A^{-2})}$, denoted ω^{-2}, provides a significantly better estimate for the true conditioning of a linear system. (Though ω^{-2} is currently for theoretical purposes only as we do not yet have an efficient way of exploiting it without calculating A^{-1}.)

1.2 Notation

We denote: $\mathbb{R}^n$ as the real Euclidean space of dimension n, and $\mathbb{R}^n_+, \mathbb{R}^n_{++}$ as the nonnegative and positive orthants, respectively; $\mathbb{R}^{m \times n}$ as the space of $m \times n$ matrices; $\mathbb{S}^n$ as the space of $n \times n$ symmetric matrices; $\mathbb{S}^n_+$ and $\mathbb{S}^n_{++}$ for the cone of positive semidefinite and positive definite $n \times n$ symmetric matrices, respectively; and $A \succeq 0$ (resp., > 0) as A is in $\mathbb{S}^n_+$ (respectively, $\mathbb{S}^n_{++}$). We use the Kronecker product and Hadamard (elementwise) product $A \otimes B, C \circ D$, respectively, with the matrix to vector columnwise vectorization $x = \text{vec}(X)$.

We use Diag $: \mathbb{R}^n \to \mathbb{R}^{n \times n}$ to denote the linear operator that maps a vector v into the diagonal matrix $\text{Diag}(v)$ whose diagonal is v. Its adjoint operator is denoted by diag $= \text{Diag}^*$.

For integers $t \geqslant s$, we let $[s, t] = \{s, s + 1, \ldots, t\}$. For a positive integer k, let $[k] = [1, k]$ and denote $t(k) = k(k + 1)/2$, *triangular number*.

---

[4]Since the original version of this paper was submitted, the recent report [15] (and many references therein) discusses numerical scalable algorithms for κ-optimal diagonal preconditioning. We have added relationships to this paper in this revised version. In particular, we present an alternative algorithm as well as illustrate that using the ω-optimal formula in the positive definite case has relatively no cost in evaluation, and is a better preconditioner.

For a differentiable function $f : \mathbb{R}^n \to \mathbb{R}$, we use ∇f for the gradient. If the dimension $n = 1$, we just write f' for the derivative of f. Given a nonempty open set $\Omega \subseteq \mathbb{R}^n$, a function $f : \Omega \to \mathbb{R}$ is said to be *pseudoconvex* on Ω if it is differentiable and

$$\nabla f(x)^T (y - x) \geqslant 0 \implies f(y) \geqslant f(x), \quad \forall x, y \in \Omega.$$

This implies that for an open convex set Ω and a *pseudoconvex function* $f : \Omega \to \mathbb{R}$, we have: $\nabla f(x) = 0$ is a necessary and sufficient condition for x to be a global minimizer of f in Ω, see, e.g., [28].

1.3 Outline and Main Results

The goal of this paper is to show that the use of ω as a condition number instead of κ is advantageous in certain aspects. In particular, we illustrate the efficacy of ω for preconditioning, error estimation, and low rank updating.

Section 2.1 introduces basic and new properties of ω. A *new* ω-optimal diagonal preconditioner, with respect to ω^{-2}, is proposed in Theorem 2.2. Efficiency and accuracy of computing ω is studied in Section 2.1.1. Indeed *the condition number of the condition number is the condition number* holds for κ but not for ω indicating that numerical calculations of ill-conditioned κ can be very inaccurate in contrast to ω.

In Section 2.2 we empirically illustrate that ω is a better indicator of conditioning for iterative solutions of linear equations. Moreover, Remark 2.5 provides the justification for using $\omega^{-2} := \sqrt{\omega(A^{-2})}$ as a measure and emphasizing the advantage over κ of not being inverse invariant. This includes empirical results for better clustering of eigenvalues, Figure 2.4 in Section 2.3. Though as mentioned above, ω^{-2} is currently considered only for theoretical purposes.

Section 2.4 provides an explicit formula for an incomplete upper triangular ω-optimal preconditioner. We include connections to preserving sparsity and to incomplete Cholesky preconditioners.

In Section 3, we derive ω-optimal conditioning for low rank updates of positive definite matrices. These updates often arise in the construction of generalized Jacobians.

Numerical results are in Section 4. We use the linear equations that involve positive definite matrices as well as the generalized Jacobians for our original motivation. We empirically illustrate that reducing the ω-condition number improves the performance of iterative methods for solving these linear systems.

Conclusions are provided in Section 5.

2 Properties and Optimal Preconditioning: ω vs κ

We now introduce basic and new properties of ω, and study the efficiency of its numerical evaluation. In addition, we empirically compare its effectiveness with κ for preconditioning, clustering of eigenvalues, and in estimating the actual conditioning of positive definite linear systems.

In particular, we derive the following explicitly found ω-optimal preconditioners (scalings): (i) diagonal (2.1); (ii) block diagonal (2.2); and (iii) incomplete upper triangular (2.12). Two further explicit formulae for ω-optimal preconditioners in the form of (iv) lower triangular with a subdiagonal, and (v) upper triangular with diagonal, are derived in the supplementary material to this paper[5]. Specifically, preconditioners (i)-(iii) maintain their sparsity properties after matrix inversion. Empirical comparisons with state-of-the-art sparse incomplete Cholesky preconditioners follow in Section 4.1.

2.1 Basic Properties and ω-Optimal Diagonal Preconditioner

For iterative solutions of linear systems a preconditioner S is often essential, e.g., for preconditioned conjugate gradients for $Ax = b, A > 0$, we solve $(S^T A S)\tilde{x} = \tilde{b} = S^T b$, $x = S\tilde{x}$, see e.g., [4, 16, 18]. Moreover, it is known that the simple scaling diagonal preconditioner using the norms of the columns of A is the optimal diagonal preconditioner with respect to ω and is efficient in practice, see [11, 32], i.e., ω validates the use of this specific diagonal preconditioner.[6] Various preconditioners based on (partial) factorizations of A, are compared in [18]. One is the QR-factorization. We note that scaling columns is an essential part of a QR-factorization. We see below that our ω-optimal preconditioners are related to a modified QR-factorization (Cholesky for positive definite systems). Moreover, convergence rates of iterative methods are correlated to clustering of eigenvalues of $A^T A$, see e.g., [19]. We see below in Section 2.2 that the ω-optimal preconditioners promote this property better than those for κ.

The optimal diagonal preconditioner is extended to the block diagonal case in [12]. We now summarize these and other basic properties of ω in the following Proposition 2.1. We include a proof of Proposition 2.1, Item 2, that is different than that provided in [11] so as to emphasize the extension to new formulae for ω-optimal preconditioners in Section 2.4.

Proposition 2.1 ([11, 12]). *The following statements hold.*

1. *ω is pseudoconvex on the cone of symmetric positive definite matrices; thus every stationary point is a global minimizer of ω.*

2. *Let V be a full rank $m \times n$ matrix, $n \leq m$. Then the optimal column scaling that minimizes ω is given by:*

$$d^* = (d_i^*) = \underset{d \in \mathbb{R}_{++}^n}{\operatorname{argmin}} \omega((V\operatorname{Diag}(d))^T (V\operatorname{Diag}(d))), \quad d_i^* = \frac{1}{\|V_{:,i}\|}, \ i \in [n], \quad (2.1)$$

where $V_{:,i}$ is the i-th column of V.

---

[5]This material can be found online in https://github.com/DavidTBelen/omega-condition-number

[6]In [15] the motivation for numerically finding diagonal κ-optimal preconditioners was the lack of theoretical validation. See also the near optimality results in [37]. Validation using ω is now provided in Proposition 2.1, Item 2. However, an improved diagonal preconditioner is provided in Theorem 2.2.

3 *Let V be a full rank $m \times n$ matrix, $n \leqslant m$, with block structure $V = \begin{bmatrix} V_1 & V_2 & \ldots & V_k \end{bmatrix}$,
$V_i \in \mathbb{R}^{m \times n_i}$. Then an* optimal corresponding block diagonal scaling

$$
D = \begin{bmatrix} D_1 & 0 & 0 & \ldots & 0 \\ 0 & D_2 & 0 & \ldots & 0 \\ \ldots & \ldots & \ldots & \ldots & \ldots \\ 0 & 0 & 0 & \ldots & D_k \end{bmatrix}, \quad D_i \in \mathbb{R}^{n_i \times n_i},
$$

that minimizes the measure ω, i.e.,

$$
\min \; \omega((VD)^T(VD)), \tag{2.2}
$$

over D block diagonal, is given by the factorization

$$
D_i D_i^T = \{V_i^T V_i\}^{-1}, \quad i \in [k].
$$

Proof. The results are proved in [11, 12]. We provide a new proof of Item 2 as it leads to different extensions below. Moreover, this proof illustrates the ease in differentiating ω and applying to derive *explicit* formulae.

Let $d := \operatorname{diag}(D), W := V^T V, w = \operatorname{diag}(W)$ and note that

$$
\begin{aligned}
\omega(d) := \omega((V\operatorname{Diag}(d))^T(V\operatorname{Diag}(d))) &= \frac{1}{n \det(V^T V)^{1/n}} \frac{\langle w, d \circ d \rangle}{\det(D)^{2/n}} \\
&=: K \frac{\sum_{i=1}^n w_i d_i^2}{\prod_{i=1}^n d_i^{2/n}} \\
&=: K \frac{f_w(d)}{g(d)},
\end{aligned}
$$

thus defining the constant $K > 0$ and functions $f_w, g : \mathbb{R}_{++}^n \to \mathbb{R}_{++}$. The reason for including this proof is to emphasize that V only appears in the numerator f_w of the function to be minimized as the denominator involves only d.

We now differentiate this pseudoconvex function with respect to d_i :

$$
\begin{aligned}
\frac{\partial \omega(d)}{\partial d_i} &= \frac{K}{g(d)^2} \left(g(d) 2 w_i d_i - f_w(d) \frac{2}{n} g(d) \frac{1}{d_i} \right) \\
&= \frac{2K}{g(d)} \left(w_i d_i - \frac{1}{n} f_w(d) \frac{1}{d_i} \right) \\
&= \frac{2K}{g(d)} \left(\frac{1}{d_i} - \frac{1}{n} f_w(d) \frac{1}{d_i} \right) \\
&= 0,
\end{aligned}
$$

since $w_i = \|V_{:,i}\|^2 = 1/d_i^2 \implies f_w(d) = n$.

∎

We now derive properties for the (square root) of the ω-condition number of A^{-2}, denoted ω^{-2}. The motivation for this is introduced below in (2.7).

8

Theorem 2.2. *Let $A \in \mathbb{S}_{++}^n$ and denote the Hadamard (elementwise) product $B := A^{-1} \circ A^{-1} \in \mathbb{S}_{++}^n \cap \mathbb{R}_+^{n \times n}$. Let $\bar{d} \in \mathbb{R}_{++}^n$, $\bar{D} = \mathrm{Diag}(\bar{d})$, be the solution of*

$$B\bar{d} = \mathrm{diag}(\bar{D}^{-1}) > 0,$$

and let

$$D = \bar{D}^{1/2}, \ d = \mathrm{diag}(D).$$

Then, $\bar{d}^T B \bar{d} = n$. Moreover, d provides the ω^{-2}-optimal scaling, i.e., the ω-optimal diagonal scaling $D = \mathrm{Diag}(d)$ of A with respect to A^{-2}:

$$d = \operatorname*{argmin}_{D=\mathrm{Diag}(d)>0} \omega(DA^{-1}DDA^{-1}D).$$

The corresponding optimal scaling (preconditioning) for solving $Ax = b$, with respect to the motivation for using A^{-2} in (2.7), is

$$(D^{-1}AD^{-1})(Dx) = D^{-1}b.$$

Proof. First note $\bar{d}^T B \bar{d} = \bar{d}^T \mathrm{diag}\, \bar{D}^{-1} = n$. From (2.7), to improve conditioning for the system $Ax = b$, we want to decrease $\omega(A^{-2})$. We restrict to a diagonal scaling and try to find:

$$
\begin{aligned}
d &= \operatorname*{argmin}_{D=\mathrm{Diag}(d)>0} \omega(DA^{-1}DDA^{-1}D) \\
&= \operatorname*{argmin}_{D=\mathrm{Diag}(d)>0} \omega(A^{-1}D^2A^{-1}D^2) \\
&= \operatorname*{argmin}_{\bar{D}=\mathrm{Diag}(\bar{d})>0} \frac{\frac{1}{n}\mathrm{tr}(A^{-1}\bar{D}A^{-1}\bar{D})}{\det\left(A^{-1}\bar{D}A^{-1}\bar{D}\right)^{1/n}}, \quad \bar{D} = D^2 \\
&= \operatorname*{argmin}_{\bar{D}=\mathrm{Diag}(\bar{d})>0} \frac{\det(A)^{\frac{2}{n}}}{n}\frac{\mathrm{tr}(A^{-1}\bar{D}A^{-1}\bar{D})}{\det\left(\bar{D}^{2/n}\right)}. \\
&= \operatorname*{argmin}_{\bar{D}=\mathrm{Diag}(\bar{d})>0} \frac{\mathrm{tr}(A^{-1}\bar{D}A^{-1}\bar{D})}{\det\left(\bar{D}^{2/n}\right)}.
\end{aligned}
\tag{2.3}
$$

We use the Kronecker product notation and Hadamard product notation $\otimes, \circ$, with the vectorization $\mathrm{vec}(\cdot)$, and obtain

$$\mathrm{tr}\, A^{-1}\bar{D}A^{-1}\bar{D} = \mathrm{vec}(\bar{D})^T A^{-1} \otimes A^{-1} \mathrm{vec}(\bar{D}) = \bar{d}^T A^{-1} \circ A^{-1}\bar{d} =: \bar{d}^T B\bar{d} =: \bar{f}(\bar{d}),$$

thus defining the positive definite matrix B and quadratic form $\bar{f}(\bar{d})$. We let

$$\bar{g}(\bar{d}) = \prod_{i=1}^n \bar{d}_i^{2/n} = \det\left(\bar{D}^{2/n}\right).$$

Then

$$\nabla \bar{f}(\bar{d}) = 2B\bar{d}, \quad \nabla \bar{g}(\bar{d}) = \left(\frac{2g(d)}{nd_i}\right) = \frac{2\bar{g}(\bar{d})}{n}\bar{D}^{-1}e.$$

9

From the minimization problem in (2.3), we want the stationary point

$$
\begin{aligned}
0 &= \tfrac{1}{\bar{g}(\bar{d})^2} \left(\bar{g}(\bar{d}) \nabla \bar{f}(\bar{d}) - \bar{f}(\bar{d}) \nabla \bar{g}(\bar{d}) \right) \\
&= 2\bar{g}(\bar{d}) B\bar{d} - 2\tfrac{\bar{f}(\bar{d})\bar{g}(\bar{d})}{n} \bar{D}^{-1} e \\
&= B\bar{d} - \tfrac{\bar{f}(\bar{d})}{n} \bar{D}^{-1} e.
\end{aligned}
$$

We normalize and get the two equations

$$
\bar{f}(\bar{d}) = n, \quad B\bar{d} = \operatorname{diag}(\bar{D}^{-1}).
$$

$\blacksquare$

Remark 2.3. *Though currently only of theoretical interest due to dependence on having* A^{-1}*, we note that solving for* $\bar{d}$ *in Theorem 2.2 can be done by e.g., Newton's method. If* B *is diagonal, then an explicit solution is* $\bar{d}_i := \frac{1}{\sqrt{B_{ii}}}$*. Note that* B *diagonal holds if, and only if,* A *is diagonal and then the optimal diagonal preconditioner is*

$$
\bar{D} = \operatorname{Diag}(\bar{d}) = \sqrt{B^{-1}} = A.
$$

Therefore, the optimal preconditioner for A *is* $\bar{D}^{-1/2} = A^{-1/2}$ *which agrees with our optimal* ω *preconditioner. In general, with* $\alpha := \bar{d}^T B \bar{d}$*, then* $\sqrt{\frac{n}{\alpha}}\bar{d}$ *provides a good starting point for Newton's method, as it is highly likely that the matrix* B *is significantly diagonally dominant. We solve*

$$
F(d) := \operatorname{Diag}(d) Bd - e = 0.
$$

The Jacobian with the matrix representation is

$$
F'(d)(\Delta d) = \operatorname{Diag}(d) B \Delta d + \operatorname{Diag}(\Delta d) Bd = \left[\operatorname{Diag}(d) B + \operatorname{Diag}(Bd) \right] (\Delta d).
$$

In our experiments Newton's method always converged in a few iterations, in fact 4 *iterations independent of* n*.*[7]

We now include the gradients of the condition numbers for use in the definitions below. In the case of κ, for simplicity and to avoid the use of subgradients, we assume that the largest and smallest eigenvalues are singletons.

Lemma 2.4. *Let* $A \in \mathbb{S}_{++}^n$ *with eigenvalues* $\lambda_1 \geqslant \lambda_2 \geqslant \ldots \geqslant \lambda_{n-1} \geqslant \lambda_n$*, with corresponding orthonormal eigenvectors* $v_1, \ldots, v_n$*. Then:*

1

$$
\begin{aligned}
\nabla \omega(A) &= \tfrac{1}{n \det(A)^{1/n}} \left(I - \tfrac{\operatorname{tr} A}{n} A^{-1} \right) \quad \text{is indefinite,} \\
&\text{with } \|\nabla \omega(A)\| = \tfrac{1}{n \det(A)^{1/n}} \max \left\{ 1 - \tfrac{\operatorname{tr} A}{n\lambda_1}, \tfrac{\operatorname{tr} A}{n\lambda_n} - 1 \right\}.
\end{aligned}
$$

---

[7]This could be a result of monotonicity arising from the nonnegativity of both F, F'. It is still an open question whether d can be found efficiently without explicitly finding A^{-1} first.

2 In addition, assume that the largest and smallest eigenvaules of A are singletons, i.e., $\lambda_1 > \lambda_2 \geqslant \ldots \geqslant \lambda_{n-1} > \lambda_n$. Then:

$$\nabla\kappa(A) = \frac{1}{\lambda_n}\left(v_1 v_1^T - \kappa(A)v_n v_n^T\right), \;\; \textit{is indefinite, with } \|\nabla\kappa(A)\| = \frac{\kappa(A)}{\lambda_n}.$$

Proof.   1 The gradient is

$$
\begin{aligned}
\nabla\omega(A) &= \frac{1}{n\det(A)^{2/n}}\left(\det(A)^{1/n}I - \frac{\operatorname{tr}A}{n}\det(A)^{\frac{1}{n}-1}\operatorname{adj}A\right) > 0;\\
&= \frac{1}{n\det(A)^{2/n}}\left(\det(A)^{1/n}I - \frac{\operatorname{tr}A}{n}\det(A)^{\frac{1}{n}-1}\operatorname{adj}A\right) > 0;\\
&= \frac{1}{n\det(A)^{1/n}}\left(I - \frac{\operatorname{tr}A}{n}A^{-1}\right),
\end{aligned}
$$

where $\operatorname{adj}A$ is the adjunct, the matrix of cofactors. The last expression follows from $A^{-1} = \frac{1}{\det(A)}\operatorname{adj}A$. The indefiniteness and norm follow from:

$$
\begin{aligned}
\lambda_{\max}(I - \tfrac{\operatorname{tr}A}{n}A^{-1}) &= \max_{\|x\|=1} x^T(I - \tfrac{\operatorname{tr}A}{n}A^{-1})x \\
&= 1 - \tfrac{\operatorname{tr}A}{n}\min_{\|x\|=1} x^T A^{-1}x \\
&= 1 - \tfrac{\operatorname{tr}A}{n\lambda_1}, && \text{with attainment at } x = v_1,\\
&> 0;
\end{aligned}
$$

$$
\begin{aligned}
\lambda_{\min}(I - \tfrac{\operatorname{tr}A}{n}A^{-1}) &= \min_{\|x\|=1} x^T(I - \tfrac{\operatorname{tr}A}{n}A^{-1})x \\
&= 1 + \tfrac{\operatorname{tr}A}{n}\min_{\|x\|=1}(-x^T A^{-1}x) \\
&= 1 - \tfrac{\operatorname{tr}A}{n}\max_{\|x\|=1} x^T A^{-1}x \\
&= 1 - \tfrac{\operatorname{tr}A}{n\lambda_n}, && \text{with attainement at } x = v_n,\\
&< 0.
\end{aligned}
$$

2 Since the eigenvalues are singletons, they are differentiable with gradients $v_1 v_1^T, v_n v_n^T$, respectively. The result follows from the definitions of the gradient of the fractional function κ, the spectral norm, and orthonormality of the eigenvectors.

■

2.1.1 Efficiency and Accuracy for Evaluating ω, κ

Since eigenvalue decompositions can be expensive, one issue with $\kappa(A)$ is how to estimate it efficiently when the size of matrix A is large. A survey of estimates and, in particular, estimates using the ℓ_1-norm, is given in [22, 24]. Extensions to sparse matrices and block-oriented generalizations are given in [21, 25]. Results from these papers form the basis of the *condest* command in MATLAB. More recently [15] deals with scalable methods for finding the κ-optimal diagonal preconditioner. This illustrates the difficulty in accurately estimating $\kappa(A)$. The main goal of this section is to discuss the difficulties that arise in the numerical evaluation of $\omega(A)$ and study how to estimate it accurately.

First, recall that the measure $\omega(A)$ can be calculated using the simple trace and determinant functions that do not require eigenvalue decompositions. Derivatives are given in Lemma 2.4 above.[8] However, for large n, the determinant is also numerically difficult to compute as it could easily result in an overflow $+\infty$ or 0 due to the limits of finite precision arithmetic, e.g., if the order of A is $n = 50$ and the eigenvalues $\lambda_i = .5, \forall i$, then the determinant $.5^n$ is zero to machine precision. A similar problem arises for e.g., $\lambda_i = 2, \forall i$ with overflow. In order to overcome this problem, we take the n-th root first and then the product, i.e., we define the value obtained from the spectral factorization as

$$\omega_{\text{eig}}(A) = \frac{\sum_{i=1}^n \lambda_i(A)/n}{\prod_{i=1}^n (\lambda_i(A)^{1/n})}.$$

We now let $A = R^T R = LUP$ denote the Cholesky and LU factorizations, respectively, with appropriate permutation matrix P. We assume that L is unit lower triangular. Therefore,

$$\det(A)^{1/n} = \det(R^T R)^{1/n} = \det(R)^{2/n} = \prod_{i=1}^n \left(R_{ii}^{2/n} \right). \tag{2.4}$$

Similarly,

$$\det(A)^{1/n} = \det(LUP)^{1/n} = \prod_{i=1}^n \left(|U_{ii}|^{1/n} \right). \tag{2.5}$$

Therefore, we find $\omega(A)$ with numerator $\text{tr}(A)/n$ and denominator given in (2.4) and (2.5), respectively:

$$\omega_R(A) = \frac{\text{tr}(A)/n}{\prod_{i=1}^n \left(R_{ii}^{2/n} \right)}, \quad \omega_{LU}(A) = \frac{\text{tr}(A)/n}{\prod_{i=1}^n \left(|U_{ii}|^{1/n} \right)}.$$

Tables 2.1 and 2.2 provide comparisons on the time and precision from the three different factorization methods. Each column presents different order of κ-condition number, while each row corresponds to different decompositions with different size n of the problem. We form the random matrix using $A = QDQ^T$ for random orthogonal Q and positive definite diagonal D. We then symmetrize $A \leftarrow (A + A^T)/2$ to avoid roundoff error in the multiplications. Therefore, we consider the evaluation using D as the *exact value* of $\omega(A)$, i.e.,

$$\omega(A) = \frac{\sum_{i=1}^n (D_{ii})/n}{\prod_{i=1}^n \left(D_{ii}^{1/n} \right)}.$$

Table 2.2 shows the absolute value of the difference between the exact ω-condition number and the ω-condition numbers obtained by making use of each factorization, namely, ω_{eig}, ω_R and ω_{LU}. Surprisingly, we see that both the Cholesky and LU decompositions give better results than the eigenvalue decomposition.

---

[8]Since the first version of this paper we have been made aware of the new CVX MATLAB function *det_rootn* that calculates $\det(A)^{1/n}$, the denominator of ω, using the Cholesky decomposition.

n	Fact.	order κ 1e2	order κ 1e3	order κ 1e4	order κ 1e5	order κ 1e6	order κ 1e7	order κ 1e8	order κ 1e9
	eig	5.5267e-02	5.7766e-02	5.2747e-02	5.9256e-02	6.0856e-02	6.2197e-02	5.5592e-02	5.7626e-02
500	R	1.1218e-02	8.0907e-03	7.5172e-03	8.4705e-03	9.2774e-03	8.5553e-03	8.1462e-03	7.9027e-03
	LU	2.2893e-02	1.8159e-02	1.8910e-02	2.0902e-02	2.0057e-02	2.0308e-02	1.9060e-02	1.8879e-02
	eig	3.0664e-01	2.8968e-01	2.6095e-01	2.7796e-01	5.7083e-01	5.9007e-01	5.8351e-01	5.9630e-01
1000	R	2.9328e-02	2.8339e-02	2.7869e-02	3.1909e-02	5.8628e-02	6.0873e-02	6.2429e-02	6.1074e-02
	LU	7.5011e-02	7.2666e-02	7.0497e-02	7.6778e-02	1.6313e-01	1.7313e-01	1.7666e-01	1.7326e-01
	eig	3.4794e+00	3.4804e+00	3.1916e+00	3.4386e+00	3.4235e+00	3.4766e+00	3.2327e+00	3.3704e+00
2000	R	3.5644e-01	3.5989e-01	2.9556e-01	3.6375e-01	3.5847e-01	3.5972e-01	3.2629e-01	3.4227e-01
	LU	9.0136e-01	9.0537e-01	7.1161e-01	8.7445e-01	8.6420e-01	8.8027e-01	8.1990e-01	8.1383e-01

Table 2.1: CPU sec. for evaluating $\omega(A)$, averaged over the same 10 random instances; eig, R, LU are eigenvalue, Cholesky, LU decompositions, respectively.

n	Fact.	order κ 1e2	order κ 1e3	order κ 1e4	order κ 1e5	order κ 1e6	order κ 1e7	order κ 1e8	order κ 1e9
	eig	1.5632e-13	2.7853e-12	2.2618e-10	1.2695e-08	8.9169e-07	5.4109e-05	2.2610e-03	1.7349e-01
500	R	1.7053e-13	2.5580e-12	1.0039e-10	1.1339e-08	4.9818e-07	2.6470e-05	1.3173e-03	1.6217e-01
	LU	1.5987e-13	2.4585e-12	1.0652e-10	1.1987e-08	5.1592e-07	2.1372e-05	1.3641e-03	1.4268e-01
	eig	2.1316e-13	2.1032e-12	8.7653e-11	4.6271e-09	3.1477e-07	1.9602e-05	9.9290e-04	7.6469e-02
1000	R	4.2633e-13	1.5632e-12	4.2235e-11	3.9297e-09	2.9562e-07	1.1498e-05	9.1506e-04	5.3287e-02
	LU	4.4054e-13	1.4850e-12	3.7858e-11	3.8287e-09	2.7390e-07	1.3820e-05	6.0492e-04	4.8568e-02
	eig	2.4336e-13	4.1780e-12	4.2019e-10	2.0080e-08	7.7358e-07	6.4819e-05	5.5339e-03	3.7527e-01
2000	R	4.3698e-13	2.0819e-12	5.0704e-11	2.3442e-09	1.8376e-07	8.9575e-06	5.5255e-04	4.8842e-02
	LU	4.3165e-13	2.2595e-12	2.3249e-11	2.5057e-09	1.5020e-07	6.0479e-06	5.4228e-04	4.4205e-02

Table 2.2: Precision of evaluation of $\omega(A)$ averaged over the same 10 random instances. eig, R, LU are eigenvalue, Cholesky, LU decompositions, respectively.

Moreover, Figure 2.1 illustrates a comparison of accuracy in evaluations of ω, κ. We use one positive definite matrix with spectral decomposition $A = QDQ^T$, and $n = 1000, density = 1e-4$ with $\kappa = 200$. We use perturbations of the eigenvalues $\|D(\epsilon) - D\|/\|D\| = 1e-8$ and reform $A(\epsilon) = QD(\epsilon)Q^T$. Figure 2.1 clearly shows that ω is calculated more accurately as the ill-conditioning grows. This relates to the *condition number of the condition number* in Remark 2.5.

Remark 2.5. *If we consider b as the input to a function G with output x, then a Taylor type argument gives to first order condition number as in* (2.9)

$$\text{cond}(G) = \frac{\|b\|}{\|G(b)\|} \frac{\|\Delta G\|}{\|\Delta b\|} \cong \|\nabla G(b)\| \frac{\|b\|}{\|G(b)\|},\tag{2.6}$$

a first order approximation for the condition number of G. Therefore, if G is one of κ, ω, we get the condition number of the condition number, see e.g., [8, 23] and the related result that for κ, the condition number of the condition number is the condition number. We have observed empirically that the condition number of ω is significantly smaller than the condition number of κ.

Let $G(\cdot) := A^{-1}(\cdot)$. Let v_i be orthonormal eigenvectors of A and $b = \sum \beta_i v_i, \Delta b = \sum \delta_i v_i$ and by abuse of notation

$$\|G(b)\|^2 = \langle \beta^2, \frac{1}{\lambda^2} \rangle = \sum_i \frac{\beta_i^2}{\lambda_i^2}, \qquad \|G(\Delta b)\|^2 = \langle \delta^2, \frac{1}{\lambda^2} \rangle = \sum_i \frac{\delta_i^2}{\lambda_i^2},$$
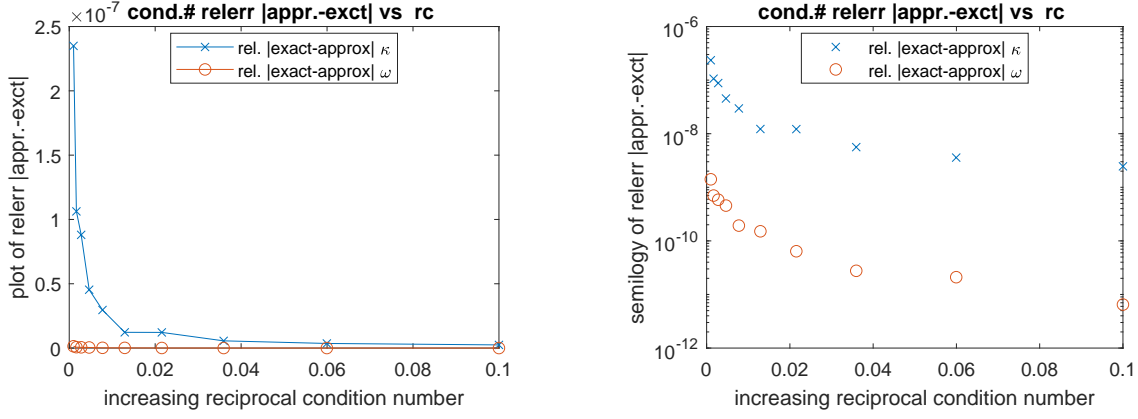
13

Figure 2.1: Comparing accuracy under perturbations for calculating ω, κ.

We now consider the conditioning for the linear system $Ax = b$. From (2.9) (squared) we have:

$$
\begin{aligned}
\text{cond}(A)^2 &= \frac{\|b\|^2}{\|G(b)\|^2} \frac{\|G(\Delta b)\|^2}{\|\Delta b\|^2} \\
&= \frac{\|b\|^2}{\|\Delta b\|^2} \frac{\langle \delta^2, \frac{1}{\lambda^2} \rangle}{\langle \beta^2, \frac{1}{\lambda^2} \rangle} \\
&\approx \frac{\|b\|^2}{\|\Delta b\|^2} \frac{\|\delta\|^2}{\|\beta\|^2} \frac{\text{tr} \, A^{-2}}{\prod_i \frac{1}{\lambda_i^{\frac{2}{n}}}} \\
&= \omega(A^{-2})
\end{aligned}
\tag{2.7}
$$

where the extended/strengthened AGM with an expected value gives the last approximation, i.e. we divide numerator and denominator by n and then apply the generalized AGM inequality, e.g., [14, page 6].

For example, to make $\omega(A^{-2})$ small we can use a diagonal scaling on left and right

$$\omega(DA^{-1}DDA^{-1}D).$$

The above suggests that we should use the measure $\sqrt{\omega(A^{-2})}$ rather than $\omega(A)$. the numerical tests appear to confirm this as well. Moreover, from [11, Prop. 2.1 (i)] and the inverse invariance of κ we have

$$1 \leqslant \sqrt{\omega(A^{-2})} \leqslant \sqrt{\kappa(A^{-2})} = \kappa(A) \leqslant \sqrt{4\omega(A^{-2})} = 2\sqrt{\omega(A^{-2})},$$

i.e., the measure $\sqrt{\omega(A^{-2})}$ is a valid condition number.

2.2 Error Analysis for Linear System $Ax = b$

We consider the linear system $Ax = b, A \in \mathbb{S}^n_{++}, b \in \mathbb{R}^n$. We are interested in understanding how small changes in the data affect the solution of the system. Let $x + \Delta x$ be a solution of the perturbed system

$$A(x + \Delta x) = b + \Delta b, \tag{2.8}$$

where $\Delta x, \Delta b \in \mathbb{R}^n$. The *condition number* aims to be a measure on how strongly a relative error in the data affects the relative error in the solution [35]. Therefore, it can be estimated as the ratio

$$\text{cond} :\approx \frac{\|\Delta x\|\|b\|}{\|x\|\|\Delta b\|} \qquad \text{(rel. error output/rel. error input).} \qquad (2.9)$$

Note that the above ratio depends on the choice of the perturbation Δb, as well as on the choice of the norms. $\kappa = \kappa(A) = \lambda_{\max}(A)/\lambda_{\min}(A)$ is taken as a *worst case* estimator of the condition number as the inequality

$$cond = \frac{\|\Delta x\|\|b\|}{\|x\|\|\Delta b\|} \leqslant \kappa(A) \qquad (\leqslant 4\omega(A)^n \ [11]),$$

holds for all $\Delta b \in \mathbb{R}^n$. See (1.2) above, and e.g., [22] or [36, Chapter 7] for further details.

In the literature, the system is termed *ill-conditioned* if κ is *large*, and termed *well-conditioned* otherwise.[9] We now study the correlation between the three estimates of cond: (i) κ, (ii) ω, and (iii)ω^{-2}, with the estimate of cond evaluated by sampling. We sample as follows:

1. Generate 200 linear systems $A_i x = b_i, i \in [200]$, where the positive definite matrices A_i are randomly generated with uniformly distributed eigenvalues in $(0,1)$ and the column vectors are set as $b_i := A_i x_i$, with x_i sampled from the standard normal distribution.

2. For each $i \in [200]$, we generate 1000 perturbations $\{\Delta b_j\}_{j \in [1000]}$ of norm 10^{-6} and set $\Delta x_j := A_i^{-1} \Delta b_j$. Then, for each $j \in [1000]$ we compute the relative residual ratio in (2.9). We then average over j to yield an estimate of the condition number, $\text{cond}(A_i)$, of the i-th system.

3. We then check the resulting correlation between the following vectors (i) $(\kappa(A_i))_{i=1}^{200}$, (ii) $(\omega(A_i))_{i=1}^{200}$, and (iii) $(\sqrt{\omega(A_i^{-2})})_{i=1}^{200}$, with $(\text{cond}(A_i))_{i=1}^{200}$, by comparing the corresponding linear regression models.

Figure 2.2, page 16, reveals a significant linear correlation between cond and ω, with a correlation coefficient of 0.9251 for ω^{-2} and of 0.9062 for ω; whereas in contrast, cond and κ are not linearly correlated as the correlation coefficient is 0.4530.

The same experiment with the eigenvalues of the matrices $\{A_i\}_{i \in [200]}$ generated from the *normal standard distribution* are displayed in Figure 2.3, page 17. We get correlation coefficients: $0.7982 > 0.4847 > 0.0295$ with $\omega^{-2}, \omega, \kappa(A)$, respectively, i.e., we cannot conclude existence of a linear relation between cond and κ, ω.

2.3 Eigenvalue Clustering

As stated above, preconditioning is essential for iterative methods for solving linear systems. And, many of the convergence analysis results depend on clustering of eigenvalues, e.g., [19].

---

[9]$\kappa(A)$ is also used to measure error that arises from perturbations in A: $\frac{\|\Delta x\|}{\|x+\Delta x\|} \leqslant \kappa(A)\frac{\|\Delta A\|}{\|A\|}$. The results are essentially equivalent.
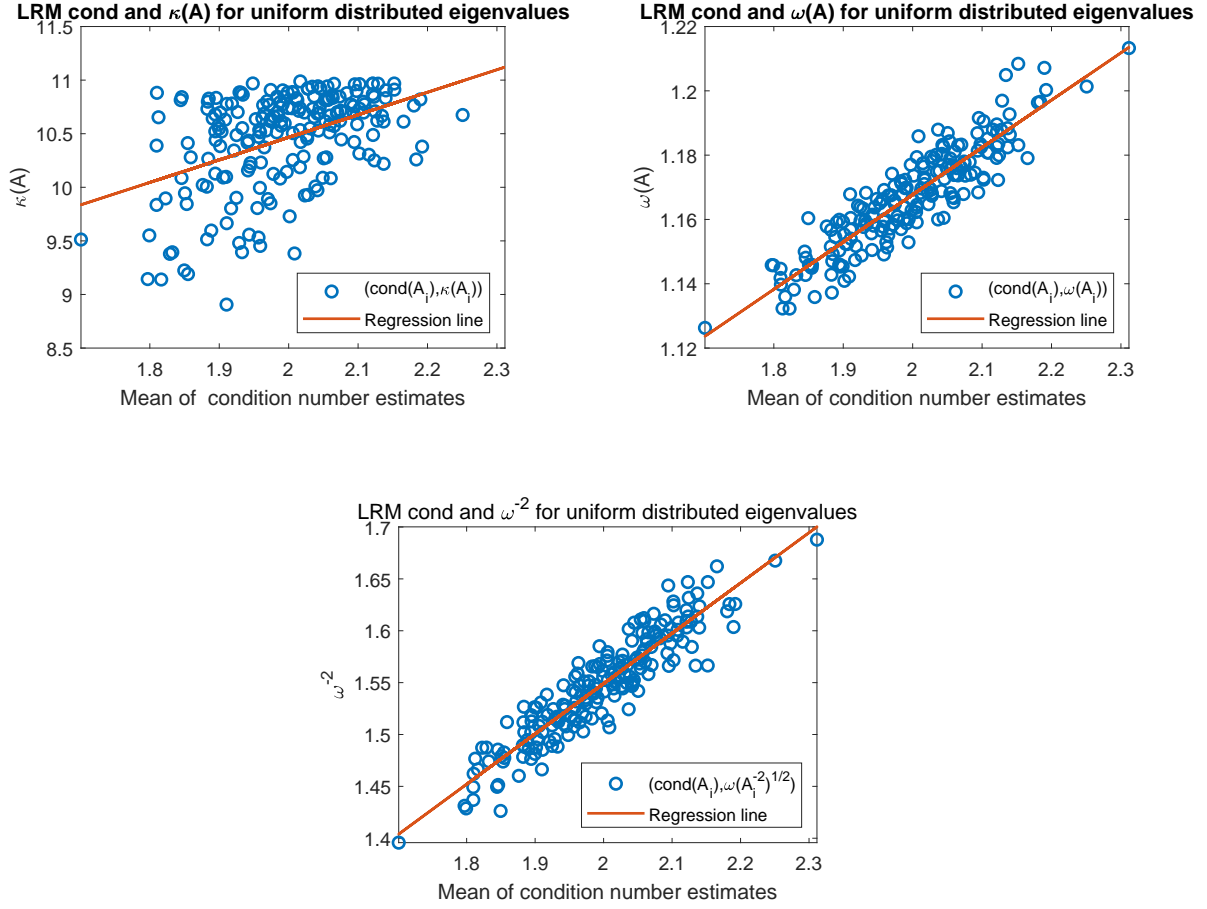
Figure 2.2: Linear regression models (LRM) between cond and: κ, ω, ω^{-2}, respectively; uniformly distributed eigenvalues.

A typical comparison for the eigenvalues of $A > 0$ after preconditioning with the optimal $\kappa, \omega, \omega^{-2}$ diagonal preconditioners is given in Figure 2.4. (The corresponding MATLAB code is available online in https://github.com/DavidTBelen/omega-condition-number file "Erroranalysistests_inv.m".) Figure 2.4 clearly shows the improved clustering of eigenvalues, as κ essentially shifts the eigenvalues to reduce the $\lambda_{\max}/\lambda_{\min}$ ratio, while both ω measures move the eigevenvalues towards 1 and promote clustering. We see this in the large number of eigenvalues that are close to the mean value for the ω optimal preconditioned matrices in the second of the two figures in Figure 2.4. In the supplementary material to this paper we include the distributions of eigenvalues pre-post (ω and κ) preconditioning for an appropriate subset of the test matrices at the SuiteSparse Matrix Collection [27]. These show similar improved clustering for ω over κ. The effect on iterations for solving linear systems is given in Section 4, below.
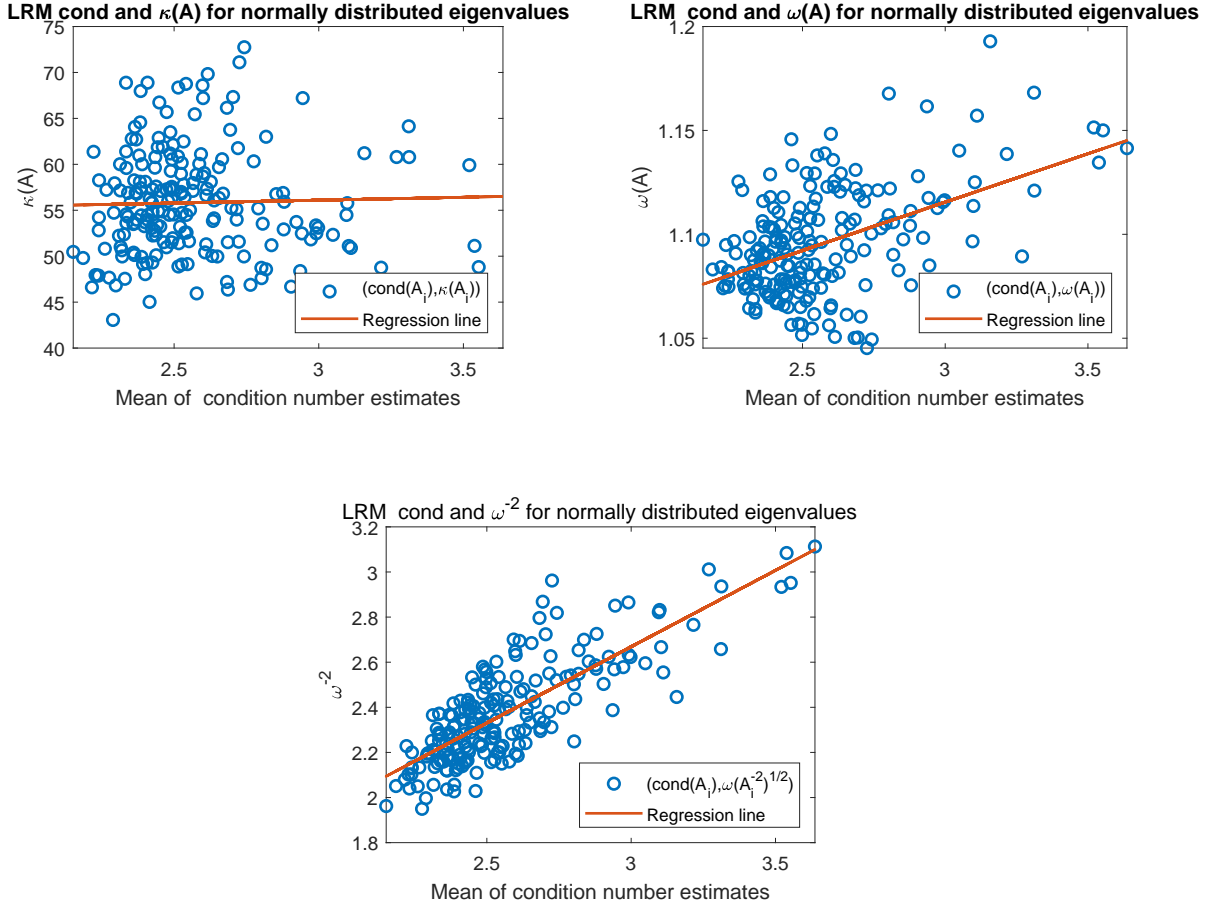
Figure 2.3: Linear regression models (LRM) between cond and: κ, ω, ω^{-2}, respectively; normally distributed eigenvalues.

Table 2.3 further empirically illustrates that the ω condition number correlates better with the number of iterations needed for **pcg**. In the table we fix κ by fixing the smallest $\lambda_{\min}$ and largest $\lambda_{\max}$ of the eigenvalues and also fix the orthogonal matrix of eigenvectors Q in the spectral decomposition $A = Q \operatorname{Diag}(\lambda) Q^T$. We generate 5 random sparse matrices whose $(n-2)$-non extremal eigenvalues have the same mean but follow different distributions, i.e., they can be either uniformly distributed or normally distributed with different standard deviations. Hence their ω values vary. We generate a column vector b which is in the range of the sum of the 5 matrices plus a random error term of order 10^{-7}. Finally, we use MATLAB's **pcg** with a tolerance set at 10^{-6} and solve the 5 linear systems given by b and each of the 5 matrices. The correlation (monotonicity) between the std, ω, iteration count in Table 2.3 is clear. We include the fifth and sixth columns to highlight the significant correlation in the percentage reduction in the values of ω and the number of iterations.
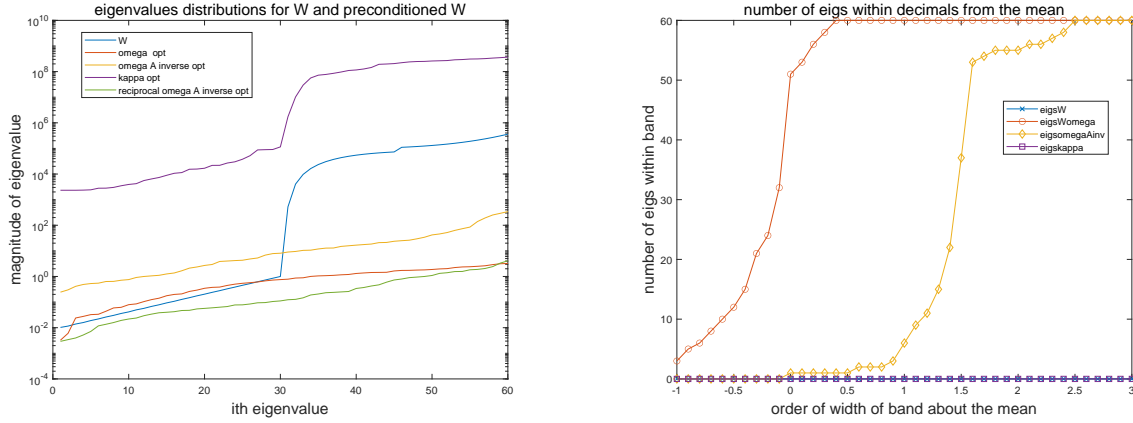
17

Figure 2.4: Comparison for clustering of eigenvalues pre-post preconditioning

Distribution	std	ω	# Iterations	% Reduct. ω	% Reduct. Iter.	Rel. Resid.
Uniform	1.163e+11	1.3682	264	-	-	9.738e-07
Normal	4.031e+10	1.0308	14	91.63	94.70	5.596e-07
Normal	2.022e+10	1.0143	8	96.11	96.97	9.834e-07
Normal	9.205e+09	1.0101	6	97.25	97.73	4.432e-07
Normal	6.095e+09	1.0095	5	97.42	98.11	6.067e-07

Table 2.3: correlations: eigenvals standard deviation (std); ω; iterations

2.4 Incomplete Upper Triangular ω-Optimal Preconditioner

Approximations of the inverse of the Cholesky decomposition are widely used as preconditioners for linear systems. It is easy to verify that the inverse of the Cholesky provides the optimal ω preconditioner. Indeed, let $W = R^T R$ be the Cholesky decomposition of W. Then $\omega(R^{-T}WR^{-1}) = \omega(I) = 1$. However, it is well-known that sparsity can be lost when finding R and R^{-1}. Therefore, permutation techniques are used when finding an incomplete Cholesky decomposition, e.g., [18].

In this section we see an interesting relationship between finding an ω-optimal *incomplete upper triangular preconditioner* and an incomplete Cholesky factorization, see Theorem 2.7. Specifically, given an integer $2 \leqslant k \leqslant n$, let $\alpha = (\alpha_{1,2}, \alpha_{1,3}, \alpha_{2,3}, \ldots, \alpha_{1,k}, \ldots, \alpha_{k-1,k}) \in \mathbb{R}^{t(k-1)}$

and $d \in \mathbb{R}^n$. We consider a preconditioner in the form of

$$\mathrm{D}_{+\mathrm{tk}}(d,\alpha) = \mathrm{Diag}(d) + \mathrm{Trir}_k(\alpha)$$

$$= \begin{pmatrix}
d_1 & \alpha_{1,2} & \alpha_{1,3} & \dots & \alpha_{1,k} & 0 & \dots & 0 \\
0 & d_2 & \alpha_{2,3} & \dots & \alpha_{2,k} & 0 & \dots & 0 \\
0 & 0 & d_3 & \ddots & \alpha_{3,k} & 0 & \dots & 0 \\
0 & \dots & \dots & \ddots & \alpha_{k-1,k} & 0 & \dots & 0 \\
\vdots & \dots & \dots & \dots & d_k & 0 & \dots & 0 \\
0 & \dots & \dots & \dots & 0 & d_{k+1} & \dots & 0 \\
0 & \dots & \dots & \dots & \dots & 0 & \ddots & 0 \\
0 & \dots & \dots & \dots & \dots & 0 & \dots & d_n
\end{pmatrix}, \tag{2.10}$$

where the linear mapping $\mathrm{Trir}_k : \mathbb{R}^{t(k-1)} \to \mathbb{R}^{n\times n}$ is defined accordingly. Its adjoint operator is $\mathrm{Trir}_k^* = \mathrm{trir}_k : \mathbb{R}^{n\times n} \to \mathbb{R}^{t(k-1)}$, $M \mapsto (M_{1,2}, M_{1,3}, M_{2,3}, \dots, M_{1,k}, \dots, M_{k-1,k})$. Moreover, the inverse maintains the same structure and sparsity pattern.

Observe that if $k = n$ then $\mathrm{D}_{+\mathrm{tk}}(d,\alpha)$ returns a complete upper triangular matrix. In that case it trivially follows that the ω-optimal preconditioner will be given by the Cholesky decomposition. In any case, even when $k < n$, the ω-optimal incomplete upper triangular preconditioner will be related to the Cholesky factorization. Therefore, we first recall the following recursive formula for computing the latter.

Remark 2.6 (Recursive formula for the Cholesky decomposition). *Let $W \in \mathbb{S}^k$ be a positive definite matrix and let $W = R^T R$ be the Cholesky decomposition of W. We recall that the upper triangular Cholesky factor R admits the following recursive columnwise construction for $j \in [k]$:*

$$R_{i,j} = \frac{1}{R_{i,i}}\left(W_{i,j} - \sum_{t=1}^{i-1} R_{t,j}\, R_{t,i} \right), \quad \text{for } i \in [j-1],$$

$$R_{j,j} = \sqrt{ W_{j,j} - \sum_{t=1}^{j-1} R_{t,j}^2 }. \tag{2.11}$$

Theorem 2.7. *Let $W \in \mathbb{S}_{++}^n$, and let $W_{1:k,1:k} = R^T R$ be the Cholesky decomposition, $k \in [n]$. The ω-optimal incomplete upper triangular preconditioner in the form of (2.10) for W, i.e.,*

$$(\bar{d}, \bar{\alpha}) := \operatorname*{argmin}_{(d,\alpha)\in\mathbb{R}_{++}^n \times \mathbb{R}^{t(k-1)}} \omega\left(\mathrm{D}_{+\mathrm{tk}}(d,\alpha)^T W\, \mathrm{D}_{+\mathrm{tk}}(d,\alpha) \right), \tag{2.12}$$

is given by

$$\begin{aligned}
\bar{d}_j &= R_{j,j}^{-1}, & \text{for } j \in [k]; \\
\bar{d}_j &= W_{j,j}^{-1/2}, & \text{for } j \in [k+1,n]; \\
\bar{\alpha}_{i,j} &= -\frac{1}{R_{i,i}}\left(\sum_{s=i+1}^{j-1} R_{i,s}\bar{\alpha}_{s,j} + R_{i,j}\bar{d}_j \right), & \text{for } k \geqslant j > i \geqslant 1.
\end{aligned} \tag{2.13}$$

19

We get, by abuse of notation with (2.13),

$$\mathrm{D}_{+tk}(\bar{d}, \bar{\alpha}) = \mathrm{blkdiag}\left(R^{-1}, \mathrm{Diag}\left(\bar{d}_{[k+1,n]}\right)\right). \tag{2.14}$$

Proof. We divide the proof into three claims.

Claim 1: The ω-optimal D_{+tk} preconditioner is obtained by $(\bar{d}, \bar{\alpha})$ solving the nonlinear system

$$\begin{bmatrix} \mathrm{diag}\, W\left(\mathrm{Diag}(\bar{d}) + \mathrm{Trir}_k(\bar{\alpha})\right) \\ \mathrm{trir}_k\, W\left(\mathrm{Diag}(\bar{d}) + \mathrm{Trir}_k(\bar{\alpha})\right) \end{bmatrix} = \begin{pmatrix} \bar{d}^{-1} \\ 0 \end{pmatrix}, \tag{2.15}$$

where $\bar{d}^{-1} = (\bar{d}_1^{-1}, \ldots, \bar{d}_n^{-1})^T$.

In order to prove this, and to ease the notation, fix W and consider the ω-condition number, f and g as functions of a pair $(d, \alpha) \in \mathbb{R}_{++}^n \times \mathbb{R}^{t(k-1)}$. Namely, we set

$$\omega_{+tk}(d, \alpha) = \frac{f_{+tk}(d, \alpha)}{g_{+tk}(d, \alpha)} := \frac{\mathrm{tr}\left(\mathrm{D}_{+tk}(d, \alpha)^T W\, \mathrm{D}_{+tk}(d, \alpha)\right)/n}{\det\left(\mathrm{D}_{+tk}(d, \alpha)^T W\, \mathrm{D}_{+tk}(d, \alpha)\right)^{1/n}}.$$

Alternatively, we can rewrite f_{+tk} as

$$\begin{aligned} f_{+tk}(d, \alpha) &= \frac{1}{n}\,\mathrm{tr}\left(\mathrm{D}_{+tk}(d, \alpha)^T W\, \mathrm{D}_{+tk}(d, \alpha)\right) \\ &= \frac{1}{n}\left\langle \mathrm{D}_{+tk}{}^* W\, \mathrm{D}_{+tk}(d, \alpha), \begin{pmatrix} d \\ \alpha \end{pmatrix} \right\rangle. \end{aligned} \tag{2.16}$$

Hence,

$$\begin{aligned} \nabla f_{+tk}(d, \alpha) &= \frac{2}{n}\,\mathrm{D}_{+tk}{}^* W\, \mathrm{D}_{+tk}(d, \alpha) \\ &= \frac{2}{n}\begin{bmatrix} \mathrm{diag}\, W\left(\mathrm{Diag}(d) + \mathrm{Trir}_k(\alpha)\right) \\ \mathrm{trir}_k\, W\left(\mathrm{Diag}(d) + \mathrm{Trir}_k(\alpha)\right) \end{bmatrix}. \end{aligned} \tag{2.17}$$

On the other hand, we have that

$$g_{+tk}(d, \alpha) = \det(W)\left(\prod_{i=1}^{n} d_i\right)^{\frac{2}{n}} \quad \text{and} \quad \nabla g_{+tk}(d, \alpha) = \frac{2}{n} g_{+tk}(d, \alpha)\begin{pmatrix} d^{-1} \\ 0 \end{pmatrix},$$

where $d^{-1} = (d_1^{-1}, \ldots, d_n^{-1})^T \in \mathbb{R}_{++}^n$.

Therefore, the optimality condition for the pseudoconvex function ω_{+t} is given by

$$\nabla \omega_{+tk}(d, \alpha) = K\left(\mathrm{D}_{+tk}{}^* W\, \mathrm{D}_{+tk}(d, \alpha) - f_{+tk}(d, \alpha)\begin{pmatrix} d^{-1} \\ 0 \end{pmatrix}\right) = 0, \tag{2.18}$$

with $K := 2/(n\, g_{+tk}(d, \alpha)) > 0$. Finally, observe that it suffices to obtain $(\bar{d}, \bar{\alpha}) \in \mathbb{R}_{++}^n \times \mathbb{R}^{t(k-1)}$ such that

$$\mathrm{D}_{+tk}{}^* W\, \mathrm{D}_{+tk}(\bar{d}, \bar{\alpha}) - \begin{pmatrix} \bar{d}^{-1} \\ 0 \end{pmatrix} = 0, \tag{2.19}$$

20

as by (2.16) this immediately implies

$$f_{+tk}(\bar{d}, \bar{\alpha}) = \frac{1}{n} \left\langle \begin{pmatrix} \bar{d}^{-1} \\ 0 \end{pmatrix}, \begin{pmatrix} \bar{d} \\ \bar{\alpha} \end{pmatrix} \right\rangle = 1,$$

which in turn would yield (2.18). Thus, (2.19) together with (2.17) concludes this part of the proof.

Claim 2: A solution $(\bar{d}, \bar{\alpha})$ to (2.15) is given by $\bar{d}_i = W_{i,i}^{-1/2}$, for $i \in [k+1, n]$, and with

$$Q := \mathrm{Diag}(\bar{d}_{1:k}) + \mathrm{Triu}(\bar{\alpha}) \tag{2.20}$$

being the inverse of the Cholesky decomposition of the matrix $W_{1:k,1:k}$.

We start by fixing notation. Let $\widehat{W} := W_{1:k,1:k}$ and $\widetilde{W} := W_{k+1:n,k+1:n}$. Recall the definition of the operator Triu which applied to a vector $\alpha = (\alpha_{1,2}, \alpha_{1,3}, \alpha_{2,3}, \dots, \alpha_{1,k}, \dots, \alpha_{k-1,k}) \in \mathbb{R}^{t(k-1)}$ returns the upper triangular matrix $\mathrm{Triu}(\alpha) = T \in \mathbb{R}^{k \times k}$ such that $T_{i,j} = \alpha_{i,j}$ if $1 \leqslant i < j \leqslant n$, and $T_{i,j} = 0$ otherwise. The adjoint of Triu is denoted as triu. Then the system (2.15) can be split into the two equations

$$\mathrm{diag}\, W \left(\mathrm{Diag}(\bar{d}) + \mathrm{Trir}_k(\bar{\alpha}) \right) = \begin{bmatrix} \mathrm{diag}\, \widehat{W} \left(\mathrm{Diag}(\bar{d}_{1:k}) + \mathrm{Triu}(\bar{\alpha}) \right) \\ \mathrm{diag}\, \widetilde{W} \, \mathrm{Diag}(\bar{d}_{k+1:n}) \end{bmatrix} = \bar{d}^{-1} \tag{2.21}$$

and

$$\mathrm{trir}_k\, W \left(\mathrm{Diag}(\bar{d}) + \mathrm{Trir}_k(\bar{\alpha}) \right) = \mathrm{triu}\, \widehat{W} \left(\mathrm{Diag}(\bar{d}_{1:k}) + \mathrm{Triu}(\bar{\alpha}) \right) = 0. \tag{2.22}$$

Observe that the variables $\bar{d}_{k+1}, \dots, \bar{d}_n$ only appear in the lower block of (2.21), that can be directly solved to obtain $\bar{d}_i = W_{i,i}^{-1/2}$ for all $i \in [k+1, n]$.

On the other hand, the variables $\bar{d}_1, \dots, \bar{d}_k$ and $\bar{\alpha}$ are present in (2.22) and the upper block of (2.21). Nonetheless, by taking into account that if $n = k$ then $\mathrm{Triu} = \mathrm{Trir}_k$, it is easy to check that these equations define the ω-optimal triangular preconditioner of the matrix $\widehat{W} \in \mathbb{S}_{++}^k$. Therefore we conclude that Q coincides with the inverse of the Cholesky factorization of $\widehat{W}$.

Claim 3: Let $Q := \mathrm{Diag}(\bar{d}_{1:k}) + \mathrm{Triu}(\bar{\alpha})$ be the inverse of the Cholesky decomposition of $\widehat{W}$. Then $(\bar{d}_{1:k}, \bar{\alpha})$ is given as in (2.13).

Let $\widehat{W} = R^T R$ be the Cholesky decomposition of $\widehat{W}$, where

$$R = \begin{pmatrix} R_{1,1} & R_{1,2} & R_{1,3} & \dots & R_{1,k} \\ 0 & R_{2,2} & R_{2,3} & \dots & R_{2,k} \\ \vdots & \dots & \dots & \dots & \vdots \\ 0 & \dots & \dots & 0 & R_{k,k} \end{pmatrix},$$

and the entries are given as in (2.11). Let $Q = R^{-1}$ be the matrix defined in (2.20). We now use the equation $RQ = \mathrm{Id}$ to obtain an expression of Q in terms of R. We have:

$$\mathrm{Id} = \begin{pmatrix} R_{1,1} & R_{1,2} & R_{1,3} & \dots & R_{1,k} \\ 0 & R_{2,2} & R_{2,3} & \dots & R_{2,k} \\ \vdots & \dots & \ddots & \dots & \vdots \\ \vdots & \dots & \dots & R_{k-1,k-1} & R_{k-1,k} \\ 0 & \dots & \dots & 0 & R_{k,k} \end{pmatrix} \begin{pmatrix} \bar{d}_1 & \bar{\alpha}_{1,2} & \bar{\alpha}_{1,3} & \dots & \bar{\alpha}_{1,k-1} & \bar{\alpha}_{1,k} \\ 0 & \bar{d}_2 & \bar{\alpha}_{2,3} & \dots & \dots & \bar{\alpha}_{2,k} \\ \vdots & \dots & \dots & \dots & \dots & \vdots \\ \vdots & \dots & \dots & \dots & \bar{d}_{k-1} & \bar{\alpha}_{k-1,k} \\ 0 & \dots & \dots & \dots & 0 & \bar{d}_k \end{pmatrix}.$$

For each column $j \in [k]$ of Q, this leads to the following linear system of j equations:

$$1 = R_{j,j}\,\bar{d}_j, \tag{2.23a}$$

$$0 = R_{j-1,j-1}\,\bar{\alpha}_{j-1,j} + R_{j-1,j}\,\bar{d}_j, \tag{2.23b}$$

$$\vdots$$

$$0 = R_{j-\ell+1,j-\ell+1}\,\bar{\alpha}_{j-\ell+1,j} + \sum_{s=j-\ell+2}^{j-1} R_{j-\ell+1,s}\,\bar{\alpha}_{s,j} + R_{j-\ell+1,j}\,\bar{d}_j, \tag{2.23c}$$

$$\vdots$$

$$0 = R_{1,1}\,\bar{\alpha}_{1,j} + \sum_{s=2}^{j-1} R_{1,s}\,\bar{\alpha}_{s,j} + R_{1,j}\,\bar{d}_j. \tag{2.23d}$$

Equation (2.23a) readily implies that $\bar{d}_j = R_{j,j}^{-1}$ for all $j \in [k]$. Moreover, for any $\ell \in [2,j]$, we can solve (2.23c) for getting an expression for $\bar{\alpha}_{j-\ell+1,j}$ in terms of $\bar{d}_j, \bar{\alpha}_{j-1,j}, \dots, \bar{\alpha}_{j-\ell+2,j}$. This yields

$$\bar{\alpha}_{j-\ell+1,j} = -\frac{1}{R_{j-\ell+1,j-\ell+1}} \left(\sum_{s=j-\ell+2}^{j-1} R_{j-\ell+1,s}\,\bar{\alpha}_{s,j} + R_{j-\ell+1,j}\,\bar{d}_j \right), \tag{2.24}$$

which concludes Claim 3 and the proof.

■

Finally, observe that (2.14) suggests that the ω-optimal incomplete upper triangular preconditioner can be efficiently computed by exploiting a Cholesky decomposition of a submatrix of the original matrix.

3 Optimal Conditioning for Generalized Jacobians

We now consider the problem of improving conditioning for low rank updates of very ill-conditioned (close to singular) positive definite matrices.

3.1 Preliminaries

More precisely, given a positive definite matrix $A \in \mathbb{S}_{++}^n$ and a matrix $U \in \mathbb{R}^{n \times t}$ with $t << n$, we aim to find $\gamma \in \mathbb{R}^t$ so as to minimize the condition number of the low rank update

$$A + U\,\mathrm{Diag}(\gamma)U^T. \tag{3.1}$$

This kind of updating arises when finding generalized Jacobians in nonsmooth optimization. We provide insight on the problem in the following Example 3.1.

Example 3.1 (Generalized Jacobians). *In many nonsmooth and semismooth Newton methods one aims to find a root of a function $F : \mathbb{R}^n \to \mathbb{R}^n$ of the form*

$$F(y) := B(v + B^T y)_+ - c,$$

where $B \in \mathbb{R}^{n \times m}$, $v \in \mathbb{R}^m$, $c \in \mathbb{R}^n$ and $(\cdot)_+$ denotes the projection onto the nonnegative orthant, e.g., [3, 26, 33]. At every iteration of these algorithms a generalized Jacobian of F of the form

$$J := \sum_{i \in \mathcal{I}_+} B_i B_i^T + \sum_{j \in \mathcal{I}_0} \gamma_j B_j B_j^T, \ \ with \ \gamma_j \in [0, 1],$$

is computed. Here B_i and B_j denote columns of B over the set of indices

$$\begin{aligned}
\mathcal{I}_+ &:= \{i \in [m] : (v + B^T y)_i > 0\}; \ and \\
\mathcal{I}_0 &:= \{j \in [m] : (v + B^T y)_j = 0 \ and \ (B_j)_{j \in \mathcal{I}_0} \ is \ a \ maximal \ linearly \ independent \ set\}.
\end{aligned}$$

The generalized Jacobian J, is usually singular. It is used to obtain a Newton direction $d \in \mathbb{R}^n$ by solving a least-square problem for the system $(J + \epsilon I) d = -F(y)$, where ϵI, with $\epsilon > 0$, is analogous to the regularization term of the well-known Levenberg–Marquardt *method. Thus, this linear system is in general very ill-conditioned. This makes preconditioning by optimal updating appropriate.*

The optimal preconditioned update can be done in our framework as we start with

$$A := \sum_{i \in \mathcal{I}_+} B_i B_i^T + \epsilon I, \quad U = [B_j]_{j \in \mathcal{I}_0}, \tag{3.2}$$

and then find an optimal low rank update as in (3.1); done with additional box constraints on γ, namely, $\gamma \in [0, 1]^t$.

Similar conditioning questions also appear in the normal equations matrix, ADA^T, in interior point methods, e.g., modifying the weights in D appropriately to avoid ill-conditioning [6, 17]. For other related work on minimizing condition numbers for low rank updates see, e.g., [5, 20].

Here, we propose obtaining optimal conditioning of the update (3.1) based on using the measure ω instead of the classic κ. As illustrated above, ω presents many advantages over κ when it comes to e.g., differentiability and avoiding both the inverse invariance and emphasis on worst case scenario. Our empirical results show a significant decrease in the number of iterations for a requested accuracy in the residual.

3.2 Optimal Conditioning for Rank One Updates

We first consider the special case where the update is rank one. Related eigenvalue results for rank one updates are well known in the quasi-Newton literature, e.g., [10, 34]. We include this special rank one case as it yields interesting results. The general rank-t update is studied in Section 3.3, below.

Theorem 3.2. *Suppose we have a given $A \in \mathbb{S}_{++}^n$ and $u \in \mathbb{R}^n$. Let $A = QDQ^T$ be the (orthogonal) spectral decomposition of A. Let $U = uu^T$ and define the rank one update*

$$A(\gamma) = A + \gamma U, \ \gamma \in \mathbb{R}.$$

Set

$$w = D^{-1/2}Q^T u, \tag{3.3}$$

and

$$\gamma^* = \frac{\mathrm{tr}(A)\|w\|^2 - n\|u\|^2}{(n-1)\|u\|^2\|w\|^2}. \tag{3.4}$$

Then, $\gamma^* \in \,]-\|w\|^{-2}, +\infty[\,$ provides the ω-optimal conditioning, i.e.,

$$\gamma^* = \operatorname*{argmin}_{A(\gamma)>0} \omega(\gamma). \tag{3.5}$$

Proof. Let

$$f(\gamma) := \mathrm{tr}(A(\gamma))/n \quad \text{and} \quad g(\gamma) := \det(A(\gamma))^{1/n}.$$

We want to find the optimal γ to minimize the condition number

$$\omega(\gamma) = f(\gamma)/g(\gamma)$$

subject to $A(\gamma)$ being positive definite. By Proposition 2.1 and item 1, $\omega : \mathbb{R} \to \mathbb{R}; \gamma \to \omega(\gamma)$ is pseudoconvex as long as $A(\gamma) > 0$. We prove that the later is true for γ belonging to an open interval in the real line. Indeed, let $A = QDQ^T$ be the spectral decomposition of A and define

$$w = D^{-1/2}Q^T u \quad \text{and} \quad W = ww^T = D^{-1/2}Q^T uu^T QD^{-1/2}. \tag{3.6}$$

Then we can rewrite

$$A(\gamma) = QD^{1/2}(I + \gamma W)D^{1/2}Q^T, \tag{3.7}$$

which is positive definite if and only if the rank one update of I, $I + \gamma W$, belongs to the cone of positive definite matrices. Now, note that the eigenvalues of this term are $\lambda_1 = 1$, with multiplicity $n - 1$, and $\lambda_2 = 1 + \gamma\|w\|^2$ with multiplicity 1. We then conclude that

$$A(\gamma) \in \mathbb{S}_{++}^n \iff \gamma \in \left]-\frac{1}{\|w\|^2}, +\infty\right[,$$

in which case $\lambda_2 > 0$. Moreover, $\omega(\gamma)$ tends to ∞ as γ approaches the extreme of the above interval. Therefore ω possesses a minimizer in the open interval, $\gamma^* \in \,]-\|w\|^{-2}, +\infty[$, that satisfies $\omega'(\gamma^*) = 0$. Note that since ω is pseudoconvex the fact that its derivative is equal to zero is also a sufficient condition for global optimality (see Fact 3.7 below).

In the following we obtain an explicit expression for the (unique) minimizer of (3.5), γ^*, by studying the zeros of ω'. Using the notation introduced in (3.6), f and its derivative are expressed as

$$f(\gamma) = \left(\mathrm{tr}(A) + \gamma\|u\|^2\right)/n \quad \text{and} \quad f'(\gamma) = \|u\|^2/n,$$

24

respectively. By making use of (3.7), g becomes

$$g(\gamma) := (\det(A) \det(I + \gamma W))^{1/n},$$

since $\det(D) = \det(A)$. As explained above the eigenvalues of $I + \gamma W$, are $\lambda_1 = 1 + \gamma \|w\|^2$, and the others are all 1, which yields that

$$g(\gamma) = \left(\det(A)(1 + \gamma \|w\|^2)\right)^{1/n} = \det(A)^{1/n}(1 + \gamma \|w\|^2)^{1/n}.$$

We get

$$g'(\gamma) = \frac{1}{n} \det(A)^{1/n} \|w\|^2 (1 + \gamma \|w\|^2)^{(1-n)/n}.$$

The derivative of ω is then obtained as follows

$$
\begin{aligned}
\omega'(\gamma) &= \frac{f'(\gamma)g(\gamma) - f(\gamma)g'(\gamma)}{g(\gamma)^2} \\
&= \frac{1}{g(\gamma)^2} \left[\frac{\|u\|^2}{n} \det(A)^{1/n} \left(1 + \gamma \|w\|^2\right)^{1/n} \right. \\
&\quad \left. - \frac{\|w\|^2}{n^2} \left(\text{tr}(A) + \gamma \|u\|^2\right) \det(A)^{1/n} \left(1 + \gamma \|w\|^2\right)^{(1-n)/n} \right] \\
&= \frac{\det(A)^{1/n}}{g(\gamma)^2 n^2} \left(1 + \gamma \|w\|^2\right)^{(1-n)/n} \left[n\|u\|^2 + (n-1)\gamma \|u\|^2 \|w\|^2 - \text{tr}(A)\|w\|^2 \right].
\end{aligned}
\tag{3.8}
$$

A simple computation shows that this derivative is 0 only when γ attains the value

$$\gamma^* = \frac{\text{tr}(A)\|w\|^2 - n\|u\|^2}{(n-1)\|u\|^2 \|w\|^2}, \tag{3.9}$$

that has to be in the open interval $]-\|w\|^{-2}, +\infty[$. Since ω is pseudoconvex, we conclude that γ^* is the ω-optimal conditioning that solves (3.5).

∎

Equivalently, we can deduce an expression for the ω-optimal conditioning by making use of the Cholesky decomposition of A instead of the spectral decomposition. This is gathered in our next corollary. The proof follows from the same calculations than Theorem 3.2 and thus is omitted.

Corollary 3.3. *Given A and U as in Theorem 3.2. Let $A = LL^T$ be the Cholesky decomposition of A. Then, the formula for the ω-optimal conditioning γ^* in (3.4) holds with the replacement*

$$w \leftarrow L^{-1}u.$$

25

As shown in Example 3.1, in some applications the preconditioner multiplier γ is required to take values in the interval $[0,1]$. In the following, we analize the optimal ω-preconditioner for the rank 1 update subject to this interval constraint.

Corollary 3.4. *Let the assumptions of Theorem 3.2 hold and let $\bar{\gamma}$ be the ω-optimal conditioning in the interval $[0,1]$, i.e.,*

$$\bar{\gamma} = \arg \min_{\substack{0 \leqslant \gamma \leqslant 1 \\ A(\gamma) > 0}} \omega(\gamma).$$

Then, if $\gamma^* \in \,]-\|w\|^2, +\infty[$ is the ω-optimal "unconstrained" conditioning obtained in Theorem 3.2, the following hold:

(i) If $\gamma^* \in [0,1] \implies \bar{\gamma} = \gamma^*$;

(ii) If $\gamma^* < 0 \implies \bar{\gamma} = 0$;

(iii) If $\gamma^* > 1 \implies \bar{\gamma} = 1$.

Proof. Item (i) In this case, since γ^* is the global optimum of ω in $\,]-\|w\|^2, +\infty[$, it would also be so in the interval $[0,1]$.

For Items (ii) and (iii), it suffices to observe that, by (3.8) and (3.9), when $\gamma^* < 0$ (respectively, $\gamma^* > 1$) the derivative of ω is monotonically increasing (respectively, decreasing) in the interval $[0,1]$.

■

3.3 Optimal Conditioning with a Low Rank Update

We now consider the case where the update is low rank. We need the following notations. For a matrix $Z \in \mathbb{R}^{n \times t}$, we use MATLAB notation and define the function $\mathrm{norms}(Z) : \mathbb{R}^{n \times t} \to \mathbb{R}^t$ as the (column) vector of column 2-norms of Z. We let $\mathrm{norms}^\alpha(Z)$ denote the vector of column norms with each norm to the power α.

Theorem 3.5 (Rank t-update). *Let $A \in \mathbb{S}_{++}^n$, $U = [u_1, \ldots, u_t] \in \mathbb{R}^{n \times t}$, be given with $n > t \geqslant 2$, and $\mathrm{norms}(U) > 0$. Set*

$$A(\gamma) = A + U \operatorname{Diag}(\gamma) U^T, \ \ for \ \gamma \in \mathbb{R}^t.$$

Let the spectral decomposition of A be given by $A = QDQ^T$, define $w_i = D^{-1/2} Q^T u_i, i \in [t]$, as in (3.3), with $W = [w_1 \ldots w_t]$. Let

$$
\begin{aligned}
K(U) &= \left[n \operatorname{Diag}\left(\mathrm{norms}^2(U)\right) - e \, \mathrm{norms}^2(U)^T \right], \\
b(U) &= \left(\mathrm{tr}(A)e - n \operatorname{Diag}\left(\mathrm{norms}^2(W)\right)^{-1} \mathrm{norms}^2(U) \right),
\end{aligned}
\tag{3.10}
$$

where e denotes the vector of all ones. Then, the ω-optimal conditioning,

$$\gamma^* = \underset{A(\gamma)>0}{\operatorname{argmin}} \, \omega(\gamma), \tag{3.11}$$

is given component-wise for $i \in [t]$ by

$$
\begin{aligned}
(\gamma^*)_i &= (K(U)^{-1}b(U))_i \\
&= \frac{1}{(n-t)\|u_i\|^2} \left(\operatorname{tr}(A) - (n-t)\frac{\|u_i\|^2}{\|w_i\|^2} - \sum_{j=1}^{t} \frac{\|u_j\|^2}{\|w_j\|^2} \right).
\end{aligned} \tag{3.12}
$$

Proof. Let $A > 0$ and

$$U = \begin{bmatrix} u_1 & \dots & u_t \end{bmatrix} \in \mathbb{R}^{n \times t}, \quad \text{with } n > t \geqslant 2.$$

We consider the update of the form

$$A(\gamma) = A + U \operatorname{Diag}(\gamma)U^T = A + \sum_{i=1}^{t} \gamma_i u_i u_i^t, \quad \gamma \in \mathbb{R}^t.$$

Same than in Theorem 3.2, we start characterizing an open subset of $\mathbb{R}^t$ where $A(\gamma)$ is positive definite. In order to do this, we again transform the problem using the spectral decomposition of A, $A = QDQ^T$, and setting

$$w_i = D^{-1/2}Q^T u_i \quad \text{and} \quad W_i = w_i w_i^T \quad \text{for } i \in [t].$$

Then, we can express $A(\gamma)$ as

$$
\begin{aligned}
A(\gamma) &= A + U \operatorname{Diag}(\gamma)U^T \\
&= QD^{1/2} \left(I + D^{-1/2}Q^T U \operatorname{Diag}(\gamma)U^T QD^{-1/2} \right) D^{1/2}Q^T \\
&= QD^{1/2} \left(I + \sum_{i=1}^{t} \gamma_i (D^{-1/2}Q^T u_i)(u_i^T QD^{-1/2}) \right) D^{1/2}Q^T \\
&= QD^{1/2} \left(I + \sum_{i=1}^{t} \gamma_i W_i \right) D^{1/2}Q^T.
\end{aligned}
$$

By repeatedly making use of the formula for the determinant of the sum of an invertible matrix and a rank one matrix (see, e.g., [29, Example 4]), we obtain the following expression for the determinant of $A(\gamma)$

$$\det\left(A(\gamma)\right) = \det(A) \prod_{i=1}^{t} (1 + \gamma_i \|w_i\|^2). \tag{3.13}$$

Consequently, $A(\gamma)$ is nonsingular and, by continuity of the eigenvalues, positive definite for γ belonging to the set

$$\Omega := \left] -\frac{1}{\|w_1\|^2}, +\infty \right[\times \left] -\frac{1}{\|w_2\|^2}, +\infty \right[\times \dots \times \left] -\frac{1}{\|w_t\|^2}, +\infty \right[. \tag{3.14}$$

27

Now, note that the constraint $A(\gamma) > 0$ is a positive definite constraint, so it is convex. Therefore, if there exists some γ outside of Ω such that $A(\gamma) > 0$, we would lose the convexity of the feasible set, since $A(\gamma)$ is singular on the boundary of Ω. This implies that

$$A(\gamma) > 0 \iff \gamma \in \Omega.$$

Moreover, since $\omega(\gamma) \to +\infty$ as γ tends to the border of Ω or to $+\infty$, we can ensure that γ has a minimizer in Ω. Since the function is pseudoconvex on this open set, the global minimum is attained at a point γ^* such that $\nabla \omega(\gamma^*) = 0$. Next, we prove that γ^* is given by (3.12).

For this, note that $f(\gamma)$ can be expressed as

$$f(\gamma) = \frac{1}{n} \operatorname{tr}(A + U \operatorname{Diag}(\gamma) U^T) = \frac{1}{n} \left(\operatorname{tr}(A) + \sum_{i=1}^{t} \gamma_i \|u_i\|^2 \right) = \frac{1}{n} \left(\operatorname{tr}(A) + \gamma^T \operatorname{norms}^2(U) \right),$$

and its gradient is $\nabla f(\gamma) = \frac{1}{n} \operatorname{norms}^2(U)$. On the other hand, by (3.13) $g(\gamma)$ can be expressed as

$$g(\gamma) = \det(A)^{1/n} \left(\prod_{i=1}^{t} (1 + \gamma_i \|w_i\|^2) \right)^{1/n}.$$

The gradient of $g(\gamma)$ is then given component-wise by

$$
\begin{aligned}
\frac{\partial g(\gamma)}{\partial \gamma_j} &= \frac{1}{n} \det(A)^{1/n} \left(\prod_{i=1}^{t} (1 + \gamma_i \|w_i\|^2) \right)^{(1-n)/n} \left(\prod_{i=1, i \neq j}^{t} (1 + \gamma_i \|w_i\|^2) \right) \|w_j\|^2 \\
&= \frac{1}{n} \det(A)^{1/n} \left(\prod_{i=1}^{t} (1 + \gamma_i \|w_i\|^2) \right)^{1/n} (1 + \gamma_j \|w_j\|^2)^{-1} \|w_j\|^2 \\
&= \frac{g(\gamma)}{n} \frac{\|w_j\|^2}{1 + \gamma_j \|w_j\|^2}, \qquad\qquad j \in [t].
\end{aligned}
$$

We make use of these expressions in order to compute the partial derivatives of ω. For every $j \in [t]$, we have

$$
\begin{aligned}
\frac{\partial \omega(\gamma)}{\partial \gamma_j} &= \frac{1}{g(\gamma)^2} \left[\frac{\partial f(\gamma)}{\partial \gamma_j} g(\gamma) - f(\gamma) \frac{\partial g(\gamma)}{\partial \gamma_j} \right] \\
&= \frac{1}{n^2 g(\gamma)} \left[n \|u_j\|^2 - \frac{\|w_j\|^2 \left(\operatorname{tr}(A) + \gamma^T \operatorname{norms}^2(U) \right)}{1 + \gamma_j \|w_j\|^2} \right].
\end{aligned}
\tag{3.15}
$$

Since $n^2 g(\gamma) > 0$, the j-th partial derivative of ω is zero if, and only if,

$$n \|u_j\|^2 - \frac{\|w_j\|^2 \left(\operatorname{tr}(A) + \gamma^T \operatorname{norms}^2(U) \right)}{1 + \gamma_j \|w_j\|^2} = 0.$$

Therefore, the minimum of the pseudoconvex function is obtained as the solution of the linear system defined by the t equations

$$(n-1)\|u_k\|^2\gamma_k - \sum_{i=1,i\neq k}^{t} \|u_i\|^2\gamma_i = \text{tr}(A) - n\frac{\|u_k\|^2}{\|w_k\|^2}, \quad k \in [t].$$

Equivalently,

$$\begin{bmatrix} (n-1)\|u_1\|^2 & -\|u_2\|^2 & \ldots & & -\|u_t\|^2 \\ -\|u_1\|^2 & (n-1)\|u_2\|^2 & -\|u_3\|^2 & \ldots & -\|u_t\|^2 \\ \ldots & & & & \\ -\|u_1\|^2 & \ldots & \ldots & -\|u_{t-1}\|^2 & (n-1)\|u_t\|^2 \end{bmatrix} \gamma = \begin{pmatrix} \text{tr}(A) - n\|u_1\|^2/\|w_1\|^2 \\ \ldots \\ \text{tr}(A) - n\|u_t\|^2/\|w_t\|^2 \end{pmatrix}.$$

This is further equivalent to

$$\left[n \,\text{Diag}(\text{norms}^2(U)) - e\,\text{norms}^2(U)^T\right]\gamma = \left(\text{tr}(A)e - n\,\text{Diag}\left(\text{norms}^2(W)\right)^{-1}\text{norms}^2(U)\right),$$

which is the system $K(U)\gamma = b(U)$ using the notation in (3.10).

Now we derive an explicit expression for the optimal γ. In order to do this, note that $K(U)$ is given as the sum of an invertible matrix, $n\,\text{Diag}(\text{norms}^2(U))$, and an outer product of vectors, $-e\,\text{norms}^2(U)^T$. By the *Sherman-Morrison formula*, this sum is invertible if and only if

$$1 - \frac{1}{n}\text{norms}^2(U)^T\,\text{Diag}(\text{norms}^2(U))^{-1}e \neq 0.$$

This is always true for $t < n$. Indeed, we have

$$1 - \frac{1}{n}\text{norms}^2(U)^T\,\text{Diag}(\text{norms}^2(U))^{-1}e = 1 - \frac{1}{n}e^T e = 1 - \frac{t}{n} > 0.$$

Moreover, we obtain the following expression for the inverse

$$\left(n\,\text{Diag}(\text{norms}^2(U)) - e\,\text{norms}^2(U)^T\right)^{-1}$$
$$= \frac{1}{n}\text{Diag}(\text{norms}^2(U))^{-1} + \frac{1}{\left(1-\frac{t}{n}\right)n^2}\text{Diag}(\text{norms}^2(U))^{-1}e\,\text{norms}^2(U)^T\,\text{Diag}(\text{norms}^2(U))^{-1}$$
$$= \frac{1}{n}\text{Diag}(\text{norms}^{-2}(U)) + \frac{1}{(n-t)n}\text{Diag}(\text{norms}^{-2}(U))ee^T.$$

Therefore, the inverse of $K(U)$ in matrix form is given by

$$K(U)^{-1} = \frac{1}{n}\begin{bmatrix} \frac{1}{\|u_1\|^2} & 0 & \ldots & 0 \\ 0 & \frac{1}{\|u_2\|^2} & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & \frac{1}{\|u_t\|^2} \end{bmatrix} + \frac{1}{(n-t)n}\begin{bmatrix} \frac{1}{\|u_1\|^2} & \frac{1}{\|u_1\|^2} & \cdots & \frac{1}{\|u_1\|^2} \\ \frac{1}{\|u_2\|^2} & \frac{1}{\|u_2\|^2} & \cdots & \frac{1}{\|u_2\|^2} \\ \vdots & \vdots & \vdots & \vdots \\ \frac{1}{\|u_t\|^2} & \frac{1}{\|u_t\|^2} & \cdots & \frac{1}{\|u_t\|^2} \end{bmatrix}.$$

Finally, we obtain γ^* by calculating the product $\gamma^* = K(U)^{-1}b(U)$ which yields

$$\gamma_i^* = \frac{1}{(n-t)\|u_i\|^2}\left(\text{tr}(A) - (n-t)\frac{\|u_i\|^2}{\|w_i\|^2} - \sum_{j=1}^{t}\frac{\|u_j\|^2}{\|w_j\|^2}\right), \tag{3.16}$$

for all $i \in [t]$. Since γ^* is the unique zero of the gradient of ω, we conclude that it belongs to Ω and solves (3.11).

■

We note that the ω-optimal conditioning for the rank one update in Theorem 3.2 is obtained from (3.12) when $t = 1$. On the other hand, we can also employ the Cholesky decomposition of A to derive the ω-optimal conditioning in Theorem 3.5. We state this in the following corollary.

Corollary 3.6. *Given A and U as in Theorem 3.5. Let $A = LL^T$ be the Cholesky decomposition of A. Then, the formula for the ω-optimal conditioning γ^* in (3.12) holds with the replacement*

$$w_i \leftarrow L^{-1}u_i, \quad i \in [t].$$

Proof. The proof follows similarly to the one of Theorem 3.5 and thus is omitted. ■

With the same assumptions as in Theorem 3.5, we now consider the problem of finding the ω-optimal conditioning in the box $[0,1]^t$, i.e.,

$$\bar{\gamma} = \arg\min_{\substack{\gamma \in [0,1]^t \\ A(\gamma) \succ 0}} \omega(\gamma). \tag{3.17}$$

For the rank one update $(t = 1)$, Corollary 3.4 shows that the solution to (3.17) can be obtained by first computing the minimum of the unconstrained problem, whose explicit expression was given in Theorem 3.2, and then projecting onto the box constraint, which in that case was the interval $[0,1]$. However, this simple projection can fail in general for the low rank update, as we now show in Example 3.8 below.

The illustration of this phenomenon will require considering a constrained pseudoconvex minimization problem. In the following Fact 3.7, see, e.g., [28, Chapter 10], we recall the sufficient optimality conditions for this class of optimization problems. We note that no constraint qualification is needed for *sufficiency.*

Fact 3.7 (Sufficient optimality conditions for pseudoconvex programming). *Let $\Omega \subseteq \mathbb{R}^n$ be nonempty open and convex. Let $f : \Omega \to \mathbb{R}$ be a pseudoconvex function and $(g_i)_{i=1}^m : \Omega \to \mathbb{R}$ a family of differentiable and quasiconvex functions. Consider the optimization problem*

$$\begin{array}{ll} \min & f(x) \\ \text{s.t.} & g_i(x) \leqslant 0, \quad i \in [m], \\ & x \in \Omega. \end{array} \tag{3.18}$$

Let $\bar{x} \in \Omega, \bar{\lambda} \in \mathbb{R}^m$, be a KKT primal-dual pair, i.e., the following KKT conditions hold:

$$\nabla f(\bar{x}) + \sum_{i=1}^{m} \bar{\lambda}_i \nabla g_i(\bar{x}) = 0$$

$$\bar{\lambda}_i \geqslant 0, \quad i \in [m], \tag{3.19}$$
$$\bar{\lambda}_i g_i(\bar{x}) = 0, \quad i \in [m],$$
$$\bar{x} \in \Omega \text{ and } g_i(\bar{x}) \leqslant 0, \quad i \in [m].$$

Then $\bar{x}$ solves (3.18).

Example 3.8 (Failure of projection for constrained problem (3.17))**.** *Let $n = 3$, $t = 2$ and consider the following initial data for the ω-minimization problem:*

$$A := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{bmatrix} \quad and \quad U := \begin{bmatrix} \frac{1}{\sqrt{2}} & 0 \\ \frac{-1}{\sqrt{2}} & 0 \\ 0 & 1 \end{bmatrix}.$$

Then, we get the following:

- *From (3.16) and Theorem 3.5, the ω-optimal preconditioner is $\gamma^* = \frac{1}{3}\begin{pmatrix} 1 \\ -1 \end{pmatrix}$;*

- *projecting onto $[0,1]^2$ yields $\gamma_p^* = \frac{1}{3}\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, where $\omega(\gamma_p^*) = 16/(9\sqrt[3]{5})$;*

- *however, with $\bar{\gamma} := \frac{1}{2}\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, we get a lower value:*

$$\omega(\bar{\gamma}) = 1/\left(3\sqrt[3]{(2/11)^2}\right) \approx 1.0386 < 1.0397 \approx 16/(9\sqrt[3]{5});$$

  and $\bar{\gamma}$ is the ω-optimal preconditioner in $[0,1]^2$, as we now show.

To prove the last statement, note that (3.17) can be written as the pseudoconvex program in (3.18) by setting $f := \omega : \mathbb{R}^2 \to \mathbb{R}$, $g_1(\gamma) = -\gamma_1$, $g_2(\gamma) = -\gamma_2$, $g_3(\gamma) = \gamma_1 - 1$, $g_4(\gamma) = \gamma_2 - 1$ and Ω defined as in (3.14). In particular, the only active constraint for $\bar{\gamma} = (1/2, 0)^T$ is $g_2(\gamma) = 0$, so the KKT conditions become

$$0 = \frac{\partial \omega(\bar{\gamma})}{\partial \gamma_1},$$
$$0 = \frac{\partial \omega(\bar{\gamma})}{\partial \gamma_2} - \bar{\lambda}_2,$$

for some $\bar{\lambda}_2 \geqslant 0$. This can be verified by simply substituting using the expressions of the partial derivatives of ω obtained in (3.15). By Fact 3.7, we conclude that for the given data, $\bar{\gamma}$ is the solution of (3.17).

As done in the previous example, obtaining the ω-optimal preconditioner in the box $[0,1]^t$ would require obtaining a KKT point for the constrained pseudoconvex problem (3.17). This is not an easy task. To the author's knowledge, closed formulas for this kind of box constrained minimization problems are not known even when the objective is a quadratic. Nevertheless, using the projection of γ^* onto $[0,1]^t$ as an approximation to $\bar{\gamma}$ appears to give good results in practice. We see this in our numerical tests in Section 4.

Finally, observe that the computation of γ^* in formula (3.12) might be as expensive as finding the Newton direction without preconditioning, since it requires the spectral or Cholesky decomposition. However, under the framework of Example 3.1, we now see in Proposition 3.9 that we can get the following inexpensive and effective approximation $\gamma = \gamma^*_{\mathrm{apr}}$, given by the expression

$$\left[\gamma^*_{\mathrm{apr}}\right]_i := \frac{\mathrm{tr}(A)}{(n-t)\|u_i\|^2}, \quad \forall i \in [t]. \tag{3.20}$$

Proposition 3.9. *Let $A \in \mathbb{S}^n_+, B \in \mathbb{R}^{n \times m}, U \in \mathbb{R}^{n \times t}$, with $t := |\mathcal{I}_0|$, be defined as in Example 3.1. For simplicity define the columns of the matrix $\bar{B} := \left[B_i\right]_{i \in \mathcal{I}_+}$ with $r := \mathrm{rank}(\bar{B}) < n$. Let $j \in \mathcal{I}_0$ and let u_j, w_j be defined as in Theorem 3.5, and let $\epsilon > 0$ be the Levenberg–Marquardt regularization parameter in (3.2). Then,*

$$u_j \notin \mathrm{range}(\bar{B}) \implies \frac{\|u_j\|^2}{\|w_j\|^2} \leqslant \epsilon \frac{\|u_j\|^2}{\mathrm{dist}(u_j, \mathrm{range}(\bar{B}))^2}.$$

We conclude that (3.20) *provides an efficient estimate of γ^* in formula* (3.12).

Proof. Let $x \in \mathrm{range}(\bar{B})$, $y \in \mathrm{null}(\bar{B}^T)$ such that $u_j = x + y$. Then,

$$\begin{aligned}
\|w_j\|^2 &= u_j^T Q D^{-1} Q^T u_j \\
&= (x+y)^T Q D^{-1} Q^T (x+y) \\
&= x^T Q D^{-1} Q^T x + 2x^T Q D^{-1} Q^T y + y^T Q D^{-1} Q^T y \\
&= \sum_{\ell=1}^r \frac{1}{\lambda_\ell + \epsilon}(q_\ell^T x)^2 + 0 + \sum_{\ell=r+1}^n \frac{1}{\epsilon}(q_\ell^T y)^2 \\
&\geqslant \sum_{\ell=r+1}^n \frac{1}{\epsilon}(q_\ell^T y)^2 \\
&= \frac{1}{\epsilon} \sum_{\ell=r+1}^n (q_\ell^T y)^2 = \frac{1}{\epsilon} \sum_{\ell=1}^n (q_\ell^T y)^2 = \frac{1}{\epsilon}\|y\|^2,
\end{aligned}$$

where q_ℓ is ℓ-th column of Q. Since $\|y\| = \mathrm{dist}(u_j, \mathrm{range}(\bar{B}))$,

$$\frac{\|u_j\|^2}{\|w_j\|^2} \leqslant \epsilon \frac{\|u_j\|^2}{\mathrm{dist}(u_j, \mathrm{range}(\bar{B}))^2}.$$

■

Remark 3.10. *We note that the approximate ω-optimal conditioning obtained in* (3.20) *is strictly related to a popular choice of γ appearing in the literature (see, e.g., [3]), namely, taking*

$$\gamma_i := \frac{1}{\|u_i\|^2}, \quad \forall i \in [t]. \tag{3.21}$$

Indeed, the updates resulting from (3.20) *and* (3.21) *only differ in a scaling given by* $\mathrm{tr}(A)/(n-t)$. *Recall from Proposition 2.1 Item 2 that the selection of* (3.21) *corresponds to the ω-optimal diagonal preconditioner of the matrix U, i.e., it aims to minimize the ω-condition number of only the update term. In contrast, our proposed approach aims to minimize ω for the whole matrix $A(\gamma) = A + U\,\mathrm{Diag}(\gamma)U^T$. Numerical comparisons between both approaches are presented in the numerics in Table 4.1 and Figure 4.2, below.*

4 Numerical Tests

We now present empirics for: the various preconditioners Section 4.1; and the optimal preconditioned low rank updates Section 4.2. The experiments were done on: Intel Core i7-12700H 2.30 GHz with 16GB RAM, under Windows 11 (64-bit). We used MATLAB version 2024a. The MATLAB source code and data of all the experiments is available at https://github.com/DavidTBelen/omega-condition-number.

4.1 Comparisons of Preconditioners; Positive Definite Systems

In this section, we analyze the performance of an iterative method for approximately solving positive definite linear systems subject to different preconditioning strategies. Specifically, we compare the ω-optimal diagonal and incomplete upper triangular ω-optimal preconditioners introduced above with state-of-the-art preconditioners, e.g., the incomplete Cholesky preconditioner. Our test enviroment follows the line of the extensive numerical comparisons presented in the survey [18].

4.1.1 Test Enviroment

The problems used in our experiment are all constructed with data from the SuiteSparse Matrix Collection [27]. We consider the symmetric positive definite matrices in this repository whose number of rows (columns) range from $5{,}000$ to $30{,}000$; but without "duplicates" (i.e., without similar matrices belonging to the same group). The right hand side of our linear system $b = e$, is always set as the vector of all ones.

As the iterative method for solving the positive definite linear systems, we consider the implementation of the *Preconditioned Conjugate Gradient Method* given by MATLAB's built-in function **pcg**. This is MATLAB's benchmark iterative solver for positive definite linear

systems. In all our experiments, our stopping criterion for **pcg** is when the relative residual reaches a tolerance smaller than 10^{-6}, i.e.,

$$\frac{\|Wx - b\|}{\|b\|} < 10^{-6}.$$

Finally, in order to avoid "trivialities", we discard matrices that generate problems that can be solved to the desired tolerance in less than 10 seconds by **pcg** with *no preconditioner*. This leaves a subset of 16 matrices whose specific characteristics are detailed in Table A.1. In the following we use P to denote the set of these 16 problems.

4.1.2 Preconditioning Strategies

We use the following strategies (with acronyms):

- **No preconditioning** (NONE).

- The ω-**optimal diagonal preconditioner** (DIAG) given by (2.1).

- The ω-**optimal incomplete upper triangular preconditioner** (ITRIU) given by (2.12). The dimension k of the triangular block is chosen according to the nonzero entries $nnz(W)$ of the matrix of interest W as

$$k = \left\lceil \frac{1}{2}\left(1 + \sqrt{1 + \frac{4}{5}nnz(W)}\right)\right\rceil + 1,$$

  where $\lceil \cdot \rceil$ is *ceiling*. The motivation on this choice resides in obtaining a preconditioner with fewer nonzero entries than in W, i.e., $t(k-1) << nnz(W)$. The last summand 1 ensures that the preconditioner is not diagonal.

- **Incomplete Cholesky factorization** (ICHOL). This preconditioning strategy consists in considering a Cholesky factorization of W, given by LL^T, but where some of the entries of L are ignored agreeing with the sparsity pattern of W. The preconditioned system then becomes
$$L^{-1}WL^{-T}y = L^{-1}b, \quad y = L^Tx.$$

  We use MATLAB's **ichol** to construct L and use the options of the **pcg** solver for solving the preconditioned system without constructing L^{-1} explicitly, as that could lead to the loss of sparsity. To ensure that the process does not break down (which can happen if a non positive pivot is encountered) we shift W and obtain an approximation of $W + \alpha \operatorname{Diag}(\operatorname{diag} W)$. We make use of two different choices for the scaling factor α. The first of them, denoted below as ICHOL(1) , uses the recommended tuning in the MATLAB Help Center. However, as can be observed in our tests below, this leads to a larger number of iterations of **pcg** than what is expected from an incomple Cholesky preconditioning. For instance, ICHOL(1) does not improve ITRIU in this aspect.

Hence, we include a second choice of α taken two orders of magnitude smaller than the recommended parameter in MATLAB Help Center. This significantly reduces the number of iterations required by **pcg**. In the experiments below, ICHOL(2) refers to this latter choice of α.

4.1.3 Performance Profile

Besides illustrating the output from the experiments as displayed in Tables A.1 to A.4, we also employ performance profile plots, e.g., [13]. These plots are constructed as follows. Let $\Gamma := \{\text{NONE, DIAG, ITRIU, ICHOL(1), ICHOL(2)}\}$ be the set of preconditioners for our comparisons. For each $p \in P$ and $\gamma \in \Gamma$, we denote as $t_{p,\gamma}$ the measure we want to compare. In particular, we will separately consider the number of iterations and the time required for solving the system (to the desired tolerance) for the preconditioned linear system described in Section 4.1.1. In the cases where we consider a preconditioned system (i.e., all except NONE), the time for computing the preconditioner is also included in $t_{p,\gamma}$, i.e.,

$$t_{p,\gamma} = \{\text{time for computing the preconditioner}\}$$
$$+ \{\text{time for solving the preconditioned problem by } \textbf{pcg}\}.$$

Then, for every problem $p \in P$ and every $\gamma \in \Gamma$, we define the performance ratio as

$$r_{p,\gamma} := \left\{ \begin{array}{ll} \frac{t_{p,\gamma}}{\min\{t_{p,\gamma} \, : \, \gamma \in \Gamma\}} & \text{if convergence test passed,} \\ +\infty & \text{if convergence test failed.} \end{array} \right.$$

In our experiments, a convergence test *passed* if it succeeded in solving the linear system with the required relative residual tolerance in less than $100,000$ iterations, and otherwise it *failed*. Note that the best performing preconditioner with respect to the measure under study (time or number of iterations), say $\tilde{\gamma}$, for problem p will have performance ratio $r_{p,\tilde{\gamma}} = 1$. In contrast, if the preconditioner γ underperforms in comparison with $\tilde{\gamma}$, but still manages to pass the test, then

$$r_{p,\gamma} = \frac{t_{p,\gamma}}{t_{p,\tilde{\gamma}}} > 1$$

is the ratio between the overall time (resp., number of iterations) required for solving the problem p for this particular choice and the time (resp., number of iterations) employed by $\tilde{\gamma}$. Consequently, the larger the value of $r_{p,\gamma}$, the worse the preconditioner γ performed for problem p.

Finally, the performance profile of $\gamma \in \Gamma$ is defined as

$$\rho_{\gamma}(\tau) := \frac{1}{|P|} \, \text{size} \, \{p \in P \, : \, r_{p,\gamma} \leqslant \tau\},$$

where $|P|$ is the number of problems in P. This can be understood as the relative portion of times that the performance ratio $r_{p,\gamma}$ is within a factor of $\tau \geqslant 1$ of the best possible performance ratio. In particular, $\rho_{\gamma}(1)$ represents the number of problems where γ is the best choice. Also, the existence of a $\tau \geqslant 1$ such that $\rho_{\gamma}(\tau) = 1$, indicates that γ passed the convergence test for every single problem in P. In Figure 4.1, we display our performance profiles, with $\log_2$ scale on τ.
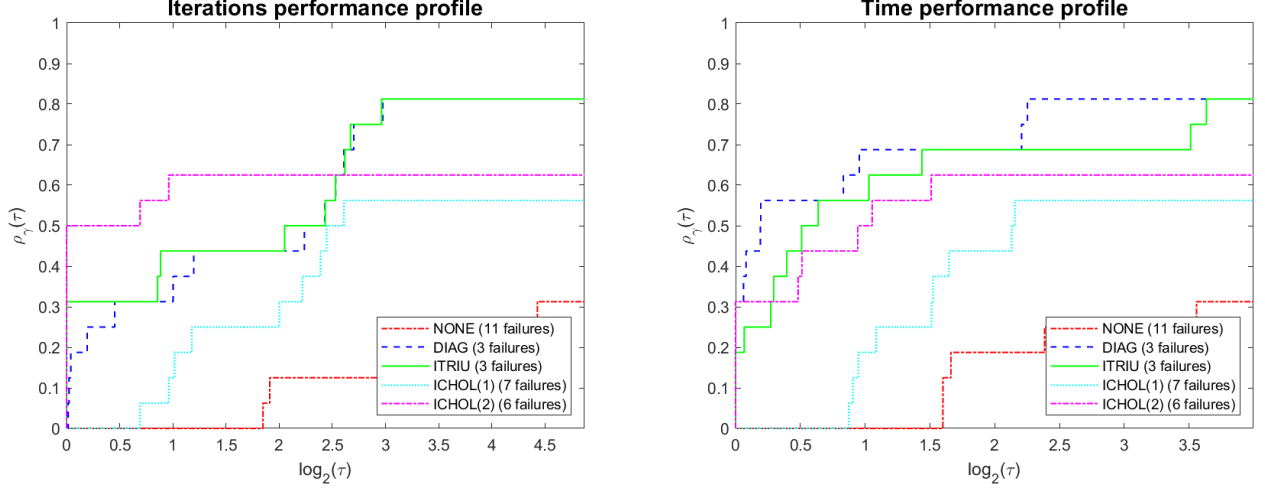
Figure 4.1: Iterations and time performance profiles for solving the system with the different choices of preconditioner.

4.1.4 Summary of the Empirics

Our empirics suggest that the diagonal (DIAG) and the incomplete upper diagonal ω-optimal (ITRIU) preconditioners have similar behaviour. More precisely, ITRIU seems to reduce the number of iterations required by **pcg** in comparison to DIAG (see Table A.1). This can be understood as a benefit of the additional reduction of the ω-condition number furnished by the incomplete upper triangular block. In general, the incomplete Cholesky (ICHOL)(2) appears to be the best solver for reducing the number of iterations, however it also fails to reach the desired relative residual accuracy more often (6 times) than the ω-optimal precon-ditioners (3 times). The residuals obtained by each one of the methods can be checked in Table A.3. Although the numerical results suggest that an appropriate incomplete Cholesky factorization provides a superior preconditioning, we note that the ω-optimal preconditioners are more *stable*, as the performance of the incomplete Cholesky is usually influenced by the choice of the scaling factor α. Indeed, a better result of **pcg** is obtained with the choice of a small scaling factor α in ICHOL, but if α is too small a non positive pivot can be encountered and ICHOL may fail to be computable. In contrast, the ω-preconditioners can always be computed without the need to manually set additional parameters.

4.2 ω-Optimal Low Rank Updates for Generalized Jacobians

We now present tests with different choices of γ for efficient iterative solutions of linear systems of the form $\boxed{A(\gamma)x = b}$, where $A(\gamma)$ is given in (4.1). We use MATLAB's builtin preconditioned conjugate gradient function **pcg**. We focus our attention on the case where $A(\gamma) \in \mathbb{S}_{++}^n$ is a low rank update that appears in choosing subgradients in nonsmooth Newton methods, see Example 3.1. Our aim is to improve conditioning to improve convergence, thus we call this γ-*conditioning*.

4.2.1 Problem Generation and Definitions of γ

Specifically, we generate random instances as follows:

- Define

$$A(\gamma) := A + \epsilon I + U \operatorname{Diag}(\gamma) U^T; \tag{4.1}$$

- ϵ is a random number in the interval $[10^{-9}, 10^{-7}]$;

- $A = A_0^T A_0$ with $A_0 \in \mathbb{R}^{r \times n}$ a normally distributed random sparse matrix with density at most $0.5/\log(n)$; $r \in [n/2 + 1, n - 1]$ is a random integer;

- $t \in [2, r/2]$ is a randomly chosen integer determining the rank of the update matrix U; $U \in \mathbb{R}^{n \times t}$ is a normally distributed random sparse matrix of density at most $1/\log(n)$;

- the right hand side, b, is chosen as the sum of two random vectors in the range of A and U, respectively. More precisely,

$$b = A b^1 + U b^2,$$

with $b^1 \in \mathbb{R}^n$ and $b^2 \in \mathbb{R}^t$ vectors randomly generated using the standard normal distribution.

As explained in Example 3.1, in this application the γ for conditioning is required to belong to the hypercube $[0, 1]^t$. Therefore, in our experiments we test the performance of five different choices of γ-conditioning:

- ($\gamma = 0$): the zero vector;

- ($\gamma = e$): the vector of ones;

- ($\gamma = u^{-2}$): projection onto $[0, 1]^t$ of the ω-optimal diagonal preconditioner for the last term, $U \operatorname{Diag}(\gamma) U^T$, of (4.1). Recall from Proposition 2.1, Item 2 that γ is given by

$$\gamma_i = \min\{1, 1/\|u_i\|^2\}, \; i \in [t],$$

where u_i denotes the ith column of U;

- ($\gamma = \gamma_p^*$) projection of γ^*, obtained in Theorem 3.5, onto $[0, 1]^t$;

- ($\gamma = \gamma_{\text{apr}}^*$): projection of the approximated ω-optimal obtained in (3.20), onto $[0, 1]^t$.

4.2.2 Descriptions of Parameters and Outputs

For each dimension $n \in \{1000, 2000, 3000, 5000\}$, we generate 10 instances of random problems and solve the corresponding systems with MATLAB's **pcg** and with the five different choices of γ-conditioning.

Table 4.1 shows the average over the 10 instances of: κ- and ω-condition numbers of every $A(\gamma)$; relative residual; number of iterations; time used by **pcg** for every choice of γ; and time for computing each (nontrivial) γ. Both the spectral and Cholesky decompositions are implemented for computing $\gamma = \gamma_p^*$ but the table only contains the time for the most efficient approach. We indicated the corresponding approach in the last column as well. We stop if a tolerance of 10^{-12} is reached or the maximum $50,000$ iterations is exceeded. We use the origin as our initial starting point.

n	γ	$\kappa(A(\gamma))$	$\omega(A(\gamma))$	Rel.Res	Iter	T. Total	T. Solve	T. γ^*
1000	0	2.3173e+09	1.0173e+04	1.5747e-01	2688.90	0.0956	0.0956	-
	e	2.3445e+11	4.7689e+03	1.5987e-06	4210.70	0.1182	0.1182	-
	u^{-2}	8.5539e+09	2.2114e+03	3.8172e-07	2730.10	0.0677	0.0675	0.0003
	γ_p^*	1.0722e+10	2.2095e+03	1.8169e-07	2788.10	0.0825	0.0691	0.0134 (C)
	γ_{apr}^*	1.0432e+10	2.2095e+03	1.5855e-07	2821.80	0.0749	0.0745	0.0004
2000	0	2.0653e+12	2.2127e+04	1.5144e-07	231.00	0.0554	0.0554	-
	e	3.3680e+12	6.5292e+02	9.2932e-13	425.30	0.0954	0.0954	-
	u^{-2}	1.8491e+12	8.9700e+02	9.1536e-13	1364.40	0.2972	0.2969	0.0003
	γ_p^*	2.0726e+12	5.6538e+02	9.1902e-13	376.20	0.3969	0.0841	0.3128 (S)
	γ_{apr}^*	2.0644e+12	5.6538e+02	9.1737e-13	376.30	0.0862	0.0843	0.0019
3000	0	4.4976e+12	1.9718e+04	5.8425e-08	587.10	0.5799	0.5799	-
	e	3.6326e+12	3.9795e+03	9.4326e-13	261.60	0.2657	0.2657	-
	u^{-2}	3.6699e+12	4.3747e+03	9.1465e-13	917.80	0.9040	0.9039	0.0001
	γ_p^*	3.6544e+12	3.9795e+03	9.4326e-13	261.60	1.4322	0.2594	1.1728 (S)
	γ_{apr}^*	3.5688e+12	3.9795e+03	9.4326e-13	261.60	0.2666	0.2616	0.0050
5000	0	9.9930e+12	3.0421e+04	2.6244e-07	698.60	2.3558	2.3558	-
	e	1.1796e+13	8.9242e+02	9.3999e-13	362.90	1.2255	1.2255	-
	u^{-2}	8.2778e+12	1.4249e+03	1.2583e-09	1563.00	5.3365	5.3344	0.0022
	γ_p^*	8.7440e+12	8.2755e+02	9.6946e-13	344.30	6.9389	1.1789	5.7600 (S)
	γ_{apr}^*	8.8089e+12	8.2755e+02	9.5456e-13	344.30	1.1824	1.1669	0.0155

Table 4.1: For different dimensions n, every choice of γ for updating, average of 10 instances: κ- and ω-condition numbers of $A(\gamma)$; residual; number of iterations; total time (in seconds); solve time; time for computing γ^*, (S) stands for spectral and (C) for Cholesky decomposition.

We also use performance profiles to compare the different choices of γ; details in Section 4.1.3. Again, let P denote the set of problems, and now set $\Gamma := \{0, e, u^{-2}, \gamma_p^*, \gamma_{apr}^*\}$ as the set of γ-conditioners. We separately consider the number of iterations and the time required for solving the system $A(\gamma)\,x = b$. We set the time

$$t_{p,\gamma_p^*} = \{\text{time for solving the system } A(\gamma)\,x = b\} + \{\text{time for computing } \gamma_p^*\}.$$

The latter quantity is taken as the minimum between the spectral and Cholesky approach. For constructing the performance ratio in this setting, we consider that a convergence test passed, rather than failed, if it succeeded in solving the linear system with the required tolerance in less than $50,000$ iterations. The output appears in Figure 4.2.
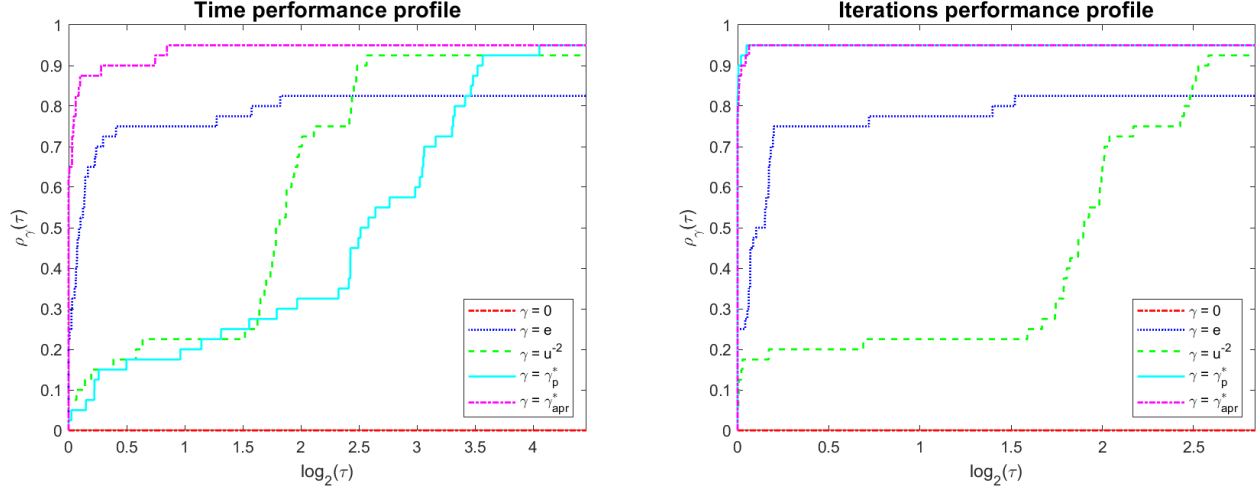

Figure 4.2: Time, iterations performance profiles; for system $A(\gamma)x = b$ with different choices of γ in Section 4.2.1; using MATLAB's **pcg**.

4.2.3 Summary of Empirics

Firstly, we observe that **pcg** with no γ-conditioning ($\gamma=0$) fails to achieve the desired residual, so the problems under consideration are sufficiently ill-conditioned. The performance profiles reveal that, in more than 90% of the tested instances, the ω-optimal conditioning leads to a problem that can be solved with the least number of iterations. However, the computation of the optimal conditioning γ_p^* is too time expensive (see Table 4.1) which does not make it advantageous in general[10]. Nonetheless, we observe that the approximated ω-optimal conditioning γ_{apr}^* maintains the benefits of γ_p^* in terms of both solve time and number of iterations. In addition, it can be computed very efficiently which makes it the best option among all γ-conditionings.

5   Conclusion

In this paper we have studied ω, a nonclassical matrix condition number formed as the ratio of the arithmetic and geometric means of eigenvalues. We have shown that ω has many

---

[10]Regarding the different approaches for computing $\gamma = \gamma_p^*$, we want to mention that although obtaining the Cholesky decomposition $A = LL^T$ is in general less costly than computing its eigenvalue decomposition, the computation of the ω-optimal conditioning in this case requires solving the system $LW = U$, see Corollary 3.6. This means that, for larger dimensions, employing the spectral decomposition for computing γ^* seems to be more time efficient.

advantages over κ, the classic condition number formed as the ratio of the largest to smallest eigenvalues, as the latter is more of a *worst case* condition number. Moreover, the fact that $\kappa(A) = \kappa(A^{-1})$, (not true for ω) is misleading as the conditioning of a linear system $\mathrm{cond}(A) \neq \mathrm{cond}(A^{-1})$, and it is the latter that in general needs reducing. In fact, for linear systems $Ax = b$, we illustrated empirically that it is $\omega^{-2} \cong \omega(A^{-2})$ that needs reducing. In addition, we found a new optimal diagonal preconditioner using this new measure, and we verified its strengths empirically. This is currently of theoretical interest only as exploiting ω^{-2} without evaluating A^{-1} first is still an open question.

We have used the differentiability and simplicity of trace and determinant in $\omega(A)$ to find optimal parameters for improving condition numbers for: low rank updates that arise in the context of nonsmooth Newton methods; and for preconditioning for linear systems. We empirically show that the ω-optimal preconditioners obtained in this work improve the performance of iterative methods.

The ω-condition number, when compared to the classical κ-condition number, is significantly more closely correlated to reducing the number of iterations and time for iterative methods for positive definite linear systems. This matches known results that show that preconditioning for clustering of eigenvalues helps in iterative methods, i.e., using all the eigenvalues rather than just the largest and smallest is desirable. This is further evidenced by the empirics that show that $\omega(A)$ is a significantly better estimate of the true conditioning of a linear system, i.e., how perturbations in the data A, b effect the solution x.

Finally, we have shown that an exact evaluation of $\omega(A)$ can be found using either the Cholesky or LU factorization. This is in contrast to the evaluation of $\kappa(A)$ that requires a spectral decomposition or a $\|A\|\|A^{-1}\|$ evaluation.

In a future study we hope to continue on exploiting the measure $\omega(A^{-2})$ by finding appropriate approximations, e.g., [38]. The complexity of the denominator is unchanged, but estimating the trace of an *inverse* is a more difficult problem. The condition number is also important in complexity analysis of optimization methods, e.g., in the convergence of conjugate gradient type methods. We hope to avoid the worst case analysis to get a more average case using ω.

Finally we note that the results we presented here can be extended beyond A positive definite by replacing eigenvalues with singular values in the definition of $\omega(A)$. The denominator is essentially unchanged but the numerator is the nuclear norm over n. The nuclear norm has been heavily studied and used in the recent literature, e.g., for low rank matrix completion problems.

Data Availability   The Matlab source code and data of all the experiments in this manuscript are available at https://github.com/DavidTBelen/omega-condition-number.

Declarations

Conflict of interest   The authors declare they have no conflict of interest.

A   Tables

We now present the tables for the empirics for the five preconditioners in Section 4.1. We use matrices from the SuiteSparse Matrix Collection. The instances in which MATLAB's **pcg** stops before converging to the desired tolerance are indicated with a symbol >. Observe that this could happen whenever two consecutive iterates are the same or when the solver detects that the problem is too ill-conditioned. In the latter case, there might be instances where not a single iteration is performed, this is indicated as − in the Table A.1.

name	n	$nnz(W)$	NONE	DIAG	ITRIU	ICHOL(1)	ICHOL(2)
mhd4800b	4800	27520	>97633	26	19	37	37
s3rmt3m3	5357	207123	>99172	14283	14134	>15207	>14300
ex15	6867	98671	>98296	46299	45029	>99102	>47275
bcsstk38	8032	355460	-	10104	8837	14264	14267
aft01	8205	125567	>8452	786	780	610	100
nd3k	9000	3279690	6012	9245	8632	>8007	1599
bloweybq	10001	49999	-	-	>11	-	-
msc10848	10848	1229776	56719	5274	4767	5328	2634
t2dah_e	11445	176117	>99495	33	29	28	7
olafu	16146	1015156	>90196	28028	22572	27670	12232
gyro	17361	1021159	28942	11605	11684	9964	1904
nd6k	18000	6897316	6589	9857	10574	8515	1831
raefsky4	19779	1316789	-	82865	81551	>87212	>17990
LFAT5000	19994	79966	-	>4984	>5037	-	-
msc23052	23052	1142686	-	>91699	>91700	>99722	>98530
smt	25710	3749582	9764	3343	3273	2803	514

Table A.1: preconditioners: number of iterations

name	n	$nnz(W)$	NONE	DIAG	ITRIU	ICHOL(1)	ICHOL(2)
mhd4800b	4800	27520	>2.70	0.00	0.00	0.01	0.00
s3rmt3m3	5357	207123	>8.54	1.76	2.16	>2.07	>2.75
ex15	6867	98671	>6.54	2.63	2.76	>14.68	>11.94
bcsstk38	8032	355460	>18.72	1.94	1.83	3.36	5.22
aft01	8205	125567	>0.66	0.07	0.07	0.15	0.03
nd3k	9000	3279690	13.13	19.78	47.35	>19.50	4.15
bloweybq	10001	49999	>5.41	>5.35	>5.97	>25.40	>24.97
msc10848	10848	1229776	38.88	3.78	3.30	6.19	6.35
t2dah_e	11445	176117	>12.74	0.01	0.01	0.02	0.01
olafu	16146	1015156	>59.21	16.55	15.86	30.59	22.62
gyro	17361	1021159	19.61	7.39	8.50	18.61	6.47
nd6k	18000	6897316	30.31	46.16	124.02	43.77	10.00
raefsky4	19779	1316789	>75.59	62.72	89.31	>223.45	>68.90
LFAT5000	19994	79966	>10.08	>10.16	>10.27	>24.52	>25.49
msc23052	23052	1142686	>78.53	>77.28	>79.37	>106.01	>173.72
smt	25710	3749582	25.31	8.63	13.13	15.18	4.85

Table A.2: preconditioners: total time

name	n	$nnz(W)$	NONE	DIAG	ITRIU	ICHOL(1)	ICHOL(2)
mhd4800b	4800	27520	-	5.860e-02	7.919e-02	5.221e-05	4.173e-05
s3rmt3m3	5357	207123	-	5.149e-02	4.339e-02	-	-
ex15	6867	98671	-	2.327e+00	2.335e+00	-	-
bcsstk38	8032	355460	-	1.148e-01	6.272e-02	8.006e-05	7.169e-05
aft01	8205	125567	-	2.083e-04	3.355e-04	8.085e-05	5.592e-05
nd3k	9000	3279690	9.084e-05	1.089e-04	5.620e-04	-	8.861e-05
bloweybq	10001	49999	-	-	-	-	-
msc10848	10848	1229776	8.705e-05	2.156e-03	1.563e-03	1.002e-04	7.700e-05
t2dah_e	11445	176117	-	1.338e-04	1.784e-03	8.938e-05	8.211e-05
olafu	16146	1015156	-	2.182e-03	7.755e-04	1.118e-04	9.864e-05
gyro	17361	1021159	1.289e-04	1.955e-04	2.167e-04	1.234e-04	1.174e-04
nd6k	18000	6897316	1.324e-04	1.602e-04	7.317e-04	1.325e-04	1.327e-04
raefsky4	19779	1316789	-	2.273e-01	2.573e-01	-	-
LFAT5000	19994	79966	-	-	-	-	-
msc23052	23052	1142686	-	-	-	-	-
smt	25710	3749582	1.450e-04	2.729e-04	2.658e-04	1.530e-04	1.310e-04

Table A.3: preconditioners: residual $\|Wx - b\|$. The symbol $-$ is used for the instances where **pcg** did not converge to the desired relative residual

name	n	$nnz(W)$	DIAG	ITRIU	ICHOL(1)	ICHOL(2)
mhd4800b	4800	27520	1.089e-03	3.193e-03	2.079e-03	7.248e-04
s3rmt3m3	5357	207123	6.965e-04	2.426e-03	1.644e-03	1.831e-03
ex15	6867	98671	4.206e-04	1.234e-03	9.327e-04	2.310e-03
bcsstk38	8032	355460	7.168e-04	5.815e-03	2.181e-03	2.136e-03
aft01	8205	125567	5.021e-04	2.030e-03	1.146e-03	1.925e-03
nd3k	9000	3279690	2.506e-03	7.995e-02	1.712e-02	1.676e-02
bloweybq	10001	49999	5.385e-04	1.365e-03	4.774e-04	4.115e-04
msc10848	10848	1229776	1.481e-03	8.259e-03	8.472e-03	2.943e-02
t2dah_e	11445	176117	7.269e-04	2.055e-03	2.599e-03	5.342e-03
olafu	16146	1015156	1.789e-03	1.623e-02	6.528e-03	1.793e-02
gyro	17361	1021159	1.730e-03	1.107e-02	1.736e-02	6.125e-02
nd6k	18000	6897316	5.121e-03	2.661e-01	3.337e-02	3.741e-02
raefsky4	19779	1316789	1.755e-03	2.275e-02	2.149e-02	6.287e-02
LFAT5000	19994	79966	8.536e-04	1.299e-03	8.915e-04	1.077e-03
msc23052	23052	1142686	2.056e-03	7.827e-03	5.936e-03	1.119e-02
smt	25710	3749582	3.886e-03	1.172e-01	5.734e-02	2.481e-01

Table A.4: Times (cpu) for computing the preconditioners

Index

References

[1] L. Bergamaschi, *A survey of low-rank updates of preconditioners for sequences of symmetric linear systems*, Algorithms, 13 (2020). 4

[2] L. Bergamaschi, R. Bru, and A. Martínez, *Low-rank update of preconditioners for the inexact Newton method with SPD Jacobian*, Mathematical and Computer Modelling, 54 (2011), pp. 1863–1873. Mathematical models of addictive behaviour, medicine & engineering. 4

[3] Y. Censor, , W. Moursi, T. Weames, and H. Wolkowicz, *Regularized nonsmooth Newton algorithms for best approximation with applications*, tech. rep., University of Waterloo, Waterloo, Ontario, 2022 under revision. 37 pages, research report, under revision. 3, 23, 33

[4] K. Chen, *Matrix preconditioning techniques and applications*, vol. 19 of Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge, 2005. 3, 7

[5] X. Chen, R. S. Womersley, and J. J. Ye, *Minimizing the condition number of a Gram matrix*, SIAM J. Optim., 21 (2011), pp. 127–148. 23

[6] S. Cipolla and J. Gondzio, *Proximal Stabilized Interior Point Methods and Low − Frequency − Update Preconditioning Techniques*, J. Optim. Theory Appl., 197 (2023), pp. 1061–1103. 23

[7] W. C. Davidon, *Optimally conditioned optimization algorithms without line searches*, Mathematical Programming, 9 (1975), pp. 1–30. 4

[8] J. Demmel, *On condition numbers and the distance to the nearest ill-posed problem*, Numer. Math., 51 (1987), pp. 251–289. 13

[9] ——, *The probability that a numerical analysis problem is difficult*, Math. Comp., 50 (1988), pp. 449–480. 4

[10] J. Dennis Jr. and R. Schnabel, *Least change secant updates for quasi-Newton methods*, SIAM Review, 21 (1979), pp. 443–459. 23

[11] J. Dennis Jr. and H. Wolkowicz, *Sizing and least-change secant methods*, SIAM J. Numer. Anal., 30 (1993), pp. 1291–1314. 3, 5, 7, 8, 14, 15

[12] X. Doan, S. Kruk, and H. Wolkowicz, *A robust algorithm for semidefinite programming*, Optim. Methods Softw., 27 (2012), pp. 667–693. 7, 8

[13] E. Dolan and J. Moré, *Benchmarking optimization software with performance profiles*, Math. Program., 91 (2002), pp. 201–213. 35

[14] W. Freedman, *Minimizing areas and volumes and a generalized AM-GM inequality*, Math. Mag., 85 (2012), pp. 116–123. 14

[15] W. Gao, Z. Qu, M. Udell, and Y. Ye, *Scalable approximate optimal diagonal preconditioning*, 2023. arXiv, 2312.15594. 5, 7, 11

[16] G. Golub and C. Van Loan, *Matrix computations*, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, fourth ed., 2013. 7

[17] J. Gondzio, S. Pougkakiotis, and J. Pearson, *General-purpose preconditioning for regularized interior point methods*, Comput. Optim. Appl., 83 (2022), pp. 727–757. 23

[18] N. Gould and J. Scott, *The state-of-the-art of preconditioners for sparse linear least-square problems*, ACM Trans. Math. Software, 43 (2017), pp. Art. 36, 35. 7, 18, 33

[19] A. Greenbaum, *Iterative methods for solving linear systems*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997. 7, 15

[20] C. Greif and J. M. Varah, *Minimizing the condition number for small rank modifications*, SIAM J. Matrix Anal. Appl., 29 (2006/07), pp. 82–97. 23

[21] R. Grimes and J. Lewis, *Condition number estimation for sparse matrices*, SIAM J. Sci. Statist. Comput., 2 (1981), pp. 384–388. 11

[22] W. W. Hager, *Condition estimates*, SIAM J. Sci. Statist. Comput., 5 (1984), pp. 311–316. 5, 11, 15

[23] D. Higham, *Condition numbers and their condition numbers*, Linear Algebra Appl., 214 (1995), pp. 193–213. 4, 13

[24] N. Higham, *A survey of condition number estimation for triangular matrices*, SIAM Rev., 29 (1987), pp. 575–596. 11

[25] N. Higham and F. Tisseur, *A block algorithm for matrix 1-norm estimation, with an application to 1-norm pseudospectra*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1185–1201 (electronic). 11

[26] H. Hu, H. Im, X. Li, and H. Wolkowicz, *A semismooth Newton-type method for the nearest doubly stochastic matrix problem*, Math. Oper. Res., 49 (2024). arxiv.org/abs/2107.09631, 35 pages. 23

[27] S. P. Kolodziej, M. Aznaveh, M. Bullock, J. David, T. A. Davis, M. Henderson, Y. Hu, and R. Sandstrom, *The suitesparse matrix collection website interface*, Journal of Open Source Software, 4 (2019), p. 1244. 16, 33

[28] O. Mangasarian, *Nonlinear Programming*, McGraw-Hill, New York, NY, 1969. 6, 30

[29] K. S. Miller, *On the inverse of the sum of matrices*, Mathematics magazine, 54 (1981), pp. 67–72. 27

[30] S. S. Oren and D. G. Luenberger, *Self-scaling variable metric (ssvm) algorithms: Part i: Criteria and sufficient conditions for scaling a class of algorithms*, Management Science, 20 (1974), pp. 845–862. 5

[31] J. Pearson and J. Pestana, *Preconditioners for Krylov subspace methods: an overview*, GAMM-Mitt., 43 (2020), pp. e202000015, 35. 4

[32] G. Pini and G. Gambolati, *Is a simple diagonal scaling the best preconditioner for conjugate gradients on supercomputers?*, Advances in Water Resources, 13 (1990), pp. 147–153. 7

[33] H. Qi and D. Sun, *A quadratically convergent Newton method for computing the nearest correlation matrix*, SIAM journal on matrix analysis and applications, 28 (2006), pp. 360–385. 23

[34] R. Schnabel, *Analysing and improving quasi-Newton methods for unconstrained optimization*, PhD thesis, Department of Computer Science, Cornell University, Ithaca, NY, 1977. Also available as TR-77-320. 23

[35] J. Stoer and R. Bulirsch, *Introduction to Numerical Analysis*, vol. 1993, Springer, 1980. 4, 15

[36] G. Strang, *Linear Algebra and Its Applications 4th ed.*, 2012. 15

[37] A. van der Sluis, *Condition numbers and equilibration of matrices*, Numer. Math., 14 (1969/1970), pp. 14–23. 7

[38] L. Wu, J. Laeuchli, V. Kalantzis, A. Stathopoulos, and E. Gallopoulos, *Estimating the trace of the matrix inverse by interpolating from the diagonal of an approximate inverse*, Journal of Computational Physics, 326 (2016), pp. 828–844. 40