A Diagonalizable Obstruction for Random Two-Step Ritz Compressions

August 3, 2026

Abstract

We consider two-step Krylov compression of an arbitrary real matrix $A \in \mathbb{R}^{n\times n}$ from a standard Gaussian starting vector $b \sim N(0, I_n)$. On the event that $\mathcal{K}_2(A, b)$ is two-dimensional, let Q denote any orthonormal basis of this space; the associated Ritz matrix is $Q^T AQ$. For every even $n \geq 4$, we construct a real diagonalizable matrix A_n with $\|A_n\|_2 \leq 1 + e^{-n}/2 < 3/2$ such that $\mathcal{K}_2(A_n, b)$ is two-dimensional almost surely and

$$\Pr\left\{\kappa_V(Q^T A_n Q) \geq \frac{e^n}{5}\right\} \geq 1 - 2e^{-n/40}.$$

For $\varepsilon_n = 5e^{-2n}$, the same family satisfies

$$\mathbb{E}\big[\operatorname{area}\Lambda_{\varepsilon_n}(Q^T A_n Q)\big] \geq \left(1 - 2e^{-n/40}\right)\frac{\pi e^{-2n}}{4}.$$

Consequently, for every positive polynomial p, the inequality $\kappa_V(Q^T A_n Q) \leq p(n)$ fails with probability tending to one exponentially fast along the even dimensions. Moreover, for every fixed $\beta > 1$ and every positive polynomial p, no upper bound of the form $p(n)\varepsilon^\beta$ holds uniformly over all dimensions, inputs, and positive tolerances under this Gaussian starting-vector model.

Note. This paper was fully AI generated by an internal harness using GPT 5.6 Sol xhigh, with a minimal prompt from Richard Peng.

Contents

1 Introduction

Krylov subspace methods replace a large matrix eigenvalue problem by a projection onto a space generated from a starting vector. Given $A \in \mathbb{C}^{n\times n}$ and $b \in \mathbb{C}^n$, the k-step space and its orthogonal compression are

$$\mathcal{K}_k(A,b) = \operatorname{span}\{b, Ab, \ldots, A^{k-1}b\}, \qquad H_k = Q^*AQ,$$

where the columns of Q form an orthonormal basis for $\mathcal{K}_k(A,b)$. The eigenvalues of the smaller matrix H_k are the Ritz values produced by the Arnoldi process, a basic tool for large nonsymmetric eigenvalue problems [Arn51,Saa80]. For a nonnormal compression, however, the location of the Ritz values is only part of the stability question. The condition number of an eigenvector matrix governs eigenvalue sensitivity through the Bauer–Fike theorem, while the pseudospectrum records the same sensitivity without requiring diagonalizability [BF60, Tre99].

Randomizing the starting vector is a natural way to avoid exceptional Krylov spaces. Such starts have long supported positive convergence guarantees for the power and Lanczos methods on symmetric positive-definite matrices [KW92]. That setting does not, though, test nonnormal eigenvector stability: if A is Hermitian, then every Q^*AQ is Hermitian, has eigenvector condition number one, and has ε-pseudospectral area at most $\pi k\varepsilon^2$. At the other extreme, deterministic Arnoldi compressions have considerable adversarial freedom. Duintjer Tebbens and Meurant showed that prescribed Ritz-value histories can be realized [DM12]; this is a result about Ritz values for chosen matrix–vector pairs, rather than about eigenvector conditioning under a random start.

Amsel et al. [ABB+26] posed the corresponding random-start stability question for an arbitrary matrix A. They asked whether $\kappa_V(H_k)$ is bounded by a polynomial in n with high probability, or alternatively whether the expected pseudospectral area admits a bound of the form

$$\mathbb{E}[\operatorname{area}\Lambda_\varepsilon(H_k)] \le \operatorname{poly}(n)\varepsilon^\beta$$

for an exponent β close to 2. Their formulation also records an important deterministic warning: for the cyclic shift and the special start $b = e_1$, the compression before full dimension is a Jordan block. The point of randomizing b is precisely to ask whether such badly aligned starts are negligible. The distribution of the random vector and the intended range of ε are not fixed in that formulation. Here we use the standard real-Gaussian model $b \sim N(0, I_n)$ and consider the literal quantification over arbitrary inputs.

The closest positive comparison changes the random object. Shah [Sha25] studies the same expected-area observable when Q spans an independent Haar-random complex subspace, proving bounds with ε-exponents 6/5, 4/3, and 2 (up to polynomial and logarithmic factors) under successively stronger hypotheses on A. A random-start Krylov space is not a Haar-random plane: its distribution is coupled to A through $b, Ab, \ldots$. Our result shows that this distinction is decisive for arbitrary inputs. Even at $k = 2$, randomizing the start does not prevent exponentially ill-conditioned Ritz eigenvectors, and the same examples rule out every uniform expected-area exponent greater than one. The precise statement is as follows.

Theorem 1.1 (Diagonalizable obstruction for two-step Ritz compressions)**.** *For every even $n \ge 4$, there exists a diagonalizable matrix $A_n \in \mathbb{R}^{n\times n}$ with $\|A_n\|_2 \le 1 + e^{-n}/2 < 3/2$ having the following properties. Let $b \sim N(0, I_n)$ and set*

$$\mathcal{K}_2(A_n, b) = \operatorname{span}\{b, A_nb\}.$$

This Krylov space is two-dimensional almost surely. On this event, let Q be any orthonormal basis of the Krylov space. Then

$$\Pr\left\{\kappa_V(Q^T A_n Q) \geq \frac{e^n}{5}\right\} \geq 1 - 2e^{-n/40}.$$

Moreover, for $\varepsilon_n = 5e^{-2n}$,

$$\mathbb{E}\big[\operatorname{area}\Lambda_{\varepsilon_n}(Q^T A_n Q)\big] \geq \left(1 - 2e^{-n/40}\right)\frac{\pi e^{-2n}}{4}.$$

Consequently, no positive polynomial $p(n)$ bounds $\kappa_V(Q^T A_n Q)$ with high probability uniformly over arbitrary inputs, and for every fixed $\beta > 1$ no upper bound of the form $p(n)\varepsilon^\beta$ holds uniformly for the expected pseudospectral area.

Several features of the theorem delimit the obstruction. First, it is not a defective-input example: A_n is diagonalizable and uniformly bounded in operator norm. Its two eigenspaces are nevertheless exponentially close. The minimal polynomial has degree two, so the two-step Krylov plane is almost surely invariant; with probability $1 - e^{-\Omega(n)}$, the two spectral components of the Gaussian start are nearly antiparallel. Orthogonal compression preserves their angle, which becomes the angle between the two Ritz eigenvectors. A direct two-dimensional Schur-form estimate then places a disk of radius $e^{-n}/2$ in the $5e^{-2n}$-pseudospectrum. Thus the conditioning and area conclusions arise from the same geometry.

Second, the expected-area conclusion is witnessed at the exponentially small tolerance $\varepsilon_n = 5e^{-2n}$. It therefore rules out estimates intended to hold uniformly for all positive ε, but does not rule out a statement restricted to inverse-polynomial tolerances. Likewise, the construction does not address normal inputs, the cyclic shift, prescribed growing Krylov dimensions, a complex-Gaussian start, or bounds allowed to depend on the conditioning of an eigenbasis for A. These are genuine structural distinctions. For example, unitary Arnoldi admits an orthogonal-polynomial description [Gra93], while known results on Ritz values of normal matrices impose geometry absent from general nonnormal inputs [CH13]. Neither line is a conditioning theorem for the random-start Krylov compression considered here.

Since nonzero rescaling of b does not change its Krylov space, our Gaussian model is equivalently a uniform random direction on the real unit sphere; it should not be confused with an independent Haar-random two-dimensional subspace. Section 2 fixes the conditioning and pseudospectral conventions. Section 3 explains the construction and its geometric mechanism, and Section 4 proves the main theorem.

2 Preliminaries

Throughout, vector norms are Euclidean and matrix norms are the induced operator norms. A standard Gaussian vector $b \sim N(0, I_n)$ has independent $N(0,1)$ coordinates. In particular, if $n = 2r$ and $b = (u, v)$ under the block decomposition $\mathbb{R}^n = \mathbb{R}^r \oplus \mathbb{R}^r$, then u and v are independent $N(0, I_r)$ vectors.

Krylov spaces and Ritz matrices. For $A \in \mathbb{R}^{n\times n}$, $b \in \mathbb{R}^n$, and $k \geq 1$, define

$$\mathcal{K}_k(A, b) = \operatorname{span}\{b, Ab, \ldots, A^{k-1}b\}.$$

If this space has dimension d and $Q \in \mathbb{R}^{n\times d}$ has orthonormal columns spanning it, then $H = Q^T A Q$ is the associated *Ritz matrix*. A different orthonormal basis replaces H by an orthogonally similar matrix. If $\mathcal{K}_k(A,b)$ is A-invariant, then

$$AQ = QH. \tag{1}$$

In particular, if $x \in \mathcal{K}_k(A,b)$ is an eigenvector of A, then $Q^T x$ is an eigenvector of H with the same eigenvalue. Moreover, Q^T preserves norms and inner products on $\mathcal{K}_k(A,b)$.

Angles and eigenvector conditioning. For nonzero $x, y \in \mathbb{R}^d$, let

$$c(x,y) = \frac{|x^T y|}{\|x\|_2 \|y\|_2}, \qquad s(x,y) = \sqrt{1 - c(x,y)^2}.$$

Thus $s(x,y)$ is the sine of the acute angle between the lines spanned by x and y, and

$$s(x,y)^2 = \frac{\|x\|_2^2 \|y\|_2^2 - (x^T y)^2}{\|x\|_2^2 \|y\|_2^2}. \tag{2}$$

If a real diagonalizable 2×2 matrix H has distinct eigenvalues, let $V = [x_+ \; x_-]$ have unit eigenvector columns and define $\kappa_V(H) = \|V\|_2 \|V^{-1}\|_2$. This definition is unchanged by signs or ordering of the two eigenvectors. Writing $c = c(x_+, x_-)$ and $s = s(x_+, x_-)$, the eigenvalues of $V^T V$ are $1 + c$ and $1 - c$, so

$$\kappa_V(H) = \sqrt{\frac{1+c}{1-c}} = \frac{1+c}{s} \geq \frac{1}{s}. \tag{3}$$

Orthogonal similarity preserves this quantity.

Pseudospectra. For a square matrix H, write $\sigma_{\min}(H)$ and $\sigma_{\max}(H)$ for its smallest and largest singular values, respectively, and define its ε-pseudospectrum by

$$\Lambda_\varepsilon(H) = \{z \in \mathbb{C} : \sigma_{\min}(zI - H) \leq \varepsilon\}.$$

Area means planar Lebesgue measure on $\mathbb{C}$. Orthogonal similarity preserves singular values and hence preserves the pseudospectrum. Finally, for every 2×2 matrix M,

$$\sigma_{\min}(M)\sigma_{\max}(M) = |\det M|. \tag{4}$$

We also use the elementary bound $\sigma_{\max}(M) = \|M\|_2 \geq |M_{ij}|$ for each entry of M.

3 Overview

The proof is a two-eigenvalue counterexample whose geometry survives Krylov compression exactly. Write $n = 2r$, set $\delta = e^{-n}$ and $\alpha = \delta/2$, and consider

$$A_n = \begin{pmatrix} \alpha I_r & -I_r \\ 0 & -\alpha I_r \end{pmatrix}. \tag{5}$$

Its eigenspaces are $E_+ = \{(x, 0) : x \in \mathbb{R}^r\}$ and $E_- = \{(x, \delta x) : x \in \mathbb{R}^r\}$, with eigenvalues α and $-\alpha$, respectively. Thus the eigenspaces are separated by an angle of order δ, but A_n is nevertheless diagonalizable and has $\|A_n\|_2 \leq 1 + \alpha < 3/2$. The choice $2\alpha = \delta$ is what reconciles these two features: the exponentially ill-conditioned eigenbasis is multiplied by an equally small eigenvalue gap, leaving the off-diagonal block of A_n at constant scale.

We take the starting vector to be $b \sim N(0, I_n)$. For this compression the model is equivalent to choosing a uniform random direction, because rescaling b does not change its Krylov space. Writing $b = (u, v)$, split it into its two spectral components

$$p = (u - \delta^{-1}v, 0), \qquad m = (\delta^{-1}v, v). \tag{6}$$

They satisfy $b = p + m$, $A_n p = \alpha p$, and $A_n m = -\alpha m$. Consequently, almost surely,

$$\mathcal{K}_2(A_n, b) = \operatorname{span}\{p, m\} \tag{7}$$

is two-dimensional and A_n-invariant. This identity is the main reduction: the Ritz matrix $H = Q^T A_n Q$ is simply the restriction of A_n to the span of these two random spectral components, expressed in orthonormal coordinates. In particular, Q^T preserves the angle between p and m.

The core technical point is to show that the particular components selected by a random start inherit the exponentially small separation of the two eigenspaces with high probability. Let s be the sine of the acute angle between p and m, and put $a = \delta u - v$. The Gram determinant of the scaled vectors $(a, 0)$ and $(v, \delta v)$ gives

$$s^2 \le \delta^2 \left(1 + \frac{\|u\|_2^2}{\|\delta u - v\|_2^2}\right). \tag{8}$$

A Gaussian norm-ratio bound shows that $\|u\|_2 \le 2\|v\|_2$ except with probability at most $2e^{-r/20}$. On this event, $\|\delta u - v\|_2 \ge \|v\|_2/2$, so $s < 5\delta$. The two unit eigenvectors of H have this same angle, and the elementary angle formula for a 2×2 eigenvector matrix yields

$$\kappa_V(H) \ge \frac{1}{s} \ge \frac{1}{5\delta} = \frac{e^n}{5}. \tag{9}$$

This proves the high-probability conditioning obstruction in Theorem 1.1.

For the pseudospectral statement, large eigenvector condition number alone is not used as a black box; the proof extracts a quantitative disk directly from the same angle. In an orthonormal basis adapted to one eigenvector, H has the real Schur form

$$T = \begin{pmatrix} \alpha & t \\ 0 & -\alpha \end{pmatrix}, \qquad |t| = \delta \frac{\sqrt{1 - s^2}}{s}. \tag{10}$$

Although the eigenvalues $\pm\alpha$ are exponentially close, the estimate $s < 5\delta$ forces $|t| \ge 1/10$. For every $z \in \mathbb{C}$, the product of the singular values of $zI - T$ is $|z^2 - \alpha^2|$, whereas the larger singular value is at least $|t|$. Hence

$$\sigma_{\min}(zI - T) \le \frac{|z^2 - \alpha^2|}{|t|} \le 5\delta^2 \qquad \text{whenever } |z| \le \alpha.$$

Thus the disk of radius $\alpha = \delta/2$ lies in $\Lambda_{5\delta^2}(H)$ on the same high-probability event.

The disk has area $\pi\delta^2/4$. Multiplying by the probability of the norm-ratio event gives the expected-area lower bound in Theorem 1.1. Finally, at $\varepsilon_n = 5e^{-2n}$ this lower bound is exponentially larger than $p(n)\varepsilon_n^\beta$ for every polynomial p and every fixed $\beta > 1$. The conditioning and pseudospectral conclusions therefore come from one obstruction: orthonormal compression preserves the nearly coincident spectral directions, and that same geometry becomes a constant Schur coupling at the chosen exponential scale.

4 Proof of the Main Theorem

We first isolate the only probabilistic estimate needed in the construction.

Lemma 4.1 (Gaussian norm-ratio concentration)**.** *Let $u, v \in \mathbb{R}^r$ be independent standard Gaussian vectors. Then*

$$\Pr\{\|u\|_2 \le 2\|v\|_2\} \ge 1 - 2e^{-r/20}.$$

Proof. For a standard Gaussian vector $g \in \mathbb{R}^r$ and $X = \|g\|_2^2$, evaluation of the one-dimensional Gaussian integral gives

$$\mathbb{E}e^{tX} = (1-2t)^{-r/2} \quad (0 < t < 1/2), \qquad \mathbb{E}e^{-tX} = (1+2t)^{-r/2} \quad (t > 0).$$

Markov's inequality, with $t = 1/4$ for the first transform and $t = 1/2$ for the second, yields

$$\Pr\{X \ge 2r\} \le e^{-r/2}2^{r/2} \le e^{-r/20}, \qquad \Pr\{X \le r/2\} \le e^{r/4}2^{-r/2} \le e^{-r/20}.$$

Apply the upper-tail estimate to $\|u\|_2^2$ and the lower-tail estimate to $\|v\|_2^2$. Outside the union of the two exceptional events, $\|u\|_2^2 \le 2r$ and $\|v\|_2^2 \ge r/2$, which imply the claimed norm ratio. $\square$

We next define the input family and use the norm-ratio event to control the angle between the two spectral components of the start. This angle is exactly the geometric quantity governing the eigenvector condition number after orthonormal compression.

Let $n = 2r \ge 4$ be even, set

$$\delta = e^{-n}, \qquad \alpha = \frac{\delta}{2},$$

and define

$$A_n = \begin{pmatrix} \alpha I_r & -I_r \\ 0 & -\alpha I_r \end{pmatrix}. \tag{11}$$

For a vector $b = (u, v) \in \mathbb{R}^r \oplus \mathbb{R}^r$, define its spectral components

$$p = (u - \delta^{-1}v, 0), \qquad m = (\delta^{-1}v, v). \tag{12}$$

Lemma 4.2 (Krylov compression from nearly coincident eigenspaces)**.** *Let A_n be defined by Equation (11) and let $b = (u, v) \sim N(0, I_n)$. Then A_n is diagonalizable, $\|A_n\|_2 \le 1 + e^{-n}/2 < 3/2$, and $\mathcal{K}_2(A_n, b)$ is two-dimensional almost surely. On this event, let $Q \in \mathbb{R}^{n\times 2}$ be any orthonormal basis of the Krylov space. Then*

$$\Pr\left\{\kappa_V(Q^T A_n Q) \ge \frac{1}{5\delta}\right\} \ge 1 - 2e^{-r/20}.$$

Proof. The subspaces

$$E_+ = \{(x, 0) : x \in \mathbb{R}^r\}, \qquad E_- = \{(x, \delta x) : x \in \mathbb{R}^r\}$$

are complementary eigenspaces of A_n with eigenvalues α and $-\alpha$, respectively. Indeed,

$$A_n = \begin{pmatrix} I_r & I_r \\ 0 & \delta I_r \end{pmatrix} \begin{pmatrix} \alpha I_r & 0 \\ 0 & -\alpha I_r \end{pmatrix} \begin{pmatrix} I_r & I_r \\ 0 & \delta I_r \end{pmatrix}^{-1}.$$

Thus A_n is diagonalizable. Also,

$$A_n = \operatorname{diag}(\alpha I_r, -\alpha I_r) + \begin{pmatrix} 0 & -I_r \\ 0 & 0 \end{pmatrix},$$

so $\|A_n\|_2 \le \alpha + 1 = 1 + e^{-n}/2 < 3/2$.

The vectors in Equation (12) satisfy

$$b = p + m, \qquad A_n p = \alpha p, \qquad A_n m = -\alpha m.$$

Consequently,

$$\mathcal{K}_2(A_n, b) = \operatorname{span}\{p, m\} \tag{13}$$

whenever p and m are nonzero. For Gaussian u, v, this occurs almost surely. The vectors are then linearly independent because they belong to the complementary eigenspaces E_+ and E_-, so the Krylov space is two-dimensional and invariant under A_n.

Let s be the sine of the acute angle between the lines spanned by p and m. Scaling both vectors by δ does not alter this angle. Put $a = \delta u - v$; the scaled vectors are $(a, 0)$ and $(v, \delta v)$. Applying Equation (2) gives

$$\begin{aligned} s^2 &= \frac{\|a\|_2^2(1+\delta^2)\|v\|_2^2 - (a^T v)^2}{\|a\|_2^2(1+\delta^2)\|v\|_2^2} \\ &= \frac{\delta^2\left(\|u\|_2^2\|v\|_2^2 - (u^T v)^2 + \|a\|_2^2\|v\|_2^2\right)}{\|a\|_2^2(1+\delta^2)\|v\|_2^2} \\ &\le \delta^2\left(1 + \frac{\|u\|_2^2}{\|a\|_2^2}\right). \end{aligned}$$

On the event $\|u\|_2 \le 2\|v\|_2$, the inequality $\delta \le 1/4$ gives

$$\|a\|_2 \ge \|v\|_2 - \delta\|u\|_2 \ge \frac{1}{2}\|v\|_2.$$

It follows that

$$s \le \sqrt{17}\,\delta < 5\delta. \tag{14}$$

By Equation (1), the unit eigenvectors of $H = Q^T A_n Q$ have the same acute angle as p and m. Therefore Equations (3) and (14) give

$$\kappa_V(H) \ge \frac{1}{s} \ge \frac{1}{5\delta}.$$

The probability estimate follows from Lemma 4.1. $\square$

It remains to turn the same small-angle event into an area lower bound. In an orthonormal basis adapted to one Ritz eigenvector, the angle controls the off-diagonal entry of a triangular form. That entry supplies the singular value estimate needed for a pseudospectral disk.

Lemma 4.3 (Pseudospectral disk from Schur coupling)**.** *Under the assumptions of Lemma 4.2, set $\varepsilon_n = 5\delta^2$. Apart from the null event on which a spectral component in Equation (12) vanishes, the event $\|u\|_2 \le 2\|v\|_2$ implies*

$$\{z \in \mathbb{C} : |z| \le \alpha\} \subseteq \Lambda_{\varepsilon_n}(Q^T A_n Q).$$

Consequently,

$$\mathbb{E}\big[\operatorname{area}\Lambda_{\varepsilon_n}(Q^T A_n Q)\big] \ge (1 - 2e^{-r/20})\frac{\pi\delta^2}{4}.$$

Proof. Write $H = Q^T A_n Q$. Its eigenvalues are α and $-\alpha$. Let s and $c = \sqrt{1-s^2}$ be the sine and cosine of the acute angle between its two unit eigenvectors. Choosing the first eigenvector and its orthogonal complement as an orthonormal basis, the two eigenvector equations show that H is orthogonally similar to

$$T = \begin{pmatrix} \alpha & t \\ 0 & -\alpha \end{pmatrix}, \qquad |t| = 2\alpha \frac{c}{s} = \delta \frac{c}{s}. \tag{15}$$

On the stated event, Equation (14) gives $s \le 5\delta$. Because $n \ge 4$ implies $5\delta < 1/2$, we have $c > 1/2$, and hence

$$|t| \ge \frac{1}{10}. \tag{16}$$

For every $z \in \mathbb{C}$, Equation (4) and the entry bound $\sigma_{\max}(zI - T) \ge |t|$ give

$$\sigma_{\min}(zI - T) \le \frac{|z^2 - \alpha^2|}{|t|}.$$

If $|z| \le \alpha$, then Equation (16) makes the right-hand side at most

$$10(|z|^2 + \alpha^2) \le 20\alpha^2 = 5\delta^2 = \varepsilon_n.$$

Orthogonal similarity preserves singular values, proving the disk inclusion. The disk has area $\pi\alpha^2 = \pi\delta^2/4$. Multiplying this deterministic lower bound by the event probability in Lemma 4.1, and using nonnegativity of area on the complementary event, proves the expectation bound. □

We now combine the geometric and pseudospectral estimates and verify both uniform impossibility statements in Theorem 1.1.

Proof of Theorem 1.1. Take the matrix A_n in Equation (11). By Lemma 4.2, it is diagonalizable, $\|A_n\|_2 \le 1 + e^{-n}/2 < 3/2$, and its two-step Gaussian Krylov space has dimension two almost surely. Since $r = n/2$ and $\delta = e^{-n}$, the same lemma gives

$$\Pr\left\{ \kappa_V(Q^T A_n Q) \ge \frac{e^n}{5} \right\} \ge 1 - 2e^{-n/40}.$$

For every fixed positive polynomial p, one has $e^n/5 > p(n)$ for all sufficiently large n. Hence $\kappa_V(Q^T A_n Q) \le p(n)$ fails with probability tending to one exponentially fast along the even dimensions.

The expectation bound in the theorem is exactly Lemma 4.3 after substituting $r = n/2$ and $\delta = e^{-n}$. Finally, fix $\beta > 1$ and a positive polynomial p. At $\varepsilon_n = 5e^{-2n}$,

$$\frac{\mathbb{E}[\text{area}\,\Lambda_{\varepsilon_n}(Q^T A_n Q)]}{p(n)\varepsilon_n^\beta} \ge \frac{(1 - 2e^{-n/40})\pi}{4 \cdot 5^\beta p(n)} e^{2(\beta-1)n}.$$

This ratio tends to infinity along the even dimensions, so no uniform $p(n)\varepsilon^\beta$ upper bound can hold. This completes the proof. □

References

[ABB$^+$26] Noah Amsel, Yves Baumann, Paul Beckman, Peter Bürgisser, Chris Camaño, Tyler Chen, Edmond Chow, Anil Damle, Michal Derezinski, Mark Embree, Ethan N. Epperly, Robert Falgout, Mark Fornace, Anne Greenbaum, Chen Greif, Diana Halikias, Zhen Huang, Elias Jarlebring, Yiannis Koutis, Daniel Kressner, Rasmus Kyng, Jörg Liesen, Jackie Lok, Raphael A. Meyer, Yuji Nakatsukasa, Kate Pearce, Richard Peng, David Persson, Eliza Rebrova, Ryan Schneider, Rikhav Shah, Edgar Solomonik, Nikhil Srivastava, Alex Townsend, Robert J. Webber, and Jess Williams. Linear systems and eigenvalue problems: Open questions from a Simons workshop, 2026. 1

[Arn51] Walter Edwin Arnoldi. The principle of minimized iterations in the solution of the matrix eigenvalue problem. *Quarterly of Applied Mathematics*, 9(1):17–29, 1951. 1

[BF60] F. L. Bauer and C. T. Fike. Norms and exclusion theorems. *Numerische Mathematik*, 2:137–141, 1960. 1

[CH13] Russell Carden and Derek J. Hansen. Ritz values of normal matrices and Ceva's theorem. *Linear Algebra and its Applications*, 438(11):4114–4129, 2013. 2

[DM12] Jurjen Duintjer Tebbens and Gérard Meurant. Any Ritz value behavior is possible for Arnoldi and for GMRES. *SIAM Journal on Matrix Analysis and Applications*, 33(3):958–978, 2012. 1

[Gra93] William B. Gragg. Positive definite Toeplitz matrices, the Arnoldi process for isometric operators, and Gaussian quadrature on the unit circle. *Journal of Computational and Applied Mathematics*, 46(1–2):183–198, 1993. 2

[KW92] Jacek Kuczyński and Henryk Woźniakowski. Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start. *SIAM Journal on Matrix Analysis and Applications*, 13(4):1094–1122, 1992. 1

[Saa80] Yousef Saad. Variations on Arnoldi's method for computing eigenelements of large unsymmetric matrices. *Linear Algebra and its Applications*, 34:269–295, 1980. 1

[Sha25] Rikhav Shah. The pseudospectrum of random compressions of matrices, 2025. 1

[Tre99] Lloyd N. Trefethen. Computation of pseudospectra. *Acta Numerica*, 8:247–295, 1999. 1