Convergence of the Arnoldi Iteration for Estimating Extreme Eigenvalues

by

Cecilia Chen

B.S. Computer Science and Engineering, Mathematics
Massachusetts Institute of Technology, 2024

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

MASTER OF ENGINEERING IN ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2025

Authored by:	Cecilia Chen Department of Electrical Engineering and Computer Science Jan 22, 2025
Certified by:	John Urschel Assistant Professor of Mathematics, Thesis Supervisor
Accepted by:	Katrina LaCurts Chair Master of Engineering Thesis Committee

Convergence of the Arnoldi Iteration for Estimating Extreme Eigenvalues

by

Cecilia Chen

Submitted to the Department of Electrical Engineering and Computer Science
on Jan 22, 2025 in partial fulfillment of the requirements for the degree of

MASTER OF ENGINEERING IN ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

ABSTRACT

Krylov subspace methods, like the Arnoldi iteration, are a powerful tool for efficiently solving high-dimensional linear algebra problems. In this work, we analyze the convergence of Krylov methods for estimating the numerical range of a matrix. Prior bounds on approximation error often depend on eigenvalue gaps of the matrix, which lead to weaker bounds than observed in practice, specifically in applications where these gaps are small. Instead, we extend a line of work proving gap-independent bounds for the Lanczos method, which depend only on the matrix dimensions and number of iterations, to the more general Arnoldi case.

Thesis supervisor: John Urschel
Title: Assistant Professor of Mathematics

Contents

Chapter 1

Introduction

1.1 Introduction

Krylov methods are used in a variety of scientific computing applications such as control theory [1] and quantum mechanics [2], to give a few examples out of many. The aim of these methods is to reduce the problem size by approximating properties of high dimensional matrices with a much smaller Krylov subspace.

Definition 1.1. *For $A \in \mathbb{C}^{n\times n}$ and starting vector $b \in \mathbb{C}^n$, the m-dimensional Krylov subspace is defined as*

$$\mathcal{K}_m(A, b) := \operatorname{span}\{b, Ab, A^2b, \ldots, A^{m-1}b\}.$$

In extremal eigenvalue problems, Krylov methods such as the Lanczos iteration and the Arnoldi iteration estimate the extremal eigenvalues of A by projecting A onto a Krylov subspace and computing the eigenvalues of the resulting smaller, well-structured matrix. The smaller the number of dimensions m (corresponding to the number of iterations), the faster the computation. Thus, it is of theoretical and practical interest to understand how quickly the approximation converges with the number of iterations m.

In this work, we will prove error bounds for eigenvalue estimates using Krylov methods when A is normal. First, we show that for any extremal eigenvalue, a maximum error of $\varepsilon \cdot \operatorname{diam}(W(A))$ can be achieved with only $m = O(\ln n/\varepsilon)$ iterations, where $W(A)$ is the numerical range of A. Then, we prove a matching lower bound, constructing a sequence of matrices $(A_m)_{m\in\mathbb{N}}$ such that we incur at least $O(1/m)$ error after m iterations. Finally, we examine how small deviations from normality can still give us similar error bounds.

The mathematical results in this thesis can be found in the preprint [3] by the author and advisor John Urschel. However, proofs in this thesis have been rewritten to be more accessible to a general audience and give additional background.

1.2 Related Work

Historically, most theoretical convergence bounds for Krylov methods depend on the eigenvalue gaps of the matrix. To give some intuition, we can consider the power iteration. The

power iteration is one of the simplest Krylov methods and is used estimate a single eigenvector. If A is normal, we can rewrite b in terms of its components along the eigenvectors $\nu_1, \dots \nu_n$ of A. Then for each successive power of A, we have

$$b = \sum_{i=1}^{n} \alpha_i \nu_i \text{ and } A^m b = \sum_{i=1}^{n} \alpha_i \lambda_i^m \nu_i.$$

As m increases, the largest eigenvalue λ_1 dominates and $A^m b$ converges to ν_1. However, the rate of convergence depends highly on the ratio of λ_1 and the next largest eigenvalue λ_2. Hence, we need a large eigenvalue gap distinguishing λ_1 from the rest.

For the Arnoldi iteration, which is a Krylov method for estimating eigenvalues for general matrices, Lemma 6.2 of [4] gives a similar type bound on the distance between the Krylov subspace and any particular eigenvector. Let Q_m be the projection onto $\mathcal{K}_m(A, b)$ and $\mathcal{P}^*_{m-1}$ be the space of polynomials of degree at most $m-1$ such that $p(\lambda_i) = 1$. Then we have

$$\|\nu_i - Q_m \nu_i\|_2 \leq \left[\sum_{j \neq i} \frac{|\alpha_j|}{|\alpha_i|}\right] \left[\min_{p \in \mathcal{P}^*_{m-1}} \max_{j \neq i} |p(\lambda_j)|\right]. \tag{1.1}$$

Here, a small gap between λ_i and λ_j implies that the right-hand term cannot be very small since $p(\lambda_j)$ will be close to $p(\lambda_i) = 1$ for any polynomial p.

In practice, the Arnoldi iteration performs much better in estimating eigenvalues, even when the eigenvalue gap is small. In fact, Musco and Musco argue in [5] that bounds on distances to any particular eigenvector such as (1.1) are overkill in cases where only eigenvalues are desired. While it is clear that having a good estimate of an eigenvector will give a good estimate of its corresponding eigenvalue, it is not true that a good eigenvalue estimate requires a good estimate of its corresponding eigenvector. For example, suppose λ_1 and λ_2 are close, and our Krylov subspace is both far from ν_1 and close to ν_2. We would have a good estimate of λ_2 instead, but that is already close to λ_1!

As a result, there is interest in proving gap-independent bounds, which depend only on the matrix size and number of iterations, better capturing the behavior of Krylov methods in practice. In [6], Kuczynski and Wozniakowski give gap-independent bounds for the Lanczos iteration, which is essentially a specialized form of the Arnoldi iteration for Hermitian matrices. Kuczynski and Wozniakowski show that for a random starting vector b, the expected relative error is upper bounded by $O(\ln^2 n/m^2)$ after m iterations. This bound was later improved by Urschel in [7] to $O(\ln n/m^2)$.

This work aims to expand on these results for Hermitian matrices by generalizing to normal matrices. To date, gap-independent bounds for normal matrices have not been shown. Compared to Hermitian matrices, the challenge of normal matrices stems from the possibility of complex eigenvalues, which we will show causes slower convergence than the Hermitian case.

1.3 Background

In this section, we discuss notation and other background.

1.3.1 Notation

For a given matrix $A \in \mathbb{C}^{n\times n}$, the Arnoldi method is a way of computing a projection Q_m such that $\text{span}(Q_m) = \mathcal{K}_m(A, \boldsymbol{b})$, and $AQ_m = Q_m H_m$. We are concerned with the projection of A onto the Krylov subspace H_m, arguing that the numerical range of H_m is close to the numerical range of A. We also note that while the Arnoldi iteration is a specific way of computing the matrix H_m, we only care about the values of the projection and not the method of computation itself. Hence, our results hold for any Krylov-based method.

We use the convention that vectors are stated in bold font. We let $\Lambda(A)$ denote the spectrum of A and $W(A)$ the numerical range of A. Furthermore let $\mathcal{P}_m$ be the set of polynomials of degree at most m. Let $\text{conv}(\cdot)$ refer to the convex hull of a set. To measure the error between numerical ranges, which are convex bodies in $\mathbb{C}$, we use the standard, one-sided Hausdorff, and Hausdorff distances between sets respectively:

$$\begin{aligned} d(S,T) &= \inf_{\substack{s\in S,\\ t\in T}} |s-t|, \\ \tilde{d}_H(S,T) &= \sup_{s\in S} d(s,T), \\ d_H(S,T) &= \max\{\sup_{s\in S} d(s,T), \sup_{t\in T} d(S,t)\}. \end{aligned}$$

Finally, we use $\text{Unif}(\mathcal{S}^{n-1})$ to denote the uniform distribution over the complex n-dimensional unit sphere and $\mathcal{N}_C(0,1)$ to denote standard complex normal (the real and imaginary parts are i.i.d. from $\mathcal{N}(0,1/2)$).

1.3.2 Eigenvalue estimates as a polynomial problem

For $A \in \mathbb{C}^{n\times n}$, the numerical range of A can be defined as

$$W(A) := \left\{ \frac{\boldsymbol{v}^* A \boldsymbol{v}}{\boldsymbol{v}^* \boldsymbol{v}} \text{ for } \boldsymbol{v} \in \mathbb{C}^n \right\},$$

noting that when A is normal, $W(A) = \text{conv}(\Lambda(A))$.

When we project A onto $\mathcal{K}_m(A, \boldsymbol{b})$, we can think of it as restricting our possible vectors v such that $v \in \mathcal{K}_m(A, \boldsymbol{b}) \subset \mathbb{C}^n$. Then for iteration m, then numerical range of H_m is simply

$$W(H_m) := \left\{ \frac{\boldsymbol{v}^* A \boldsymbol{v}}{\boldsymbol{v}^* \boldsymbol{v}} \text{ for } \boldsymbol{v} \in \mathcal{K}_m(A, \boldsymbol{b}) \right\}.$$

We want to say that the numerical ranges are close if for any $\lambda \in \Lambda(A)$, there exists some point in $W(H_m)$ close to λ, i.e. $d_H(W(A), W(H_m))$ is small. It is clear that $W(H_m) \subseteq W(A)$, which means it is enough to show that any point in $W(A)$, specifically the extremal eigenvalues of A, are close to some point in $W(H_m)$.

A useful property of Krylov subspaces is that any vector $v \in \mathcal{K}_m(A, \boldsymbol{b})$ can be rewritten with a polynomial $p \in \mathcal{P}_m$ such that $v = p(A)\boldsymbol{b}$. Hence, we have

$$W(H_m) = \left\{ \frac{[p(A)\boldsymbol{b}]^* A [p(A)\boldsymbol{b}]}{[p(A)\boldsymbol{b}]^* [p(A)\boldsymbol{b}]} \text{ for } p \in \mathcal{P}_m \right\},$$

and for any extremal eigenvalue λ, we want to minimize

$$\min_{p \in \mathcal{P}_m} \left| \lambda - \frac{[p(A)\boldsymbol{b}]^* A [p(A)\boldsymbol{b}]}{[p(A)\boldsymbol{b}]^* [p(A)\boldsymbol{b}]} \right|.$$

At this point, it becomes clear that we can shift/scale the eigenvalues of A without loss of generality since we can arbitrarily shift/scale the polynomial coefficients to compensate. This fact will be used in Chapter 2, where we assume the desired extreme eigenvalue is zero and that the rest of the spectrum is contained in the half-disk $R_0 := \{z \in \mathbb{C} : |z| \leq 0, \arg z \in [\pi/2, 3\pi/2]\}$. Our goal will be to find polynomials which minimize the quantity above.

1.3.3 Randomization and numerical range estimation

As it stands, it is possible for an adversarial starting vector $\boldsymbol{b}$ to prevent $W(H_m)$ from converging on the desired eigenvalue λ. In particular, this problem occurs if $\boldsymbol{b}$ starts orthogonal or extremely close to orthogonal to the desired eigenvector. Luckily, it is well known that this pitfall can be avoided by introducing randomization on $\boldsymbol{b}$. When $\boldsymbol{b}$ is chosen uniformly randomly from the unit sphere, its component in any fixed direction is concentrated around $1/n$, making it highly unlikely for the component to be too small.

Thus, error bounds in this work and previous literature ([8], [7]) bound the probability that the error is small given $\boldsymbol{b} \sim \text{Unif}(\mathbb{S}^{n-1})$ such that the probability approaches 1 as m increases. Because the failure probability vanishes, we are also able to make statements about estimating many extremal eigenvalues at once, allowing a single matrix H_m to approximate the entire numerical range.

In fact, a common method of estimating the numerical range (see [9]) is to project the numerical range in a direction of the complex plane by computing the extremal eigenpairs of $A_\theta = \frac{1}{2}(e^{i\theta}A + e^{-i\theta}A)$. Crucially, A_θ is Hermitian, which means it is more efficient to compute its eigenvalues (and using the Krylov methods discussed in this work, one can achieve much better convergence than the general normal case). However, each A_θ only reveals two points, which means many different angles θ and corresponding matrices must be analyzed to get a picture of the entire numerical range. The advantage of our methods here is that only a single or very few matrices H_m needs to be considered.

Chapter 2

Upper Bounds on Approximation Error

In this section, we upper bound the approximation error of estimating any extreme eigenvalue using the mth Krylov subspace with arbitrary starting vector $\boldsymbol{b}$. Beyond an interesting theoretical observation about a popular algorithm, this bound also has practical implications for users looking to estimate the computation required to approximate extremal eigenvalues to a given degree of accuracy.

In Section 2.1 we go over the polynomials we will use to minimize the error. Then in Section 2.2, we will prove an upper bound on approximation for any single extreme eigenvalue. Finally in Section 2.3, we consider randomized starting vectors and approximate the entire numerical range at once.

2.1 Polynomial bounds and Remez-type inequalities

Remez-type inequalities are useful for choosing polynomials which produce large gaps, which will help us separate the eigenvalues. As discussed previously, we will assume without loss of generality that the desired extreme eigenvalue is zero and $W(A) \subseteq \{z \in \mathbb{C} : |z| \leq 0, \arg z \in [\pi/2, 3\pi/2]\}$.

Let us pick a parameter δ describing the distance at which eigenvalues are no longer "close enough". We want a polynomial which is small over the region of all possible far eigenvalues

$$R_\delta = \{z : |z| \in [\delta, 1], \arg z \in [\pi/2, 3\pi/2]\}.$$

To do so, we consider the following polynomial bound constructed by Erdelyi, Li, and Saff based on Chebshev polynomials [10].

Proposition 2.1 ([10, proof of Theorem 2.6])**.** *For any $\varepsilon > 0$ and $m \in \mathbb{N}$, we define the region*

$$D_\varepsilon := \{z \in \mathbb{C} : |z| \leq 1, \arg z \in [\varepsilon, \pi - \varepsilon] \cup [\pi + \varepsilon, 2\pi - \varepsilon]\} \cup \{z \in \mathbb{C} : |z| \leq 1 - \varepsilon/8\}.$$

Then, there exists $Q_{m,\epsilon} \in \mathcal{P}_{3m}$ such that

$$Q_{m,\epsilon}(1) \geq \exp\left(\frac{m\epsilon}{8}\right) \max_{z \in D_\epsilon} |Q_{m,\epsilon}(z)|.$$

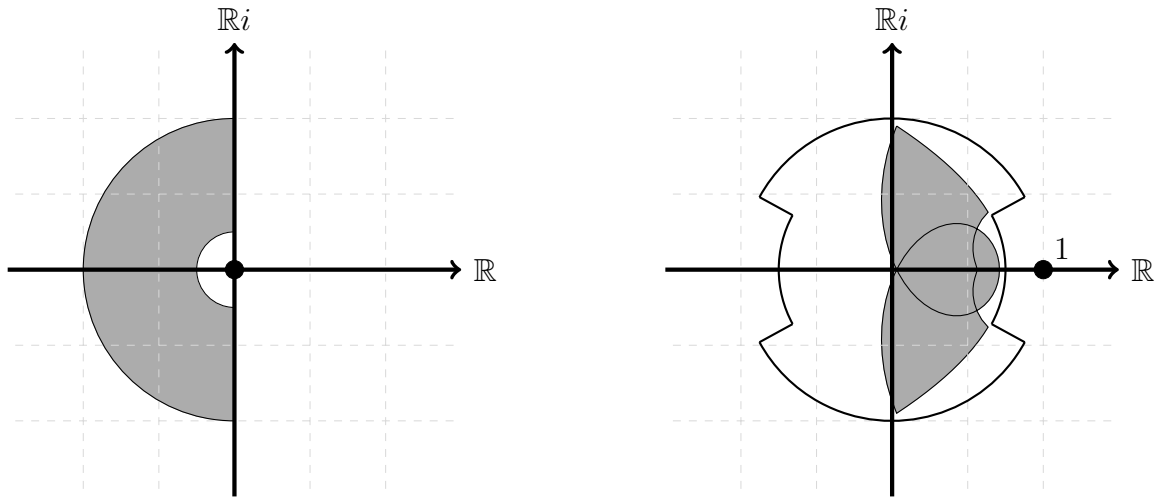

(a) Region of possible eigenvalues $R_\delta = \{z : |z| \in [\delta, 1], \arg z \in [\pi/2, 3\pi/2]\}$

(b) Mapping into D_ε via $z \mapsto (1-\delta/4)z^2 + (1-\delta/8)z + 1$

Figure 2.1: The shaded region represents the region of possible eigenvalues at least δ away from $\lambda = 0$ before and after the mapping $P(z) = (1-\delta/4)z^2 + (1-\delta/8)z + 1$. The outlined shape in Subfigure (b) is the bounded region D_ε. For illustrative purposes, certain numbers in the figure have been fudged.

As stated, the previous bound is not small over R_δ. However, we can use the function $P(z) = (1-\delta/4)z^2 + (1-\delta/8)z + 1$ to map R_δ into D_ε when $\varepsilon = \frac{2}{3}\delta$, ensuring that the desired eigenvalue 0 is mapped to 1. Composing the polynomial from Proposition 2.1 and P together produces the following bound:

Lemma 2.2. *For any $0 \le \delta < 1$ and $m \in \mathbb{N}$, there exists a polynomial $\widetilde{P}_{m,\delta} \in \mathcal{P}_{6m}$ such that*

$$\widetilde{P}_{m,\delta}(0) \ge \exp\left(\frac{m\delta}{12}\right) \max_{z \in R_\delta} |\widetilde{P}_{m,\delta}(z)|. \tag{2.1}$$

The computation showing that P properly maps R_δ into D_ε is rather tedious and unenlightening, so we refer the reader to Appendix A of [3] for the calculations. See Figure 2.1 for an illustration of the mapping.

2.2 A Deterministic Upper Bound for a Single Extremal Eigenvalue

Given the polynomial bounds from our previous section, we can now bound the approximation error for any fixed starting vector $\boldsymbol{b}$.

Theorem 2.3. *Given $m, n \in \mathbb{N}$, normal $A \in \mathbb{C}^{n\times n}$, and a starting vector $b \in \mathbb{C}^n$, let $\boldsymbol{\varphi}$ be a unit eigenvector with extreme eigenvalue λ. Then,*

$$d\left(\lambda, W(H_{6m+1})\right) \le \frac{6}{m} \ln\left(\frac{em\|\boldsymbol{b}\|_2^2}{6|\langle \boldsymbol{b}, \boldsymbol{\varphi}\rangle|^2}\right) \operatorname{diam}(W(A)).$$

Proof. Recall that we may shift and scale A such that $\lambda = 0$ and $\Lambda(A) \subseteq R_0$. For each eigenvalue λ_i with corresponding unit eigenvector $\boldsymbol{\varphi}_i$, let $\boldsymbol{\alpha_i} = \boldsymbol{\varphi}_i^* \boldsymbol{b}$. Then,

$$d(\lambda, W(H_{6m+1})) = \min_{\boldsymbol{x} \in \mathcal{K}_{6m+1}(A,\boldsymbol{b})} \left| \frac{\boldsymbol{x}^* A \boldsymbol{x}}{\boldsymbol{x}^* \boldsymbol{x}} \right| = \min_{p \in \mathcal{P}_{6m}} \left| \frac{[p(A)\boldsymbol{b}]^* A [p(A)\boldsymbol{b}]}{[p(A)\boldsymbol{b}]^* [p(A)\boldsymbol{b}]} \right| = \min_{p \in \mathcal{P}_{6m}} \left| \frac{\sum_{j=1}^n \lambda_j |p(\lambda_i)^2| |\boldsymbol{\alpha}_j|^2|}{\sum_{j=1}^n |p(\lambda_i)^2| |\boldsymbol{\alpha}_j|^2|} \right|.$$

Let δ control the distance at which we consider eigenvalues close or far from λ, and let $S_\delta = \{j : |\lambda_j| \geq \delta\}$ be the set of indices of far eigenvalues. Then, we can break into two terms

$$\begin{aligned} \min_{x \in \mathcal{K}_{6m+1}(A,b)} \left| \frac{\boldsymbol{x}^* A \boldsymbol{x}}{\boldsymbol{x}^* \boldsymbol{x}} \right| &\leq \min_{p \in \mathcal{P}_{6m}} \left| \frac{\sum_{j \in S_\delta} |\boldsymbol{\alpha}_j|^2 |p(\lambda_j)|^2 \lambda_j}{\sum_{j=1}^n |\boldsymbol{\alpha}_j|^2 |p(\lambda_j)|^2} \right| + \left| \frac{\sum_{j \notin S_\delta} |\boldsymbol{\alpha}_j|^2 |p(\lambda_j)|^2 \lambda_j}{\sum_{j=1}^n |\boldsymbol{\alpha}_j|^2 |p(\lambda_j)|^2} \right| \\ &\leq \min_{p \in \mathcal{P}_{6m}} \frac{\sum_{j \in S_\delta} |\boldsymbol{\alpha}_j|^2 |p(\lambda_j)|^2}{|\langle \boldsymbol{b}, \boldsymbol{\varphi} \rangle|^2 |p(0)|^2} + \delta. \end{aligned}$$

Clearly, the term containing close eigenvalues is bounded by δ, leaving us to handle the farther eigenvalues with the polynomial $\widetilde{P}_{m,\delta}$ from Lemma 2.2. Hence,

$$\min_{p \in \mathcal{P}_{6m}} \frac{\sum_{j \in S_\delta} |\boldsymbol{\alpha}_j|^2 |p(\lambda_j)|^2}{|\langle \boldsymbol{b}, \boldsymbol{\varphi} \rangle|^2 |p(0)|^2} \leq \frac{\sum_{j \in S_\delta} |\boldsymbol{\alpha}_j|^2 |\widetilde{P}_{m,\delta}(\lambda_j)|^2}{|\langle \boldsymbol{b}, \boldsymbol{\varphi} \rangle|^2 |\widetilde{P}_{m,\delta}(0)|^2} \leq \frac{\|\boldsymbol{b}\|_2^2}{|\langle \boldsymbol{b}, \boldsymbol{\varphi} \rangle|^2} \exp\left(-\frac{m\delta}{6}\right).$$

Finally, let $\delta = 6m^{-1} \ln\left(\frac{m\|\boldsymbol{b}\|_2^2}{6|\langle \boldsymbol{b}, \boldsymbol{\varphi} \rangle|^2}\right)$ to obtain the desired bound

$$\frac{\|\boldsymbol{b}\|_2^2}{|\langle \boldsymbol{b}, \boldsymbol{\varphi} \rangle|^2} \exp\left(-\frac{m\delta}{6}\right) + \delta \leq \frac{6}{m}\left(1 + \ln\left(\frac{m\|\boldsymbol{b}\|_2^2}{6|\langle \boldsymbol{b}, \boldsymbol{\varphi} \rangle|^2}\right)\right).$$

□

2.3 A Probabilistic Upper Bound for the Numerical Range

In this section, we will bound the error of estimating the numerical range with high probability when $\boldsymbol{b}$ is chosen uniformly from the unit sphere.

The key idea we use here is that a random direction is rarely close to orthogonal to any fixed eigenvector $\boldsymbol{\varphi}$. This is a standard result in high dimensional probability, which we state here.

Proposition 2.4. *Let* $\boldsymbol{b} \sim \mathrm{Unif}(\mathbb{S}^{n-1})$, $\boldsymbol{\varphi} \in \mathbb{C}^n$ *with* $\|\boldsymbol{\varphi}\| = 1$, *and* $t \in (0, 1)$. *Then*

$$\mathbb{P}\left[\frac{\|\boldsymbol{b}\|_2^2}{|\langle \boldsymbol{b}, \boldsymbol{\varphi} \rangle|^2} \leq \frac{n}{t}\right] \geq 1 - et.$$

Proof. Follows from [11, Lemma 2.2]. Note that, due to complex values, $k = 2$ in our case. □

Directly applying Proposition 2.4 to Theorem 2.3 gives us a probabilistic bound for estimating a single eigenvalue. Taking this a step further, we union bound over all extreme eigenvalues to obtain the following theorem:

Theorem 2.5. *Given $m, n \in \mathbb{N}$, $\alpha > 0$, normal $A \in \mathbb{C}^{n\times n}$, and $\boldsymbol{b} \sim \mathrm{Unif}(\mathbb{S}^{n-1})$, then*

$$\mathbb{P}\left[\frac{d_H\big(W(H_{6m+1}), W(A)\big)}{\mathrm{diam}(W(A))} \le \frac{6(2+\alpha)\ln n}{m}\right] \ge 1 - \frac{5m}{4n^\alpha}.$$

Proof. Let $\{\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_n\}$ be an orthonormal eigenbasis of A. Then we can take the maximum error over all extremal eigenvalues, applying Theorem 2.3 to obtain

$$d_H\big(W(H_{6m+1}), W(A)\big) \le \frac{6}{m} \ln\left(\frac{em\|\boldsymbol{b}\|_2^2}{6\min_{j\in[n]} |\langle \boldsymbol{b}, \boldsymbol{\varphi}_j\rangle|^2}\right) \mathrm{diam}(W(A)).$$

Then, applying Proposition 2.4 and taking a union bound produces

$$\mathbb{P}\left[\frac{\|\boldsymbol{b}\|_2^2}{\min_{j\in[n]} |\langle \boldsymbol{b}, \boldsymbol{\varphi}\rangle|^2} \le \frac{n}{t}\right] \ge 1 - ent,$$

and choosing $t = em/(6n^{\alpha+1})$ gives the desired bound. □

Chapter 3

Lower Bounds

In the last chapter, we proved upper bounds on the approximation error after m iterations. What we want to know now is whether these bounds are actually tight. In this chapter, we will construct a sequence of normal matrices which have an approximation error of $\Omega(1/m)$ (see Figure 3.1 for an illustration). We will first prove a lower bound for these matrices using a fixed starting vector and then make slight modifications to prove a similar bound for randomized starting vectors.

We define the first sequence of matrices here.

Definition 3.1. *For $m \in \mathbb{N}$, let $n = m^2$ and $\omega = \exp\left(\frac{2\pi i}{m}\right)$. Then let $\boldsymbol{r} = \frac{1}{\sqrt{m}}(\sqrt{1}, \sqrt{2}, \ldots, \sqrt{m})^T \in \mathbb{C}^m$ and $\boldsymbol{\theta} = (\omega, \omega^2, \ldots, \omega^m)^T \in \mathbb{C}^m$, and we define $\widehat{A}_m = \operatorname{diag}(\boldsymbol{r} \otimes \boldsymbol{\theta}) \in \mathbb{C}^{n \times n}$.*

With this definition of $\widehat{A}_m$, we are able to compute its approximation after m iterations, H_m, explicitly to bound its numerical range.

Lemma 3.2. *Given $m \geq 15$, $\widehat{A}_m \in \mathbb{C}^{n \times n}$ from Definition 3.1, and starting vector $\boldsymbol{b} = (1, \ldots, 1)^T \in \mathbb{C}^n$, let $\widehat{H}_m$ be the projection of $\widehat{A}_m$ onto $\mathcal{K}_m(\widehat{A}_m, \boldsymbol{b})$. Then,*

$$d(W(\widehat{H}_m), \partial W(\widehat{A}_m)) \geq \frac{1}{14m} \operatorname{diam}(W(\widehat{A}_m))$$

where ∂ denotes the boundary of the set.

Proof. It is simple to construct $\widehat{H}_m$ explicitly since $\boldsymbol{b}, \widehat{A}_m\boldsymbol{b}, \ldots, (\widehat{A}_m)^m\boldsymbol{b}$ are orthogonal. In fact, we have

$$\begin{aligned}
\langle \hat{A}^j \tilde{\boldsymbol{b}}, \hat{A}^k \tilde{\boldsymbol{b}} \rangle &= \sum_{p=1}^{m} \sum_{q=1}^{m} \boldsymbol{\theta}_p^j \boldsymbol{r}_q^j \boldsymbol{\theta}_p^{-k} \boldsymbol{r}_q^k \\
&= \sum_{p=1}^{m} e^{2\pi i (j-k)\frac{p}{m}} \sum_{q=1}^{m} \left(\frac{q}{m}\right)^{\frac{j+k}{2}} \\
&= \begin{cases} m \sum_{q=1}^{m} \left(\frac{q}{m}\right)^j & \text{if } j = k \\ 0 & \text{otherwise .} \end{cases}
\end{aligned}$$

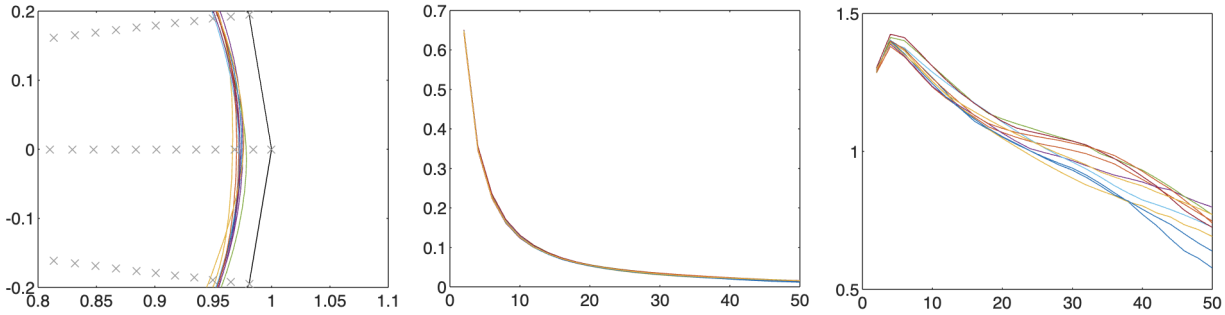

(a) $W(\tilde{A})$, $\Lambda(\tilde{A})$ and $W(\tilde{H}_{30})$ for ten different $\boldsymbol{b} \sim \text{Unif}(\mathbb{S}^{n-1})$

(b) $\tilde{d}_H(W(\tilde{A}), W(\tilde{H}_m))$ for $m = 1, \ldots, 50$ and ten different $\boldsymbol{b} \sim \text{Unif}(\mathbb{S}^{n-1})$

(c) $m \times \tilde{d}_H(W(\tilde{A}), W(\tilde{H}_m))$ for $m = 1, \ldots, 50$ and ten different $\boldsymbol{b} \sim \text{Unif}(\mathbb{S}^{n-1})$

Figure 3.1: Numerical range estimates of $\tilde{A}$ as in Definition 3.3. Whereas n changes as a function of m to simplify the proof, the essence of the statement holds for each matrix individually. Here we fix $\tilde{A} \in \mathbb{C}^{n \times n}$ with $n = 10240$ with eigenvalue multiplicity $\ell = 10$ and run the Arnoldi iteration for m up to 50. In Subfigure (a), we plot part of the numerical range of A along with the estimated numerical range for 10 random starting vectors $\boldsymbol{b}$ and $m = 50$. The eigenvalues themselves are shown by gray $\times$s. In Subfigure (b), we plot the Hausdorff distance between the estimated and actual numerical ranges as m increases. To better see the $1/m$ behavior respectively, we multiply the error by m in Subfigure (c).

Stepping through the Arnoldi iteration, we see most entries of $\widehat{H}_m$ are 0. Specifically by entry, we have

$$[\widehat{H}_m]_{jk} = \begin{cases} \frac{\|\widehat{A}_m^k \boldsymbol{b}\|}{\|\widehat{A}_m^{k-1} \boldsymbol{b}\|} & \text{if } j = k+1 \\ 0 & \text{otherwise .} \end{cases}$$

With some arithmetic, we can show for $m \geq 15$ that

$$\begin{aligned} \frac{\|\widehat{A}_m^k \boldsymbol{b}\|}{\|\widehat{A}_m^{k-1} \boldsymbol{b}\|} &= \left(\frac{\sum_{q=1}^m q^j}{m \sum_{q=1}^m q^{j-1}} \right)^{1/2} = \left(1 - \frac{m \sum_{q=1}^m q^{j-1} - \sum_{q=1}^m q^j}{m \sum_{q=1}^m q^{j-1}} \right)^{1/2} \\ &= \left(1 - \frac{\sum_{q=1}^m (m-q) q^{j-1}}{m \sum_{q=1}^m q^{j-1}} \right)^{1/2} \geq \left(1 - \frac{1}{m} \cdot \frac{(m-1)^{j-1}}{m^{j-1} + (m-1)^{j-1}} \right)^{1/2} \\ &\geq \left(1 - \frac{1}{m} \cdot \frac{1}{(1 + \frac{1}{m-1})^{j-1} + 1} \right)^{1/2} \geq \left(1 - \frac{1}{m(e+1)} \right)^{1/2} \geq 1 - \frac{1}{7m}. \end{aligned}$$

Then Gershgorin's Circle Theorem immediately implies that the eigenvalues of $\widehat{H}_m$ are contained in the closed disc of radius $(1 - \frac{1}{7m})$ centered at zero. However, by construction, the boundary of $W(\widehat{A}_m)$ lies at distance 1, implying an error of at least $\frac{1}{7m}$. □

While the previous construction provides a lower bound for a fixed starting vector $\boldsymbol{b} = (1, \ldots, 1)^T$, we really want to handle the case for $\boldsymbol{b} \sim \text{Unif}(\mathbb{S}^{n-1})$. The key idea for the following proof is to increase the multiplicity of all eigenvalues such that $\boldsymbol{b}$ looks uniform

with high probability. Then, we bound the effect of perturbing $\boldsymbol{b}$ on the resulting matrix $\widehat{H}_m$. We define our new matrix here:

Definition 3.3. *For $m \in \mathbb{N}$, let $\ell \geq 4800m^2 \ln m$, $n = \ell m^2$, and $\omega = \exp\left(\frac{2\pi i}{m}\right)$. Then let $\mathbf{1} = (1, \ldots, 1)^T \in \mathbb{C}^\ell$, $\boldsymbol{r} = \frac{1}{\sqrt{m}}(\sqrt{1}, \sqrt{2}, \ldots, \sqrt{m})^T \in \mathbb{C}^m$ and $\boldsymbol{\theta} = (\omega, \omega^2, \ldots, \omega^m)^T \in \mathbb{C}^m$, and we define $\tilde{A}_m = \operatorname{diag}(\boldsymbol{r} \otimes \boldsymbol{\theta} \otimes \mathbf{1}) \in \mathbb{C}^{n \times n}$.*

The following proposition gives a standard concentration bound, which will be used to control the sums of entries of $\boldsymbol{b}$.

Proposition 3.4 ([12, Example 2.11])**.** *Let $X \sim \chi^2_k$ be a chi-squared random variable with k degrees of freedom. Then $\mathbb{P}[|X - k| \geq kt] \leq 2e^{-kt^2/8}$ for all $t \in (0, 1)$.*

Then, we will use the following proposition to argue that the projections produced by a random starting vector $\boldsymbol{b}$ are close to the case with uniform starting vector $\mathbf{1}$.

Proposition 3.5. *Let $V \in \mathbb{C}^{m \times n}$ satisfy $V^*V = I$ with $m \geq n$, $D \in \mathbb{C}^{m \times m}$ be diagonal with $\|D\|_2 < 1$, and QH be the polar decomposition of $(I + D)V$. Then*

$$\|V^*AV - Q^*AQ\|_2 \leq \frac{4\|A\|_2\|D\|_2}{1 - \|D\|_2}$$

for all $A \in \mathbb{C}^{m \times m}$.

Proof. Because $\|D\|_2 < 1$, the matrix $(I + D)V$ has full column rank, and so

$$Q = (I + D)V[V^*(I + D^*)(I + D)V]^{-1/2} \in \mathbb{C}^{m \times n}$$

in the polar decomposition of $(I + D)V$ is unique [13, Theorem 7.3.1]. Letting $E = V - Q$,

$$\begin{aligned}\|V^*AV - Q^*AQ\|_2 &= \|E^*AV + Q^*AE\|_2 \\ &\leq \|E\|_2\|A\|_2\|V\|_2 + \|Q\|_2\|A\|_2\|E\|_2 \\ &= 2\|E\|_2\|A\|_2.\end{aligned}$$

What remains is to bound $\|E\|_2$. Using a similar technique,

$$\begin{aligned}\|E\|_2 &= \| - DV + (I + D)V\big(I - [V^*(I + D^*)(I + D)V]^{-1/2}\big)\|_2 \\ &\leq \|D\|_2 + (1 + \|D\|_2)\|I - [V^*(I + D^*)(I + D)V]^{-1/2}\|_2.\end{aligned}$$

The matrix $[V^*(I+D^*)(I+D)V]$ has singular values in the interval $[(1-\|D\|_2)^2, (1+\|D\|_2)^2]$, and so

$$\|I - [V^*(I + D^*)(I + D)V]^{-1/2}\|_2 \leq \frac{1}{1 - \|D\|_2} - 1 = \frac{\|D\|_2}{1 - \|D\|_2}.$$

Combining all these bounds, we obtain

$$\|V^*AV - Q^*AQ\|_2 \leq 2\|A\|_2\left(\|D\|_2 + \frac{(1 + \|D\|_2)\|D\|_2}{1 - \|D\|_2}\right) = \frac{4\|A\|_2\|D\|_2}{1 - \|D\|_2}.$$

□

Putting all this together, we reach the main theorem of the chapter.

Theorem 3.6. *Given $m \geq 15$, $\tilde{A}_m \in \mathbb{C}^{n\times n}$ from Definition 3.3, and $\tilde{\boldsymbol{b}} \sim \mathrm{Unif}(\mathbb{S}^{n-1})$, let $\tilde{H}_m$ be the projection of $\tilde{A}_m$ onto $\mathcal{K}_m(\tilde{A}_m, \tilde{\boldsymbol{b}})$. Then,*

$$\mathbb{P}\left[\frac{d\big(W(\tilde{H}_m), \partial W(\tilde{A}_m)\big)}{\mathrm{diam}(W(\tilde{A}_m))} \geq \frac{1}{63m}\right] \geq 1 - \frac{2}{m}.$$

Proof. Without loss of generality, we can scale $\tilde{\boldsymbol{b}}$ such that the entries of $\tilde{\boldsymbol{b}}$ are unit complex normal: $\tilde{\boldsymbol{b}}_i \sim \mathcal{N}_C(0,1)$. With the increased multiplicity for each eigenvalue by a factor of ℓ, we can group the entries of $\tilde{\boldsymbol{b}}$ which multiply against the same eigenvalue together. Then for $\mathbf{1} = (1,\ldots,1)^T \in \mathbb{C}^{m^2}$, we see that $\tilde{\boldsymbol{b}}^* \tilde{A}_m \tilde{\boldsymbol{b}}$ behaves just like $\mathbf{1}^* \widehat{A}_m \mathbf{1}$, except that each eigenvalue λ_j is given weight $\sum_{k=1}^{\ell} |\tilde{\boldsymbol{b}}_{\ell j+k}|^2 \sim \frac{1}{2}\chi^2_{2\ell}$ instead of 1. (Note that we obtain 2ℓ degrees of freedom due to $\tilde{\boldsymbol{b}}$ being complex valued). Hence, we can equivalently consider the numerical range of $\widehat{A}_m$ with a different distribution of starting vector.

Let $\boldsymbol{b} \in \mathbb{C}^{m^2}$ have entries $\boldsymbol{b}_j = \sqrt{\frac{1}{\ell}\sum_{k=1}^{\ell} |\tilde{\boldsymbol{b}}_{\ell j+k}|^2}$. Then $\mathbb{E}[\boldsymbol{b}_j] = 1$ and applying Proposition 3.4 with $t = \sqrt{12\ell^{-1}\ln m} < 1$ and union bounding over m^2 events gives us

$$\mathbb{P}\left[\max_{j=1,\ldots,m^2} |\boldsymbol{b}_j^2 - 1| \geq \sqrt{\frac{12\ell \ln m}{\ell}}\right] \leq \frac{2}{m}.$$

Setting $\ell = 4800m^2 \ln m$ then gives us with probability at least $1 - 2/m$,

$$\max_{j=1,\ldots,m^2} \boldsymbol{b}_j \leq 1 - \sqrt{1 - \frac{1}{20m}}.$$

Furthermore, we can think of the span of $\mathcal{K}_m(\widehat{A}_m, \boldsymbol{b})$ in terms of $\widehat{Q}_m = [\mathbf{1}, \widehat{A}_m\mathbf{1}, \ldots, (\widehat{A}_m)^m\mathbf{1}] \in \mathbb{C}^{m^2\times m^2}$ where $[\boldsymbol{b}, \widehat{A}_m\boldsymbol{b}, \ldots, (\widehat{A}_m)^m\boldsymbol{b}] = (I+D)\widehat{Q}_m$ and D is diagonal with $D_{jj} = \widehat{\boldsymbol{b}}_j - 1$. To produce a projection, we will consider the polar decomposition QH of $(I+D)\widehat{Q}_m$ and argue that the polar Q is close to $\widehat{Q}_m$.

Let H_m be the projection of $\widehat{A}_m$ onto $\mathcal{K}_m(\widehat{A}_m, \boldsymbol{b})$ and $\widehat{H}_m$ be the projection of $\widehat{A}_m$ onto $\mathcal{K}_m(\widehat{A}_m, \mathbf{1})$. By Proposition 3.5,

$$\|\widehat{H}_m - H_m\|_2 = \|\widehat{Q}_m^* \widehat{A}_m \widehat{Q}_m - Q^* \widehat{A}_m Q\|_2 \leq \frac{4\left(1 - \sqrt{1 - \frac{1}{20m}}\right)}{\sqrt{1 - \frac{1}{20m}}}.$$

By inspection for $m \geq 15$, the right-hand side is again upper bounded by $\frac{1}{9m}$. Combining this result with Lemma 3.2 gives us the desired bound. With probability at least $1 - 2/m$,

$$\max|W(H_m)| \leq \max|W(\widehat{H}_m)| + \frac{1}{9m} \leq 1 - \frac{2}{63m},$$

which means the distance to any of the eigenvalues of modulus one is at least $\frac{2}{63m}$. □

Remark 3.7. *While the statement and proof of Theorem 3.6 construct different matrices $\tilde{A}_m$ for each number of iteration m, we can actually observe the same $\Omega(1/m)$ behavior with a fixed matrix $\tilde{A}$ and increasing the iterations m. See Figure 3.1, which illustrates this idea.*

Chapter 4

Error Bounds for Non-Normal Matrices

While we have restricted ourselves to normal matrices in the previous sections, we will now consider a parameterization of non-normality that produces similar, albeit weaker, convergence bounds. In Section 4.1 we discuss why non-normality is an issue. Then in Section 4.2 we prove several probabilistic statements that will be used to prove Lemma 4.7 and the main theorem 4.11 in Section 4.3.

4.1 The Problem of Non-Normality

Before we approach our error bounds for non-normal matrices, we will first discuss why non-normality is a significant concern for numerical range estimation. In the normal case, the numerical range is the convex hull of eigenvalues. This fact is due to the orthogonality of eigenvectors, preventing eigenvectors from combining in any direction. When we lift the restriction on normality, even matrices with constant eigenvalue condition number lead to a breakdown in estimation for the numerical range (see Section 4 of [3]).

As a result, the bounds in this section will not deal with estimating the numerical range, but rather the convex hull of eigenvalues. We can see the separation between the numerical range and convex hull of eigenvalues in the following example, with $W(H_m)$ providing a much better approximation of $\text{conv}(\Lambda(A))$.

Example 4.1. *Let $A = V\Lambda V^{-1} \in \mathbb{C}^{n\times n}$, with $V^*V = (1 - n^{-2/3})I + n^{-2/3}\mathbf{1}\mathbf{1}^T$ and Λ as defined in Theorem 3.6. Figure 4.1 demonstrates that $\text{conv}(\Lambda(A))$ can be approximated with order $1/m$ convergence, as the condition $|[V^*V]_{jk}| \leq n^{-2/3}$ prevents the eigenvectors from clustering.*

4.2 Fun Probability Facts

Here we state and prove several probabilistic inequalities, which will be useful for the proofs later in this chapter. Proposition 4.3 and Corollary 4.4 are essentially the Hanson-Wright inequality for complex matrices (with constants).

Proposition 4.2. *Let $\boldsymbol{b} \sim \mathcal{N}_{\mathcal{C}}(0, I_n)$ and $M \in \mathbb{C}^{n\times n}$ be a normal matrix. Then, $\mathbb{E}[\boldsymbol{b}^*M\boldsymbol{b}] = \text{trace}(M)$ and $\text{Var}[\boldsymbol{b}^*M\boldsymbol{b}] = \|M\|_F^2$.*

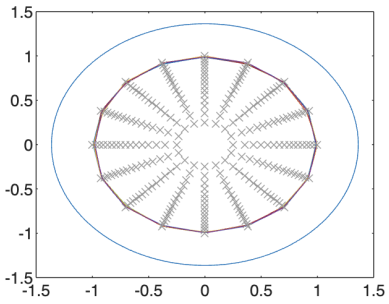

(a) $W(A)$, $\Lambda(A)$ and $W(H_{30})$ for ten different $\boldsymbol{b} \sim \text{Unif}(\mathbb{S}^{n-1})$

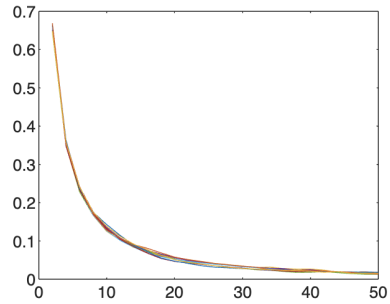

(b) $\tilde{d}_H(\text{conv}(\Lambda(A)), W(H_m))$ for $m = 1, \ldots, 50$ and ten different $\boldsymbol{b} \sim \text{Unif}(\mathbb{S}^{n-1})$

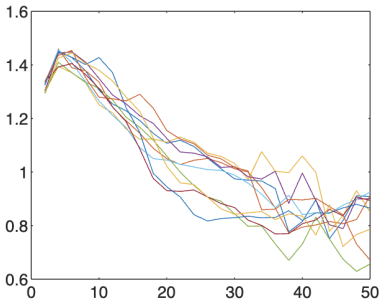

(c) $\tilde{d}_H(\text{conv}(\Lambda(A)), W(H_m))\times m$ for $m = 1, \ldots, 50$ and ten different $\boldsymbol{b} \sim \text{Unif}(\mathbb{S}^{n-1})$

Figure 4.1: Example 4.1. Here we let $n = 2560$, $m = 16$, and $\ell = 10$. In Subfigure (a), the numerical range is shown by the outer blue circle while eigenvalues are plotted as gray crosses. Note how $W(H_m)$ closely approximates the convex hull of eigenvalues. Subfigures (b) and (c) show the approximation error, and (c) demonstrates the $O(1/m)$ behavior by multiplying by m.

Proof. We provide a proof only because of the involvement of complex random variables; see [14, Lemma 3.2b.2] for well-known results for the real symmetric case. When M is normal, the invariance of the distribution of $\boldsymbol{b}$ under rotation allows us to rewrite $\boldsymbol{b}^* M \boldsymbol{b} \overset{d}{=} \sum_{j=1}^n \lambda_j |\boldsymbol{b}_j|^2$, where $\{\lambda_j\}_{j=1}^n$ are the eigenvalues of M. For $\boldsymbol{b}_j \sim \mathcal{N}_{\mathcal{C}}(0,1)$, we have $\mathbb{E}[|\boldsymbol{b}_j|^2] = 1$ and $\text{Var}[|\boldsymbol{b}_j|^2] = 1$, and so

$$\mathbb{E}\left[\sum_{j=1}^n \lambda_j |\boldsymbol{b}_j|^2\right] = \sum_{j=1}^n \lambda_j \mathbb{E}\left[|\boldsymbol{b}_j|^2\right] = \sum_{j=1}^n \lambda_j$$

and

$$\text{Var}\left[\sum_{j=1}^n \lambda_j |\boldsymbol{b}_j|^2\right] = \sum_{j=1}^n |\lambda_j|^2 \text{Var}[|\boldsymbol{b}_j|^2] = \sum_{j=1}^n |\lambda_j|^2.$$

□

Proposition 4.3 ([12, Example 2.8, Proposition 2.9])**.** *Let $\boldsymbol{b} \sim \mathcal{N}_{\mathcal{C}}(0, I_n)$ and $M \in \mathbb{C}^{n\times n}$ be a Hermitian or skew-Hermitian matrix. Let $X = \boldsymbol{b}^* M \boldsymbol{b}$. Then,*

$$\mathbb{P}\left[|X - \mathbb{E}[X]| \geq t\right] \leq 2\exp\left[-\min\left(\frac{t^2}{4\|M\|_F^2}, \frac{t}{4\|M\|_2}\right)\right].$$

Proof. Again rewrite $\boldsymbol{b}^* M \boldsymbol{b} \overset{d}{=} \sum_{j=1}^n \lambda_j |\boldsymbol{b}_j|^2$ where λ_j are the eigenvalues of M. Since M is Hermitian or skew-Hermitian, the eigenvalues are purely real or imaginary. Each term $|\lambda_j||\boldsymbol{b}_j|^2$ is sub-exponential with parameters $(\sqrt{2}|\lambda_j|, 2|\lambda_j|)$; thus, $\sum_{j=1}^n |\lambda_j||\boldsymbol{b}_j|^2$ is sub-exponential with parameters $(\sqrt{2}\|M\|_F, 2\|M\|_2)$. □

Corollary 4.4. *Let $M \in \mathbb{C}^{n\times n}$, $M_1 = \frac{1}{2}(M+M^*)$, and $M_2 = \frac{1}{2}(M-M^*)$, with corresponding random variables $X = \boldsymbol{b}^*M\boldsymbol{b}$, $X_1 = \boldsymbol{b}^*M_1\boldsymbol{b}$, and $X_2 = \boldsymbol{b}^*M_2\boldsymbol{b}$. Then*

$$\mathrm{Var}[X] = \|M_1\|_F^2 + \|M_2\|_F^2 \text{ and } \quad \mathbb{P}\left[|X - \mathbb{E}[X]| \geq t\right] \leq 4\exp\left[-\min\left(\frac{t^2}{8\mathrm{Var}[X]}, \frac{t}{4\sqrt{2}\|M\|_2}\right)\right].$$

Proof. Note that $X = X_1 + X_2$, X_1 is purely real, and X_2 is purely imaginary. For the first part, rewrite

$$\mathrm{Var}[X] = \mathrm{Var}[\mathrm{Re}(X)] + \mathrm{Var}[\mathrm{Im}(X)] = \mathrm{Var}[X_1] + \mathrm{Var}[X_2].$$

For the second, note that

$$\mathbb{P}\left[|X - \mathbb{E}[X]| \geq t\right] \leq \mathbb{P}\left[|X_1 - \mathbb{E}[X_1]| \geq t/\sqrt{2}\right] + \mathbb{P}\left[|X_2 - \mathbb{E}[X_2]| \geq t/\sqrt{2}\right].$$

The given bound is not tight but suffices for our purposes. □

Proposition 4.5. *Let $\boldsymbol{b} \sim \mathcal{N}_{\mathbb{C}}(0, I_n)$ be a random vector and $M \in \mathbb{C}^{n\times n}$ be a positive-definite Hermitian matrix. Then, $\mathbb{P}\left[\boldsymbol{b}^*M\boldsymbol{b} \leq t \cdot \mathrm{trace}(M)\right] \leq et$ for all $t > 0$.*

Proof. The proof is a standard Chernoff bound adapted from [15] for complex normals. Again rewrite $\boldsymbol{b}^*M\boldsymbol{b} \overset{d}{=} \sum_{j=1}^n \lambda_j|\boldsymbol{b}_j|^2$ where λ_j are the eigenvalues of M. Assume $t < 1$, otherwise the statement is trivially true. Then, for $c = (t^{-1} - 1)/\mathrm{trace}(M)$,

$$\begin{aligned}
\mathbb{P}\left[\boldsymbol{b}^*M\boldsymbol{b} \leq t \cdot \mathrm{trace}(M)\right] &= \mathbb{P}\left[t\sum_{j=1}^n \lambda_j - \sum_{j=1}^n \lambda_j|\boldsymbol{b}_j|^2 \geq 0\right] \\
&= \mathbb{P}\left[\prod_{j=1}^n e^{c\lambda_j(t-|\boldsymbol{b}_j|^2)} \geq 1\right] \\
&\leq \prod_{j=1}^n \mathbb{E}\left[e^{c\lambda_j(t-|\boldsymbol{b}_j|^2)}\right] \\
&= \prod_{j=1}^n \frac{e^{c\lambda_j t}}{1+c\lambda_j} \\
&\leq e^{c\cdot\mathrm{trace}(M)t}(1 + c\cdot\mathrm{trace}(M))^{-1} = e^{1-t}t.
\end{aligned}$$

□

4.3 Estimates for Non-Normal Matrices with Repulsive Eigenvectors

Now we arrive at the main purpose of this chapter. While estimating the numerical range breaks down quickly, it is possible to estimate the convex hull of eigenvalues well when the eigenvectors tend to stay away from each other. The following definition parameterizes our notion of eigenvector repulsion, preventing eigenvectors from adversarially combining in any single direction.

Definition 4.6. *We say a matrix $V \in \mathbb{C}^{n\times n}$ is β-normal for some $\beta > 0$ if its columns have norm one and, for $V = [\boldsymbol{v}_1 \ \dots \ \boldsymbol{v}_n]$ and $(V^{-1})^* = [\boldsymbol{w}_1 \ \dots \ \boldsymbol{w}_n]$,*

$$\sum_{\ell\in[n]\setminus k} |\boldsymbol{v}_k^*\boldsymbol{v}_\ell|(|\boldsymbol{w}_\ell^*\boldsymbol{w}_j| + |\boldsymbol{w}_k^*\boldsymbol{w}_j|) \le n^{-\beta}\|\boldsymbol{w}_j\|^2 \qquad \text{for all} \quad j,k \in [n]. \tag{4.1}$$

When this β-normal condition holds, we can effectively treat the additional difficulty of the non-orthogonal eigenbasis, which appears in two places. Let $A = V\Lambda V^{-1}$ where Λ is diagonal and V has columns of norm one. The issues become clear if we expand out the Rayleigh quotients for $x \in \mathcal{K}_m(A, \boldsymbol{b})$,

$$\frac{\boldsymbol{x}^*A\boldsymbol{x}}{\boldsymbol{x}^*\boldsymbol{x}} = \frac{[p(A)\boldsymbol{b}]^*A[p(A)\boldsymbol{b}]}{[p(A)\boldsymbol{b}]^*[p(A)\boldsymbol{b}]} = \frac{[p(\Lambda)V^{-1}\boldsymbol{b}]^*V^*V\Lambda[p(\Lambda)V^{-1}\boldsymbol{b}]}{[p(\Lambda)V^{-1}\boldsymbol{b}]^*V^*V[p(\Lambda)V^{-1}\boldsymbol{b}]}.$$

First, we have the problem of the V^*V term in the middle, which will be dealt with in the following lemma. Then, we will consider the second problem, which is the distribution of $V^{-1}\boldsymbol{b}$ compared to $\boldsymbol{b}$. Once we handle these issues, the resulting Rayleigh quotient looks the same as the normal case, and we can apply our theorems from previous chapters.

Lemma 4.7. *Let $n \ge 2$, $\beta > 2\ln(2\ln(en))/\ln(n)$, $V \in \mathbb{C}^{n\times n}$ be β-normal (see Definition 4.6), and $D, P \in \mathbb{C}^{n\times n}$ be diagonal with $\|D\|_2 \le 1$. Then, for $\boldsymbol{b} \sim \mathrm{Unif}(\mathcal{S}^{n-1}(\mathbb{C}))$ and $\varepsilon = 2n^{-\beta/2}\ln(en)$,*

$$\left|\frac{[PV^{-1}\boldsymbol{b}]^*V^*VD[PV^{-1}\boldsymbol{b}]}{[PV^{-1}\boldsymbol{b}]^*V^*V[PV^{-1}\boldsymbol{b}]}\right| \ge \left(\frac{1}{1-\varepsilon}\right)\left(\left|\frac{[PV^{-1}\boldsymbol{b}]^*D[PV^{-1}\boldsymbol{b}]}{[PV^{-1}\boldsymbol{b}]^*[PV^{-1}\boldsymbol{b}]}\right| + \varepsilon\right)$$

with probability at most

$$en^{-\beta/2} + 2n^{-1/2} + 4n^{-1/4}.$$

Proof. The method for this proof is rather brute force. We will compute the norms of the difference between the numerators and denominators and bound them using concentration inequalities.

For simplicity, let $W = (V^{-1})^*$, and, without loss of generality, suppose that $|P_j| \le |P_k|$ for $j < k$. We first treat the denominator $\|VPW^*\boldsymbol{b}\|^2$; the numerator will be bounded in a similar manner. For the desired denominator $\|PW^*\boldsymbol{b}\|^2$, let $\phi := \mathbb{E}[\|PW^*\boldsymbol{b}\|^2]$ denote its expectation. Then, by Proposition 4.2,

$$\phi = \mathbb{E}[\|PW^*\boldsymbol{b}\|^2] = \mathrm{trace}(WP^*PW^*) = \mathrm{trace}(P^*PW^*W) = \sum_{j=1}^n |P_j|^2\|\boldsymbol{w}_j\|^2.$$

The expectation of $\|VPW^*\boldsymbol{b}\|^2$ is given by

$$\mathbb{E}[\|VPW^*\boldsymbol{b}\|^2] = \mathrm{trace}(WP^*V^*VPW^*) = \mathrm{trace}(V^*VPW^*WP^*) = \sum_{j=1}^n\sum_{k=1}^n \boldsymbol{v}_j^*\boldsymbol{v}_kP_k\boldsymbol{w}_k^*\boldsymbol{w}_j\overline{P_j}.$$

By the ordering $|P_j| \le |P_k|$ for $j < k$ and Inequality 4.1 (with $j = k$),

$$\begin{aligned}
\mathbb{E}[\|VPW^*\boldsymbol{b}\|^2] &= \sum_{j=1}^{n} |P_j|^2 \|\boldsymbol{w}_j\|^2 + 2\sum_{j=1}^{n}\sum_{k=1}^{j-1} \mathrm{Re}(P_k\overline{P_j}\,\boldsymbol{v}_j^*\boldsymbol{v}_k\,\boldsymbol{w}_k^*\boldsymbol{w}_j) \\
&\ge \sum_{j=1}^{n} |P_j|^2 \|\boldsymbol{w}_j\|^2 - 2\sum_{j=1}^{n} |P_j|^2 \sum_{k=1}^{j-1} |\boldsymbol{v}_j^*\boldsymbol{v}_k||\boldsymbol{w}_k^*\boldsymbol{w}_j| \\
&\ge (1 - 2n^{-\beta})\phi.
\end{aligned}$$

Now, let $M = WP^*(V^*V - I)PW^*$, and consider the error term $\|VPW^*\boldsymbol{b}\|^2 - \|PW^*\boldsymbol{b}\|^2 = \boldsymbol{b}^*M\boldsymbol{b}$. By Proposition 4.2, its variance is given by

$$\begin{aligned}
\|M\|_F^2 &= \mathrm{trace}(([WP^*][V^*V - I][PW^*])^2) \\
&= \mathrm{trace}(([V^*V - I][PW^*][WP^*])^2) \\
&= \sum_{j=1}^{n}\sum_{k=1}^{n} \left(\sum_{\ell\in[n]\setminus j} \boldsymbol{v}_j^*\boldsymbol{v}_\ell\, P_\ell\overline{P_k}\,\boldsymbol{w}_\ell^*\boldsymbol{w}_k \right)\left(\sum_{m\in[n]\setminus k} \boldsymbol{v}_k^*\boldsymbol{v}_m\, P_m\overline{P_j}\,\boldsymbol{w}_m^*\boldsymbol{w}_j \right).
\end{aligned}$$

Rearranging the index pairs j, ℓ and k, m, and applying Inequality 4.1, we obtain

$$\begin{aligned}
\|M\|_F^2 &= \sum_{j=1}^{n}\sum_{k=1}^{n} \left(\sum_{\ell=1}^{j-1} \boldsymbol{v}_j^*\boldsymbol{v}_\ell\, P_\ell\overline{P_k}\,\boldsymbol{w}_\ell^*\boldsymbol{w}_k + \boldsymbol{v}_\ell^*\boldsymbol{v}_j\, P_j\overline{P_k}\,\boldsymbol{w}_j^*\boldsymbol{w}_k \right) \\
&\qquad \cdot \left(\sum_{m=1}^{k-1} \boldsymbol{v}_k^*\boldsymbol{v}_m\, P_m\overline{P_j}\,\boldsymbol{w}_m^*\boldsymbol{w}_j + \boldsymbol{v}_m^*\boldsymbol{v}_k\, P_k\overline{P_j}\,\boldsymbol{w}_k^*\boldsymbol{w}_j \right) \\
&\le \sum_{j=1}^{n}\sum_{k=1}^{n} |P_j|^2|P_k|^2 \left(\sum_{\ell=1}^{j-1} |\boldsymbol{v}_j^*\boldsymbol{v}_\ell| \big(|\boldsymbol{w}_\ell^*\boldsymbol{w}_k| + |\boldsymbol{w}_j^*\boldsymbol{w}_k|\big) \right)\left(\sum_{m=1}^{k-1} |\boldsymbol{v}_k^*\boldsymbol{v}_m| \big(|\boldsymbol{w}_m^*\boldsymbol{w}_j| + |\boldsymbol{w}_k^*\boldsymbol{w}_j|\big) \right) \\
&\le n^{-2\beta} \sum_{j=1}^{n}\sum_{k=1}^{n} |P_j|^2|P_k|^2 \|\boldsymbol{w}_k\|^2\|\boldsymbol{w}_j\|^2 \\
&\le n^{-2\beta}\phi^2.
\end{aligned}$$

Now, let $X = \|PW^*\boldsymbol{b}\|^2$, $Y = \|VPW^*\boldsymbol{b}\|^2$, and $t \ge 1$. By Proposition 4.3 (using the inequality $\|M\|_2 \le \|M\|_F$) and Proposition 4.5,

$$\begin{aligned}
\mathbb{P}\left[Y \le (1 - 2n^{-\beta/2}t)X\right] &= \mathbb{P}\left[X - Y \ge 2n^{-\beta/2}tX\right] \\
&\le \mathbb{P}\left[X \le n^{-\beta/2}\phi\right] + \mathbb{P}\left[X - Y \ge 2n^{-\beta}t\phi\right] \\
&\le en^{-\beta/2} + 2\exp\left(-\tfrac{t-1}{2}\right).
\end{aligned}$$

Now we turn to the numerator. Assume, without loss of generality, that $\|P\|_2 \le 1$, and let $Z = [PV^{-1}\boldsymbol{b}]^*(V^*V - I)D[PV^{-1}\boldsymbol{b}]$ equal the difference in numerators. Its expectation is given by

$$\mathbb{E}[Z] = \mathrm{trace}(DPW^*WP^*(V^*V - I)) = \sum_{j=1}^{n} D_jP_j \sum_{k\in[n]\setminus j} \boldsymbol{w}_j^*\boldsymbol{w}_k\,\overline{P_k}\,\boldsymbol{v}_k^*\boldsymbol{v}_j.$$

Using Inequality 4.1 and the property $|P_j| \le |P_k|$ for $j < k$, $|\mathbb{E}[Z]|$ can be bounded as follows:

$$\left|\mathbb{E}[Z]\right| \le \sum_{j=1}^{n} \sum_{k\in[n]\setminus j} |P_j||P_k||\boldsymbol{w}_j^*\boldsymbol{w}_k|\,|\boldsymbol{v}_k^*\boldsymbol{v}_j| \le 2\sum_{j=1}^{n} |P_j|^2 \sum_{k=1}^{j-1} |\boldsymbol{w}_j^*\boldsymbol{w}_k|\,|\boldsymbol{v}_k^*\boldsymbol{v}_j| \le 2n^{-\beta}\phi.$$

Bounding the variance is slightly more difficult, since $N := [PV^{-1}]^*(V^*V - I)D[PV^{-1}]$ is not necessarily normal. Instead, we bound $\|N + N^*\|_F^2$ and $\|N - N^*\|_F^2$ separately and make use of Corollary 4.4. The procedure is very similar to that of the denominator, save for the term $\max_{j,k} |D_j + \overline{D_k}| \le 2$. We sketch the estimate for $\|N + N^*\|_F^2$ ($\|N - N^*\|_F^2$ is nearly identical):

$$\begin{aligned}
&\|N + N^*\|_F^2 = \operatorname{trace}\left(([D^*(V^*V - I) + (V^*V - I)D]PW^*WP^*)^2\right)\\
&= \sum_{j=1}^{n}\sum_{k=1}^{n}\left(\sum_{\ell\in[n]\setminus j} (\overline{D_j} + D_\ell)\boldsymbol{v}_j^*\boldsymbol{v}_\ell\, P_\ell\overline{P_k}\,\boldsymbol{w}_\ell^*\boldsymbol{w}_k\right)\left(\sum_{m\in[n]\setminus k} (\overline{D_k} + D_m)\boldsymbol{v}_k^*\boldsymbol{v}_m\, P_m\overline{P_j}\,\boldsymbol{w}_m^*\boldsymbol{w}_j\right)\\
&\le 4\sum_{j=1}^{n}\sum_{k=1}^{n}\left(\sum_{\ell\in[n]\setminus j} |P_\ell||\overline{P_k}|\,|\boldsymbol{v}_j^*\boldsymbol{v}_\ell|\,|\boldsymbol{w}_\ell^*\boldsymbol{w}_k|\right)\left(\sum_{m\in[n]\setminus k} |P_m||\overline{P_j}|\,|\boldsymbol{v}_k^*\boldsymbol{v}_m|\,|\boldsymbol{w}_m^*\boldsymbol{w}_j|\right)\\
&\le 4\sum_{j=1}^{n}\sum_{k=1}^{n}|P_j|^2|P_k|^2\left(\sum_{\ell=1}^{j-1}|\boldsymbol{v}_j^*\boldsymbol{v}_\ell|\left(|\boldsymbol{w}_\ell^*\boldsymbol{w}_k| + |\boldsymbol{w}_j^*\boldsymbol{w}_k|\right)\right)\left(\sum_{m=1}^{k-1}|\boldsymbol{v}_k^*\boldsymbol{v}_m|\left(|\boldsymbol{w}_m^*\boldsymbol{w}_j| + |\boldsymbol{w}_k^*\boldsymbol{w}_j|\right)\right)\\
&\le 4n^{-2\beta}\sum_{j=1}^{n}\sum_{k=1}^{n}|P_j|^2|P_k|^2\|\boldsymbol{w}_k\|^2\|\boldsymbol{w}_j\|^2\\
&\le 4n^{-2\beta}\phi^2.
\end{aligned}$$

Hence, $\operatorname{Var}[Z] \le 2n^{-2\beta}\phi^2$ and, by Corollary 4.4,

$$\mathbb{P}\left[|Z| \ge 2n^{-\beta}t\phi\right] \le \mathbb{P}\left[|Z - \mathbb{E}[Z]| \ge 2n^{-\beta}(t-1)\phi\right] \le 4\exp\left(-\tfrac{t-1}{4}\right).$$

To complete the proof, we union bound the cases that the denominator is too small and the numerator too big. Let $Z_x = [PV^{-1}\boldsymbol{b}]^*D[PV^{-1}\boldsymbol{b}]$ and $Z_y = [PV^{-1}\boldsymbol{b}]^*V^*VD[PV^{-1}\boldsymbol{b}]$. Setting $t = 1 + \ln n$ (note that $2n^{-\beta/2}t < 1$ for $\beta > 2\ln(2\ln(en))/\ln(n)$), we have

$$\begin{aligned}
&\mathbb{P}\left[\frac{|Z_y|}{Y} \ge \frac{1}{1 - 2n^{-\beta/2}t}\left(\frac{|Z_x|}{X} + \frac{2t}{n^{\beta/2}}\right)\right]\\
&\le \mathbb{P}\left[Y \le (1 - 2n^{-\beta/2}t)X\right] + \mathbb{P}\left[|Z_x| - |Z_y| \ge 2n^{-\beta/2}tX\right]\\
&\le \mathbb{P}\left[X \le n^{-\beta/2}\phi\right] + \mathbb{P}\left[X - Y \ge 2n^{-\beta}t\phi\right] + \mathbb{P}\left[|Z| \ge 2n^{-\beta}t\phi\right]\\
&\le en^{-\beta/2} + 2\exp\left(-\tfrac{t-1}{2}\right) + 4\exp\left(-\tfrac{t-1}{4}\right)\\
&= en^{-\beta/2} + 2n^{-1/2} + 4n^{-1/4}.
\end{aligned}$$

□

The above lemma immediately implies an upper bound for estimating extremal eigenvalues.

Corollary 4.8. *Let $n \geq 2$ and $\beta > 2\ln(2\ln(en))/\ln(n)$, $A \in \mathbb{C}^{n\times n}$ be diagonalizable with eigendecomposition $A = V\Lambda V^{-1}$, where V is β-normal (see Definition 4.6), and $\Lambda_k \in \Lambda(A)$ be an extreme point of* $\operatorname{conv}(\Lambda(A))$. *Then, for $\boldsymbol{b} \sim \operatorname{Unif}(\mathbb{S}^{n-1})$,*

$$\mathbb{P}\left[\frac{d(\Lambda_k, W(H_{6m+1}))}{\operatorname{diam}(\operatorname{conv}(\Lambda(A)))} \geq \frac{\frac{6}{m}\ln\left(\frac{em\|V^{-1}\boldsymbol{b}\|_2^2}{6|[V^{-1}\boldsymbol{b}]_k|^2}\right) + \frac{2\ln(en)}{n^{\beta/2}}}{1 - \frac{2\ln(en)}{n^{\beta/2}}}\right] \leq en^{-\beta/2} + 2n^{-1/2} + 4n^{-1/4}.$$

Proof. By Lemma 2.2, there exists a polynomial $p \in \mathcal{P}_{6m}$ such that

$$\left|\Lambda_k - \frac{[p(\Lambda)V^{-1}\boldsymbol{b}]^*\Lambda[p(\Lambda)V^{-1}\boldsymbol{b}]}{[p(\Lambda)V^{-1}\boldsymbol{b}]^*[p(\Lambda)V^{-1}\boldsymbol{b}]}\right| \leq \frac{6}{m}\ln\left(\frac{em\|V^{-1}\boldsymbol{b}\|_2^2}{6|[V^{-1}\boldsymbol{b}]_k|^2}\right)\operatorname{diam}(\operatorname{conv}(\Lambda(A))).$$

Applying Lemma 4.7 with $D = (\Lambda_k I - \Lambda)$ and $P = p(\Lambda)$ completes the proof. □

Now, similar to Chapter 2, we want to take the bound for a single extremal eigenvalue and consider the entire convex hull of eigenvalues as well as randomizing the starting vector. Since we have weaker concentration of variables due to the eigenbasis, we are not actually able to guarantee good estimates for all eigenvalues at once. Instead, we estimate a fraction of the eigenvalues, using the following proposition for estimating convex bodies to show that a good choice of eigenvalues exists.

Proposition 4.9 ([16, Section 4],[17])**.** *Let $\mathcal{U}$ be a convex body in $\mathbb{C}$ with boundary length L. There exist n points S on the boundary of $\mathcal{U}$ such that the Hausdorff distance*

$$d_H(\mathcal{U}, \operatorname{conv}(S)) \leq \frac{L}{2n}\tan\frac{\pi}{n}.$$

We will also use the following proposition to handle the distribution of $V^{-1}\boldsymbol{b}$,

Proposition 4.10. *Let $M \in \mathbb{C}^{n\times n}$ be invertible, $\boldsymbol{b} \sim \operatorname{Unif}(\mathbb{S}^{n-1})$, and $t \in (0,1)$. Then*

$$\mathbb{P}\left[\min_{j\in[n]}\frac{|[M\boldsymbol{b}]_j|^2}{\|M\boldsymbol{b}\|_2^2} \geq \frac{t}{n^2\kappa^2(M)}\right] \geq 1 - et.$$

Proof. By [11, Lemma 2.2], $\mathbb{P}\left[|\langle \boldsymbol{b}, \boldsymbol{\nu}\rangle|^2 \leq \frac{\tau}{n}\right] \leq e\tau$ for $\tau \in [0,1)$ and $\|\boldsymbol{\nu}\|_2 = 1$. We have

$$\mathbb{P}\left[|\langle M\boldsymbol{b}, \boldsymbol{\nu}\rangle|^2 \leq \frac{t}{n^2\|M^{-1}\|_2^2}\right] \leq \mathbb{P}\left[|\langle \boldsymbol{b}, M^*\boldsymbol{\nu}\rangle|^2 \leq \frac{t\|M^*\boldsymbol{\nu}\|_2^2}{n^2}\right] \leq \frac{et}{n}.$$

Applying a union bound for $\boldsymbol{e}_j$, $j \in [n]$, and noting that $\|M\boldsymbol{b}\|_2^2 \leq \|M\|_2^2$ implies our desired result. □

Finally, we reach our main theorem:

Theorem 4.11. *Let $n \geq 2$ and $\beta > 2\ln(2\ln(en))/\ln(n)$, $A \in \mathbb{C}^{n\times n}$ be diagonalizable with eigendecomposition $A = V\Lambda V^{-1}$, where V is β-normal (see Definition 4.6). Then, for*

$\boldsymbol{b} \sim \mathrm{Unif}(\mathbb{S}^{n-1})$ *and any* $0 < \gamma < \min\{\beta/2, 1/4\}$*, the one-sided Hausdorff distance between* $\mathrm{conv}(\Lambda(A))$ *and* $W(H_{6m+1})$ *is bounded by*

$$\frac{\tilde{d}_H(\mathrm{conv}(\Lambda(A)), W(H_{6m+1}))}{\mathrm{diam}(\mathrm{conv}(\Lambda(A)))} \leq \frac{\frac{6}{m}\ln(n^{2+\alpha}\kappa^2(V)) + \frac{2\ln(en)}{n^{\beta/2}}}{1 - \frac{2\ln(en)}{n^{\beta/2}}} + \frac{\pi}{n^\gamma}\tan\frac{2\pi}{n^\gamma}$$

with probability at least

$$1 - \frac{e^2 m}{6n^\alpha} - \frac{e}{n^{\beta/2-\gamma}} - \frac{2}{n^{1/2-\gamma}} - \frac{4}{n^{1/4-\gamma}}.$$

Proof. Let $\mathcal{U} = \mathrm{conv}(\Lambda(A))$, and $L = \pi\,\mathrm{diam}(\mathcal{U})$. By Proposition 4.9, we can pick $n^\gamma/2$ vertices S to create an interior polytope $\mathcal{V} = \mathrm{conv}(S)$ approximating the convex hull with error at most $\frac{L}{n^\gamma}\tan\frac{2\pi}{n^\gamma}$. The optimal polytope $\mathcal{V}$ must have each element of S on the boundary of $\mathcal{U}$. For each $s \in S$, let $\lambda_s, \mu_s \in \Lambda(A)$ be the endpoints of the boundary segment containing s. Let $W = \bigcup_{s\in S}\{\lambda_s, \mu_s\}$ and $\mathcal{W} = \mathrm{conv}(W)$. Since $\mathcal{V} \subseteq \mathcal{W} \subseteq \mathcal{U}$, $\mathcal{W}$ approximates $\mathcal{U}$ at least as well as $\mathcal{V}$.

By Corollary 4.8 applied to the eigenvalues in W,

$$\mathbb{P}\left[\frac{\tilde{d}_H(\mathcal{U}, W(H_{6m+1}))}{\mathrm{diam}(\mathrm{conv}(\Lambda(A)))} \geq \frac{\frac{6}{m}\ln\left(\frac{em\|V^{-1}\boldsymbol{b}\|_2^2}{6\min_k|[V^{-1}\boldsymbol{b}]_k|^2}\right) + \frac{2\ln(en)}{n^{\beta/2}}}{1 - \frac{2\ln(en)}{n^{\beta/2}}} + \frac{\pi}{n^\gamma}\tan\frac{2\pi}{n^\gamma}\right]$$

$$\leq \frac{e}{n^{\beta/2-\gamma}} + \frac{2}{n^{1/2-\gamma}} + \frac{4}{n^{1/4-\gamma}}.$$

By Proposition 4.10 with $M = V^{-1}$ and $t = em/(6n^\alpha)$, we have

$$\mathbb{P}\left[\frac{\min_{k\in[n]}|[V^{-1}\boldsymbol{b}]_k|^2}{\|V^{-1}\boldsymbol{b}\|_2} \leq \frac{em}{6n^{2+\alpha}\kappa^2(V)}\right] \leq \frac{e^2 m}{6n^\alpha},$$

completing the proof. □

References

[1] Y. Saad. *Overview of Krylov subspace methods with applications to control problems*. Tech. rep. Research Inst. for Advanced Computer Science Moffett Field, CA, United States, May 1989.

[2] W. Mück and Y. Yang. "Krylov complexity and orthogonal polynomials". In: *Nuclear Physics B* 984 (2022), p. 115948. ISSN: 0550-3213. DOI: https://doi.org/10.1016/j.nuclphysb.2022.115948. URL: https://www.sciencedirect.com/science/article/pii/S0550321322002991.

[3] C. Chen and J. Urschel. *Estimating the numerical range with a Krylov subspace*. 2024. arXiv: 2411.19165 [math.NA]. URL: https://arxiv.org/abs/2411.19165.

[4] Y. Saad. *Numerical Methods for Large Eigenvalue Problems: Revised Edition*. Society for Industrial and Applied Mathematics, 2011, pp. 125–162. DOI: 10.1137/1.9781611970739.ch6. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611970739.ch6. URL: https://epubs.siam.org/doi/abs/10.1137/1.9781611970739.ch6.

[5] C. Musco and C. Musco. "Randomized Block Krylov Methods for Stronger and Faster Approximate Singular Value Decomposition". In: *Advances in Neural Information Processing Systems*. Ed. by C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett. Vol. 28. Curran Associates, Inc., 2015. URL: https://proceedings.neurips.cc/paper_files/paper/2015/file/1efa39bcaec6f3900149160693694536-Paper.pdf.

[6] J. Kuczyński and H. Woźniakowski. "Estimating the Largest Eigenvalue by the Power and Lanczos Algorithms with a Random Start". In: *SIAM Journal on Matrix Analysis and Applications* 13.4 (1992), pp. 1094–1122. DOI: 10.1137/0613066. eprint: https://doi.org/10.1137/0613066. URL: https://doi.org/10.1137/0613066.

[7] J. C. Urschel. "Uniform Error Estimates for the Lanczos Method". In: *SIAM Journal on Matrix Analysis and Applications* 42.3 (2021), pp. 1423–1450. DOI: 10.1137/20M1331470. eprint: https://doi.org/10.1137/20M1331470. URL: https://doi.org/10.1137/20M1331470.

[8] J. Kuczyński and H. Woźniakowski. "Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start". In: *SIAM journal on matrix analysis and applications* 13.4 (1992), pp. 1094–1122.

[9] N. J. Higham. *The matrix computation toolbox*.

[10] T. Erdélyi, X. Li, and E. Saff. "Remez-and Nikolskii-type inequalities for logarithmic potentials". In: *SIAM Journal on Mathematical Analysis* 25.2 (1994), pp. 365–383.

[11] S. Dasgupta and A. Gupta. "An elementary proof of a theorem of Johnson and Lindenstrauss". In: *Random Structures & Algorithms* 22.1 (2003), pp. 60–65.

[12] M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Vol. 48. Cambridge university press, 2019.

[13] R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge university press, 2012.

[14] A. Mathai and S. Provost. "Quadratic Forms in Random Variables". In: (Jan. 1992).

[15] G. L. (https://mathoverflow.net/users/1004/george-lowther). *Anti-concentration of Gaussian quadratic form*. MathOverflow. URL:https://mathoverflow.net/q/95108 (version: 2012-04-25). eprint: https://mathoverflow.net/q/95108. URL: https://mathoverflow.net/q/95108.

[16] E. M. Bronstein. "Approximation of convex sets by polytopes". In: *Journal of Mathematical Sciences* 153.6 (2008), pp. 727–762. DOI: 10.1007/s10958-008-9144-x. URL: https://doi.org/10.1007/s10958-008-9144-x.

[17] V. A. Popov. "Approximation of convex sets". In: *Bull. Inst. Math. Bulgar. Acad. Sci. I* 1 (1970), pp. 67–79.