

Robust subspace designs and the power of a unique small quantum witness

Simon Apers¹, Roman Edenhofer¹, Benjamin Mathieu-Bloise¹, and Partha Mukhopadhyay²

¹Université Paris Cité, CNRS, IRIF, Paris, France

²Chennai Mathematical Institute, Chennai, India

October 5, 2026

Abstract

The concept of subspace designs was introduced by Guruswami and Xing (STOC'13), and explicit constructions were given by Guruswami and Kopparty (FOCS'13). These are families of subspaces that have small intersection with any given subspace of a fixed dimension. We introduce *robust subspace designs*. Informally, these are a quantitative extension in which we demand that not too many subspaces of the family contain directions that lie *close* to any given subspace of a fixed dimension. We give a probabilistic construction of such a robust subspace design of polynomial size, as well as a non-trivial explicit construction of superpolynomial size.

Our main application of this new concept is a quantum space-bounded variant of the Valiant-Vazirani theorem (Theor. Comput. Sci.'86), which shows that restricting NP-complete problems to instances with at most one accepting witness preserves hardness under randomized reductions. For quantum witnesses, the analogous quantity is the dimension of an accepting witness subspace. We use our probabilistic construction of robust subspace designs to isolate a unique witness for space-bounded quantum Merlin-Arthur protocols with perfect completeness and an acceptance gap outside their perfectly accepting subspace.

As a further application, we give a randomized reduction of well-conditioned nullity testing to space-bounded quantum Merlin-Arthur protocols with perfect completeness. Using a similar idea, we find that ordinary subspace designs allow us to recover the classical $C=L$ containment of Allender, Beals, and Ogihara (STOC'96) for general nullity testing through a simpler proof.

Contents

1	Introduction and summary	1
1.1	Robust subspace designs	1
1.2	Isolating a unique small quantum witness	2
1.3	Further applications	4
1.4	Probabilistic construction of a robust subspace design	4
1.5	Explicit construction of a robust subspace design	5
1.6	Open problems and future directions	6
1.7	Paper overview	7
2	Preliminaries	7
3	Robust subspace designs and kernel peeling	8
3.1	Definitions	9
3.2	Kernel peeling	10
3.3	Explicit non-robust construction based on Wronskian evaluation maps	13
3.4	Subspace designs need not be robust	13
4	Probabilistic robust construction	14
4.1	The Grassmannian manifold of subspaces	14
4.2	Random Gaussian matrix theory	15
4.3	The Gaussian construction	15
4.4	Computability in randomized logspace	17
5	Explicit robust constructions	20
5.1	Explicit construction based on coordinate projections	20
5.2	Improved explicit construction using sketching	20
5.2.1	Robust sketching family	21
5.3	Composing families	23
6	Isolating a unique small quantum witness and other applications	25
6.1	Quantum space-bounded complexity classes	25
6.2	Isolating a unique small quantum witness	28
6.3	Other applications	29
6.3.1	Classical upper bounds for nullity testing and feasibility	30
6.3.2	Quantum upper bounds for nullity testing and feasibility	32
7	Acknowledgements	33
A	Explicit construction based on Wronskian evaluation maps	37
B	Proofs of Gaussian matrix concentration	39
B.1	Background: Laguerre unitary ensemble and determinantal point processes	39
B.2	Singular value concentration	40
C	Proofs for complexity-theoretic applications	41
C.1	Constant pc-NULL reduces to pc-SINGULAR	41
C.2	Completeness of well-conditioned singularity	43

1 Introduction and summary

1.1 Robust subspace designs

Subspace designs were introduced by Guruswami and Xing [GX13], who also gave a probabilistic polynomial-time construction. These are families of subspaces that have limited intersections with every subspace of a fixed dimension. Originally, the notion was defined over finite fields, mainly motivated by applications in coding theory. However, subspace designs can be defined over any field [BCDZ26].

In the weak version, a family of subspaces $\mathcal{H} = \{H_1, \dots, H_m\}$ is a (k, Δ) -weak subspace design if, for every k -dimensional subspace W ,

$$|\{i \in [m] : H_i \cap W \neq \{0\}\}| \leq \Delta.$$

The family is a (k, Δ) -strong subspace design if, for every k -dimensional subspace W ,

$$\sum_{i=1}^m \dim(H_i \cap W) \leq \Delta.$$

Guruswami and Xing [GX13] used subspace designs mainly in connection with list decoding. Subsequently, Guruswami and Kopparty explicitly constructed subspace designs with near-optimal parameters [GK16]. Since then, subspace designs have become useful objects in coding theory and algebraic pseudorandomness [FG15]. In a recent breakthrough, subspace designs were indirectly used to obtain the first NC algorithm for bipartite matching [CGG⁺26]. Soon after that, they were also used to show that $\text{NL} \subseteq \oplus\text{L}$ [App26].

In this work, we formulate and study a robust extension of this notion over the field of complex numbers $\mathbb{C}$. For a linear map $Q : \mathbb{C}^n \rightarrow \mathbb{C}^k$ and a subspace W of dimension k , it is based on the smallest singular value of Q restricted to W ,

$$\sigma_{\min}(Q|_W).$$

If Q is a co-isometry, then this is the sine of the smallest principal angle between $\ker(Q)$ and W , or alternatively the minimum distance between any unit vector from W to $\ker(Q)$. Evidently, $\sigma_{\min}(Q|_W) = 0$ if and only if $\ker(Q) \cap W \neq \{0\}$.

Now fix a robustness parameter $\eta \in [0, 1]$. We say that a family of linear maps $\mathcal{Q} = \{Q_1, \dots, Q_m\}$ from $\mathbb{C}^n$ to $\mathbb{C}^k$, with bounded operator norms $\|Q_i\|_{\text{op}} \leq 1$, is a (k, η, Δ) -weak robust subspace design if, for every k -dimensional subspace W ,

$$|\{i \in [m] : \sigma_{\min}(Q_i|_W) \leq \eta\}| \leq \Delta.$$

Similarly, the family $\mathcal{Q}$ is a (k, η, Δ) -strong robust subspace design if, for every k -dimensional subspace W ,

$$\sum_{i=1}^m |\{j \in [k] : \sigma_j(Q_i|_W) \leq \eta\}| \leq \Delta.$$

It is straightforward to observe that if $\mathcal{Q}$ is a (k, η, Δ) -weak/strong robust subspace design, then the family of subspaces

$$\mathcal{H} = \{\ker(Q) : Q \in \mathcal{Q}\}$$

is an ordinary (k, Δ) -weak/strong subspace design. In some sense, the converse is not true: for every $\eta = \eta(n) \in (0, 1]$, there is a $(1, n - 1)$ subspace design whose robustness parameter is smaller than η . Conceptually, this is illustrated in Figure 1.




Figure 1: A schematic robust and non-robust family of planes in $\mathbb{R}^3$. No direction can be too close to all three coordinate planes on the left, whereas the direction w is simultaneously close to all three planes on the right.

Alongside applications of robust subspace designs in space-bounded complexity, we prove the following:

1. We use random matrix theory to argue that, with constant probability, a polynomial-sized ensemble of Gaussian matrices forms a strong subspace design with inverse polynomial robustness. This implies a randomized-logspace-constructible robust subspace design.
2. We combine ideas from low-rank matrix approximation (specifically, sketching and max-volume submatrices) to give a deterministic space-efficient weak robust subspace design with inverse polynomial robustness of size $O(n^{O(1)}k^{O(k)})$.

Before detailing these constructions in [Sections 1.4](#) and [1.5](#), we first turn to our main application, which also serves as a concrete motivation for studying robust subspace designs.

1.2 Isolating a unique small quantum witness

The Valiant-Vazirani theorem [\[VV86\]](#) shows that randomness can isolate a unique witness among several classical witnesses. In complexity theoretic terms, it establishes $\text{NP} \subseteq \text{RP}^{\text{promiseUP}}$. Aharonov, Ben-Or, Brandão and Sattath [\[ABOBS22\]](#) extended this result to bounded-error quantum verifiers with classical witnesses, showing that $\text{QCMA} \subseteq \text{RP}^{\text{UQCMA}}$.

In our work, we consider quantum witnesses instead. In that case the perfectly accepting witnesses form a subspace, and their number is naturally measured by its dimension.

Now suppose that a quantum verifier's accepting subspace has dimension k , with witnesses on w qubits. Classically, if k is at most polynomial in n , a prover can just send all k witnesses in a canonical order to obtain a unique combined witness and let a polynomial time verifier check their distinctness and verify each one. A natural quantum analogue is to have the prover send their antisymmetric product. For an orthonormal basis $|\psi_1\rangle, \dots, |\psi_k\rangle$ of the accepting subspace, this is

$$\frac{1}{\sqrt{k!}} \sum_{\sigma \in S_k} \text{sgn}(\sigma) |\psi_{\sigma(1)}\rangle \otimes |\psi_{\sigma(2)}\rangle \otimes \dots \otimes |\psi_{\sigma(k)}\rangle,$$

where S_k is the collection of all permutations $\sigma : [k] \rightarrow [k]$. A quantum verifier can test whether a provided state lies in the antisymmetric subspace, i.e. whether it is spanned by antisymmetric products of k orthogonal states (ensuring that the prover cannot send k times the same witness state) and it can verify acceptance on each register. These conditions leave a one-dimensional subspace. Jain, Kerenidis, Kuperberg, Santha, Sattath and Zhang [JKK⁺11] use this observation to reduce polynomially many quantum witnesses to a unique one, showing that $\text{FewQMA} \subseteq \text{P}^{\text{UQMA}}$. This unique witness, however, occupies k registers and therefore $k \cdot w$ qubits in total.

This raises a space-bounded version of the isolation question: Can we obtain a unique quantum witness without asymptotically increasing the size of the witness register? We answer this question affirmatively for space-bounded unitary quantum verifiers with perfect completeness and an acceptance gap outside their perfectly accepting subspace. It should be observed that for bounded-error quantum verifiers this is almost immediate, because, using space-efficient error reduction, they can actually remove the prover entirely by sampling from the maximally mixed state instead (at the cost of a runtime blowup), see [FL18, Theorem 35]. In complexity-theoretic language, we prove:

Theorem 1.1. $\text{FewQUMAL}_1 \subseteq \text{RL}^{\text{UQUMAL}_1}$.

Here FewQUMAL_1 and UQUMAL_1 are quantum Merlin-Arthur protocols on $O(\log n)$ qubits for the witness and auxiliary work register with perfect completeness and, on yes-instances, a perfectly-accepting witness subspace of dimension $\text{poly}(n)$ and 1, respectively. Both require an acceptance gap on the orthogonal complement of the accepting subspace, that is, if a state is orthogonal to the subspace of states that are accepted with probability 1, then it is accepted with probability at most $1/3$. In the case of FewQUMAL_1 the polynomial bound on the witness space already follows from the fact that the verifier acts on $O(\log n)$ qubits only. The acceptance gap is the actual non-trivial limitation.

Our primary tool used to establish this quantum witness isolation is what we call *kernel peeling*, which relies on efficiently-constructible robust subspace designs:

Lemma 1.2 (Implicit from Lemma 3.6 and Remark 3.7). *Let $\eta, \gamma \in (0, 1]$ and $k \geq k'$. Given a (k', η, Δ) -weak robust subspace design $\mathcal{Q} = \{Q_1, \dots, Q_m\} \subseteq \mathbb{C}^{k' \times N}$ and a positive semi-definite matrix $H \in \mathbb{C}^{N \times N}$ with kernel dimension k and smallest non-zero eigenvalue $\lambda_{\min}^+(H) \geq \gamma$, there exists an index $i \in [m]$ such that*

$$\dim \ker(H + Q_i^* Q_i) = k - k', \quad \text{and} \quad \lambda_{\min}^+(H + Q_i^* Q_i) \geq \eta^2 \gamma / 3.$$

To see how we use this, recall that by the standard circuit-to-Hamiltonian construction due to Kitaev [KSV02], we find that for every FewQUMAL_1 protocol for a length n input x , there is an efficiently constructible positive semi-definite $\text{poly}(n) \times \text{poly}(n)$ matrix $H = H(x)$ whose kernel dimension is equal to the number of perfectly accepting witnesses. Furthermore, from the acceptance gap on yes-instances and bounded acceptance probability on no-instances, one can derive an inverse polynomial spectral gap of the constructed Hamiltonian, $\lambda_{\min}^+(H) \geq 1/\text{poly}(n)$.

If the witness subspace initially is of dimension $k = \dim \ker(H)$, then by taking a $(k - 1, \eta, \Delta)$ -weak robust subspace design $\{Q_1, \dots, Q_m\}$ with robustness parameter $\eta \geq 1/\text{poly}(n)$, there will be some Q_i so that adding this perturbation to H the kernel dimension becomes

$$\dim \ker(H + Q_i^* Q_i) = 1, \quad \text{and} \quad \lambda_{\min}^+(H + Q_i^* Q_i) \geq 1/\text{poly}(n).$$

Using the phase estimation based singularity testing algorithm from [FR21, Lemma 35], we find that this yields a unique witness protocol with acceptance gap scaling polynomially with the

spectral gap of the given matrix, where the prover precisely provides the isolated $O(\log n)$ -qubit kernel state of $H + Q_i^* Q_i$.

Of course we do not know the dimension of the witness subspace in advance. Our reduction relies on efficiently constructing (k', η, Δ) -weak robust subspace designs for every $k' \leq 2^w = \text{poly}(n)$ using randomness, and queries whether $H + Q_i^* Q_i$ has a unique kernel state using the just mentioned unique witness kernel verification procedure from [FR21, Lemma 35]. If any query answer returns yes, we accept. Once we reach $k' = \dim \ker(H) - 1$, we will find some matrix Q_i so that the kernel dimension of $H + Q_i^* Q_i$ is indeed one and which preserves the inverse polynomial spectral gap, yielding acceptance. If on the other hand, there is no accepting witness, then H will be regular to begin with and adding the positive semi-definite perturbations $Q_i^* Q_i$ will keep the matrix regular and only increase the spectral gap, so that no kernel verification procedure will give a positive answer for any k' . The resulting protocol is in the complexity class $\text{RL}^{\text{UQU}^{\text{MAL}}_1}$.

1.3 Further applications

As further applications of subspace designs and kernel peeling, we find alternative simpler proofs of known classical containment results, as well as new quantum containment results for linear algebraic problems such as testing the nullity of a given matrix or the feasibility of a system of linear equations.

For our classical applications, we consider the nullity testing problem NULL_k , which asks to decide whether the kernel dimension of a given matrix A is at least k or less than k . With a similar kernel peeling idea as for our quantum witness isolation, we find that a $(k - 1, \Delta)$ “non-robust” subspace design allows us to reduce the problem to NULL_1 which corresponds to deciding whether a given matrix is singular, and which is known to be in the exact counting class C=L . We then notice that the known Wronskian-based (non-robust) subspace design constructions, as presented in [GK16, BCDZ26], can be implemented in deterministic logspace. Combining these ideas yields an alternative proof for $\text{NULL}_k \in \text{C=L}$, originally proven in [ABO99] using a line of reductions including Berkowitz’s algorithm [Ber84]. By a reduction from linear system feasibility, FSLE , to nullity testing (also from [ABO99]) this yields an alternative simpler reduction path towards $\text{FSLE} \in \text{L}^{\text{C=L}}$.

In the quantum setting, we consider a well-conditioned version of the nullity testing problem, pc-NULL_k . Here we also wish to decide if the kernel dimension of a given matrix is at least k , but we are promised that if the kernel dimension is below k , then the k -th smallest singular value is at least inverse polynomial, $\sigma_k(A) \geq 1/\text{poly}(n)$. An analogous argument as in the classical case, together with a careful analysis of the robustness, shows that we can use our probabilistic robust subspace design construction to put pc-NULL_k in $\text{coRL}^{\text{QU}^{\text{MAL}}_1}$. Further, we find that a well-conditioned version of linear system feasibility, pc-FSLE is contained in $\text{BPL}^{\text{QU}^{\text{MAL}}_1}$. This extends our understanding of the complexity of well-conditioned linear algebra.

Table 1 gives a summary of these and related results about the complexity of fundamental linear algebra problems.

1.4 Probabilistic construction of a robust subspace design

We next summarize our probabilistic construction of a robust subspace design. It is based on complex Gaussian matrices and uses various tools from Gaussian random matrix theory [For10, HKPV06, Sza98, Ver18].

A first key fact that we use is that if Q is a $k \times n$ complex Gaussian matrix and a subspace $W \in \text{Gr}(k, n)^1$ is fixed, then the restriction $Q|_W$ of Q to W is a square Gaussian matrix. Now, for a

¹Here $\text{Gr}(k, n)$ is the complex Grassmannian, which is the set of all k -dimensional subspaces of $\mathbb{C}^n$.

Problem	General	Poly-conditioned
$\det(A) \stackrel{?}{=} 0$	C=L-complete [ABO99]	QUMAL_1 -complete [FR21]
$\text{nullity}(A) \stackrel{?}{\geq} k$	C=L-complete [ABO99]	QUMAL_1 -complete for $k = O(1)$ [Theorem C.1] $\in \text{coRL}^{\text{QUMAL}_1}$ [Theorem 6.11]
Feasible $Ax \stackrel{?}{=} b$	L ^{C=L} -complete [ABO99]	$\in \text{BPL}^{\text{QUMAL}_1}$ [Theorem 6.13]
Solve $Ax = b$	DET-complete [Coo85]	BQL-complete [TS13, FL18, FR21]

Table 1: Computational complexity of different linear algebra problems. We distinguish the general case from the poly-conditioned case, where the input instances are promised to satisfy appropriate well-conditionedness assumptions. Details are given in Section 6.

weak robust design, we need to control the probability that the least singular value of $Q|_W$ is small. For a strong robust design, however, we must also control how many singular values of $Q|_W$ can be simultaneously close to zero.

To do so, we use that the singular values of a square complex Gaussian matrix have a well-understood joint distribution. In particular, the expected number of singular values below $u/\sqrt{k}$ of a $k \times k$ Gaussian matrix is $O(u^2)$, and we can prove strong tail bounds for this number. Applying these estimates independently to the maps in our family gives a strong bound for every fixed subspace W . Finally, we use a union bound on a net over the complex Grassmannian $\text{Gr}(k, n)$ so that the same family works for every k -dimensional subspace simultaneously. This gives the following theorem.

Theorem 1.3 (Informal version of Theorem 4.5). *For every $1 \leq k < n$, there is a distribution over*

$$m = O(k(n - k) \log(nk))$$

maps $Q_i : \mathbb{C}^n \rightarrow \mathbb{C}^k$ such that, with high probability, $\{Q_i\}$ is a (k, η, Δ) -strong robust subspace design for $\eta = \Theta(1/\sqrt{nk})$ and $\Delta = m/4$.

We furthermore find that the above construction can be approximated to finite-precision by a randomized algorithm that uses only $O(\log n)$ workspace, by presenting a simple space-bounded Gaussian scalar sampler.

1.5 Explicit construction of a robust subspace design

We next turn to explicit constructions of weak robust subspace designs. We first give a simple construction based on coordinate projections. For $I \subseteq [n]$, let $P_I : \mathbb{C}^n \rightarrow \mathbb{C}^{|I|}$ denote the projection onto the coordinates indexed by I .

The basic idea is already visible when $k = 1$. If $v \in \mathbb{C}^n$ is a unit vector, then at least one coordinate of v has magnitude at least $1/\sqrt{n}$. Thus, among the coordinate projections, there is always one that is non-negligible on the line spanned by v ; the uniform vector shows that the scale $1/\sqrt{n}$ is the correct one in general. For a k -dimensional subspace, an analogous argument based on a carefully chosen $k \times k$ coordinate minor yields robustness at scale $1/\sqrt{nk}$, leading to the following construction.

Theorem 1.4 (Projection-based robust subspace design). *The family*

$$\mathcal{P}_k = \{P_I : I \subseteq [n], |I| = k\}$$

is a (k, η, Δ) -weak robust subspace design with $\eta = \frac{1}{2\sqrt{nk}}$ and $\Delta = \binom{n}{k} - 1$.

Thus the robustness parameter is inverse polynomial, but the family has size $\binom{n}{k}$, growing roughly as n^k . To reduce the size, we first use a sketching idea to compress the ambient space using a small family of sign matrices.

More precisely, we construct an explicit family of linear maps

$$S : \mathbb{C}^n \rightarrow \mathbb{C}^r, \quad r = \Theta(k^2),$$

such that for every k -dimensional subspace W , many of the maps S embed W with inverse-polynomial minimum singular value. The corresponding sketching analysis argues that this holds for any sample space with 4-wise independent, unbiased signs. We can thus use a polynomial-size, logspace-enumerable 4-wise independent sample space based on a construction of Alon–Babai–Itai [ABI86].

The sketching step reduces the ambient dimension from n to $r = \Theta(k^2)$, but its output dimension is still larger than k . We therefore compose each sketch S with one of the coordinate projections

$$P_I : \mathbb{C}^r \rightarrow \mathbb{C}^k, \quad |I| = k.$$

Since there are only

$$\binom{r}{k} = k^{O(k)}$$

such projections at this point, this yields the following explicit construction.

Theorem 1.5 (Explicit robust subspace design). *For every $1 \leq k \leq n$, there is an explicitly computable family of linear maps $\mathbb{C}^n \rightarrow \mathbb{C}^k$ forming a (k, η, Δ) -weak robust subspace design with*

$$m = O\left(n^4 k^{O(k)}\right), \quad \eta = \Omega\left(\frac{1}{k^{3/2} \sqrt{n}}\right), \quad \Delta < m.$$

Moreover, the family can be enumerated in space $O(\log n + k \log k)$, and its entries have polynomial bit complexity.

In particular, when

$$k = O\left(\frac{\log n}{\log \log n}\right),$$

we have $k^{O(k)} = n^{O(1)}$, and enumeration uses only $O(\log n)$ space. Hence in this regime the construction gives a polynomial-size, deterministic logspace-enumerable robust subspace design.

1.6 Open problems and future directions

We finish the introduction by discussing some questions that our work evokes.

First and foremost is the question of an *explicit* (e.g. logspace-uniform) *robust subspace design* of polynomial size and inverse polynomial robustness. While our probabilistic construction shows the non-explicit existence of such a design, our complexity-theoretic applications would benefit from a suitable derandomization. As an example, by [Corollary 6.6](#) such a design would improve the containment $\text{FewQUMAL}_1 \subseteq \text{RL}^{\text{UQUMAL}_1}$ to $\text{FewQUMAL}_1 \subseteq \text{L}^{\text{UQUMAL}_1}$.

One promising approach to such a derandomization is to study the connection with dimension expanders [LZ08] and quantum expanders [BASTS08], both of which are linear-algebraic variants of expander graphs. Following the discussion e.g. in [LQW⁺25], a quantum expander can be thought of as a “robust” version of a dimension expander. While prior work has used subspace designs to design dimension expanders [FG15], existing constructions of explicit quantum expanders may be useful to constructing robust subspace designs.

Further, we wonder whether our techniques may also be useful for quantum witness isolation in the time-bounded regime. A key challenge is to handle exponentially large witness spaces without losing efficiency or the spectral gap needed for verification. The quantum-oracle barrier of Anshu, Haferkamp, Hwang, and Nguyen [AHHN26] also suggests that such an extension would need to exploit information beyond black-box access to the verifier.

Finally, we mention a natural *geometric viewpoint* of subspace designs over $\mathbb{C}$. For a fixed subspace $H \subseteq \mathbb{C}^n$ of dimension $n - k$, the set of k -dimensional subspaces W satisfying

$$W \cap H \neq \{0\}$$

forms a Schubert variety in the Grassmannian $\text{Gr}(k, n)$. Thus, an ordinary subspace design controls how often W lies on these bad varieties. This viewpoint is exploited in recent work of Brakensiek, Chen, Dhar, and Zhang [BCDZ26] who sharpen the Guruswami-Kopparty bound [GK16], and use a Grassmannian/Schubert-calculus viewpoint to show that the improved bound is essentially optimal (over algebraically closed fields). It is an interesting direction to explore this geometric viewpoint for robust subspace designs, which require that W stay quantitatively away from them.

1.7 Paper overview

After the preliminaries in Section 2, we formally define robust subspace designs and introduce the kernel peeling technique in Section 3. In the same section, we recall the classical construction of subspace designs based on Wronskian evaluation maps which can have arbitrarily bad robustness parameter η . In Section 4, we describe our probabilistic construction of robust subspace designs based on Gaussian matrices. In Section 5, we present our explicit deterministic constructions. Finally, in Section 6, we show our application to isolate a unique small quantum witness as well as the linear algebra applications to nullity testing and linear system feasibility.

2 Preliminaries

For every $n \in \mathbb{N}$, we write $[n] := \{1, \dots, n\}$ for the first n natural numbers.

Throughout, all vector spaces we consider are finite-dimensional over the complex numbers $\mathbb{C}$, and equipped with the standard inner product and Euclidean norm $\|\cdot\|$. For a subspace $W \subseteq \mathbb{C}^n$, we write $W^\perp$ for its orthogonal complement and Π_W for the orthogonal projector onto W . We further write

$$\text{dist}(x, W) := \min_{w \in W} \|w - x\| = \|\Pi_{W^\perp} x\|$$

for the minimal distance of x to W . If W has dimension k , then $N_W \in \mathbb{C}^{n \times k}$ denotes an arbitrary matrix whose columns form an orthonormal basis of W . As such,

$$N_W^* N_W = I_k, \quad N_W N_W^* = \Pi_W.$$

We write $\text{Gr}(k, n)$ for the complex Grassmannian of k -dimensional subspaces of $\mathbb{C}^n$.

For a linear map $A : \mathbb{C}^n \rightarrow \mathbb{C}^m$, we denote by A^* its Hermitian adjoint, by $\ker(A)$ and $\text{im}(A)$ its kernel and image, and by $\text{rank}(A)$ its rank. We freely identify a linear map by its matrix representation, where the underlying bases will always be clear from context or mentioned explicitly. The operator and Frobenius norms are

$$\|A\|_{\text{op}} := \sup_{x \in \mathbb{C}^n \setminus \{0\}} \frac{\|Ax\|}{\|x\|}, \quad \|A\|_F := \sqrt{\text{Tr } A^*A} = \sqrt{\sum_{i,j} |A(i,j)|^2},$$

and satisfy $\|A\|_{\text{op}} \leq \|A\|_F$.

For a Hermitian matrix $H \in \mathbb{C}^{n \times n}$, we list its eigenvalues in nondecreasing order,

$$\lambda_1(H) \leq \dots \leq \lambda_n(H).$$

For $A \in \mathbb{C}^{m \times n}$, set $r := \min\{m, n\}$. We likewise list the singular values of A in non-decreasing order

$$\sigma_1(A) \leq \dots \leq \sigma_r(A).$$

Equivalently,

$$\sigma_j(A)^2 = \lambda_{n-r+j}(A^*A) = \lambda_{m-r+j}(AA^*), \quad \forall j \in [r].$$

By the Courant-Fischer theorem,

$$\sigma_j(A) = \sup_{\substack{U \subseteq \mathbb{C}^n \\ \dim U = r-j+1}} \inf_{x \in U \setminus \{0\}} \frac{\|Ax\|}{\|x\|}.$$

We write $\sigma_{\min}(A) := \sigma_1(A)$, $\sigma_{\max}(A) := \sigma_r(A) = \|A\|_{\text{op}}$, and $\sigma_{\min}^+(A)$ for the smallest non-zero singular value of A , with the convention $\sigma_{\min}^+(0) = \infty$. For Hermitian matrices H , we use the analogous notation and convention for $\lambda_{\min}(H)$, $\lambda_{\max}(H)$ and $\lambda_{\min}^+(H)$. We will also use the standard singular value perturbation bound

$$|\sigma_j(A) - \sigma_j(B)| \leq \|A - B\|_{\text{op}}$$

for matrices A and B of the same dimensions.

Finally, we record how singular values behave under restriction to a subspace. Let $Q \in \mathbb{C}^{d \times n}$ with $d \leq n$, let $W \subseteq \mathbb{C}^n$ have dimension $k \leq d$, and let $N_W \in \mathbb{C}^{n \times k}$ as above. Since $N_W : \mathbb{C}^k \rightarrow W$ is unitary, QN_W is a matrix representation of $Q|_W$. Extending N_W to a unitary $U = [N_W \ N_W^\perp] \in \mathbb{C}^{n \times n}$ gives

$$Q\Pi_W U = [QN_W \ 0].$$

Hence, for every $j \in [k]$,

$$\sigma_j(Q|_W) = \sigma_j(QN_W) = \sigma_{d-k+j}(Q\Pi_W U).$$

Throughout, computational inputs are matrices over $\mathbb{Q}(i)$ with polynomial-bit entries. Arithmetic outputs are streamed bit by bit to a write-only unidirectional output tape by machines that may use less space than the output size.

3 Robust subspace designs and kernel peeling

We first recall the definition of weak and strong robust subspace designs from the introduction. After that, we show that robust subspace designs give rise to our kernel peeling technique, the main technical ingredient for our applications. Finally, we remark that the ordinary subspace design constructions based on Wronskian evaluation maps from [GK16, BCDZ26] can be implemented space-efficiently. However, we find they may have arbitrarily small robustness parameter $\eta \in (0, 1]$, making them useful only to our classical complexity applications, and not for quantum.

3.1 Definitions

For all definitions, let $\eta \in [0, 1]$, $n \in \mathbb{N}$ and $k \leq n$, $\Delta \geq 0$.

Definition 3.1 (Weak robust subspace design). A family of linear maps $Q_1, \dots, Q_m : \mathbb{C}^n \rightarrow \mathbb{C}^k$ is a (k, η, Δ) -weak robust subspace design if $\Delta < m$, $\|Q_i\|_{\text{op}} \leq 1$ for every i , and for every k -dimensional subspace $W \subseteq \mathbb{C}^n$,

$$|\{i \in [m] : \sigma_{\min}(Q_i|_W) \leq \eta\}| \leq \Delta.$$

Informally, we call a map Q_i *bad* for W if it contributes to the above count, i.e. if $\sigma_{\min}(Q_i|_W) \leq \eta$.

Definition 3.2 (Strong robust subspace design). A family of linear maps $Q_1, \dots, Q_m : \mathbb{C}^n \rightarrow \mathbb{C}^k$ is a (k, η, Δ) -strong robust subspace design if $\Delta < mk$, $\|Q_i\|_{\text{op}} \leq 1$ for every i , and for every k -dimensional subspace $W \subseteq \mathbb{C}^n$,

$$\sum_{i=1}^m |\{j \in [k] : \sigma_j(Q_i|_W) \leq \eta\}| \leq \Delta.$$

Remark 3.3 (Co-domain dimension condition). We restrict the linear maps in a robust subspace design to have co-domain of dimension k , because this is what we require for our linear algebraic applications. We do later consider the relaxation to larger co-domain maps, but then we call the family a sketching family (see [Definition 5.3](#)) because of their relation to sketching.

As in the non-robust world: Every strong robust design with $\Delta < m$ is also a weak robust design with the same parameters, because every bad map contributes at least one to the above sum. Further, as noted in the introduction, it is straightforward to observe that

$$\mathcal{H} = \{\ker(Q_i) : i \in [m]\}$$

is a (k, Δ) -weak/strong subspace design if the underlying Q_i 's are a (k, η, Δ) -weak/strong robust subspace design for any $\eta \in [0, 1]$. Finally, we will sometimes also refer to non-robust subspace designs as families of maps, rather than as families of subspaces, by which we will mean the corresponding kernel-families.

Remark 3.4 (Geometric interpretation). Robust subspace designs admit a natural geometric interpretation in terms of approximate intersections of subspaces. For subspaces $W, H \subseteq \mathbb{C}^n$, define the η -robust intersection dimension by

$$\text{rdim}_\eta(W, H) := \max \left\{ \dim V : V \subseteq W, \text{dist}(v, H) \leq \eta \|v\| \text{ for all } v \in V \right\}.$$

Thus $\text{rdim}_0(W, H) = \dim(W \cap H)$, while for $\eta > 0$ the quantity also counts directions in W that are close to lying in H .

This notion is directly characterized by the singular values of the restricted projection. Indeed, if $\dim W = k$, then the Courant-Fischer min-max characterization of singular values gives

$$\text{rdim}_\eta(W, H) = |\{j \in [k] : \sigma_j(\Pi_{H^\perp}|_W) \leq \eta\}|.$$

More generally, let $Q : \mathbb{C}^n \rightarrow \mathbb{C}^k$ satisfy $\|Q\|_{\text{op}} \leq 1$ and let $H = \ker(Q)$. Since $Qx = Q\Pi_{H^\perp}x$, we have

$$\|Qv\| \leq \|\Pi_{H^\perp}v\| = \text{dist}(v, H) \quad \text{for every } v \in \mathbb{C}^n.$$

Consequently,

$$\sigma_j(Q|_W) \leq \sigma_j(\Pi_{H^\perp}|_W) \quad \text{for every } j \in [k].$$

Hence a large robust intersection between W and $\ker(Q)$ necessarily produces at least as many small singular values of $Q|_W$.

When Q is a co-isometry, i.e. in particular $Q^*Q = \Pi_{\ker(Q)^\perp}$, then the above inequality becomes an equality. Thus, in the special case where Q is a co-isometry, our definition of a robust subspace design is exactly the geometric notion obtained by replacing ordinary intersections by robust intersections.

Although this geometric viewpoint is intuitive, the formulation in terms of general linear maps is more flexible and better suited to our applications.

3.2 Kernel peeling

In this subsection we show that an explicit (robust) subspace design allows us to apply a kernel peeling procedure, that is it allows us to compute a family of rank k' perturbations so that for every matrix A with kernel dimension $k \geq k'$, there is some perturbation in the family so that adding it to A reduces its kernel dimension by exactly k' , and preserves a spectral gap.

Our main technical tool is the following lemma which states that if two maps Q_R and Q_L are invertible on the right- and left-kernel of A , respectively, then $A + Q_L^*Q_R$ must be invertible on the full space. The lemma also proves a lower bound on the resulting minimal singular value of $A + Q_L^*Q_R$ depending on the smallest non-zero singular value of A .

Lemma 3.5 (Kernel removal). *Let $\eta, \gamma \in (0, 1]$, $1 \leq k \leq n$ and let $A \in \mathbb{C}^{n \times n}$ satisfy*

$$\dim \ker A = k, \quad \sigma_{\min}^+(A) \geq \gamma.$$

Further, let $Q_R, Q_L \in \mathbb{C}^{k \times n}$ satisfy $\|Q_R\|_{\text{op}}, \|Q_L\|_{\text{op}} \leq 1$ and

$$\sigma_{\min}(Q_R|_{\ker A}), \sigma_{\min}(Q_L|_{\ker A^*}) \geq \eta.$$

*Then $A + Q_L^*Q_R$ is invertible and*

$$\sigma_{\min}(A + Q_L^*Q_R) \geq \frac{\eta^2\gamma}{3}.$$

Proof. If $k = n$, then $A = 0$ and $\sigma_{\min}(Q_L^*Q_R) \geq \eta^2 \geq \eta^2\gamma/3$, so assume $k < n$. With respect to orthonormal bases adapted to

$$\mathbb{C}^n = \ker(A)^\perp \oplus \ker(A) \quad \text{and} \quad \mathbb{C}^n = \ker(A^*)^\perp \oplus \ker(A^*),$$

we find unitarily equivalent representations of A and $B := Q_L^*Q_R$ as

$$A \sim \begin{pmatrix} A_{11} & 0 \\ 0 & 0 \end{pmatrix}, \quad B \sim \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix},$$

where $\sigma_{\min}^+(A) = \sigma_{\min}(A_{11}) \geq \gamma$ and $B_{22} = (Q_L|_{\ker A^*})^*(Q_R|_{\ker A})$ such that $\sigma_{\min}(B_{22}) \geq \eta^2$. In particular B_{22} is invertible with $\text{rank}(B_{22}) = k$. However, since $\text{rank}(B) = k$ already, we find that we can express the first $n - k$ columns of the above block representation of B by linear combinations of the last k columns, i.e. there exists some matrix $X \in \mathbb{C}^{k \times (n-k)}$ such that

$$\begin{pmatrix} B_{11} \\ B_{21} \end{pmatrix} = \begin{pmatrix} B_{12} \\ B_{22} \end{pmatrix} X.$$

From the lower block we obtain that $X = B_{22}^{-1}B_{21}$ and then from the top block we get

$$B_{11} = B_{12}B_{22}^{-1}B_{21}. \quad (1)$$

Moreover,

$$\|B_{22}^{-1}B_{21}\|_{\text{op}}, \|B_{12}B_{22}^{-1}\|_{\text{op}} \leq 1/\eta, \quad (2)$$

because, respectively,

$$B_{22}^{-1}B_{21} = (Q_R|_{\ker A})^{-1}Q_R|_{\ker(A)^\perp}, \quad B_{12}B_{22}^{-1} = \left(Q_L|_{\ker(A^*)^\perp}\right)^* \left((Q_L|_{\ker A^*})^*\right)^{-1}.$$

Using Equation (1), we find the following block factorization of $A + B$:

$$A + B \sim \begin{pmatrix} A_{11} + B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix} = \begin{pmatrix} I & B_{12}B_{22}^{-1} \\ 0 & I \end{pmatrix} \begin{pmatrix} A_{11} & 0 \\ 0 & B_{22} \end{pmatrix} \begin{pmatrix} I & 0 \\ B_{22}^{-1}B_{21} & I \end{pmatrix}.$$

As each block triangular matrix on the factorization on the right is regular, we find that also $A + B$ must be regular. A simple calculation verifies that its inverse can be factorized as

$$\begin{aligned} (A + B)^{-1} &\sim \begin{pmatrix} I & 0 \\ -B_{22}^{-1}B_{21} & I \end{pmatrix} \begin{pmatrix} A_{11}^{-1} & 0 \\ 0 & B_{22}^{-1} \end{pmatrix} \begin{pmatrix} I & -B_{12}B_{22}^{-1} \\ 0 & I \end{pmatrix} \\ &= \begin{pmatrix} I & 0 \\ -B_{22}^{-1}B_{21} & I \end{pmatrix} A_{11}^{-1} (I \quad -B_{12}B_{22}^{-1}) + \begin{pmatrix} 0 \\ I \end{pmatrix} B_{22}^{-1} (0 \quad I). \end{aligned}$$

Since the block matrix representation of $(A + B)^{-1}$ above is unitarily equivalent to $(A + B)^{-1}$, we can bound its operator norm by

$$\begin{aligned} \|(A + B)^{-1}\|_{\text{op}} &\leq \sqrt{1 + \|B_{22}^{-1}B_{21}\|_{\text{op}}^2} \|A_{11}^{-1}\|_{\text{op}} \sqrt{1 + \|B_{12}B_{22}^{-1}\|_{\text{op}}^2} + \|B_{22}^{-1}\|_{\text{op}} \\ &\leq \frac{1 + \eta^{-2}}{\gamma} + \frac{1}{\eta^2} \leq \frac{3}{\eta^2\gamma}, \end{aligned}$$

where we used Equation (2) to establish the second inequality. This gives

$$\sigma_{\min}(A + B) = \frac{1}{\|(A + B)^{-1}\|_{\text{op}}} \geq \frac{\eta^2\gamma}{3}$$

and finishes the proof. $\square$

Lemma 3.6 (Kernel peeling). *Let $1 \leq k' \leq k \leq n$, $\eta, \gamma \in (0, 1]$, and let $A \in \mathbb{C}^{n \times n}$ satisfy*

$$\dim \ker A = k, \quad \sigma_{\min}^+(A) \geq \gamma.$$

Let $Q_1, \dots, Q_m : \mathbb{C}^n \rightarrow \mathbb{C}^{k'}$ be a (k', η, Δ) -weak robust subspace design. Then there exist $i, j \in [m]$ such that

$$\dim \ker(A + Q_i^*Q_j) = k - k', \quad \sigma_{\min}^+(A + Q_i^*Q_j) \geq \frac{\eta^2\gamma}{3}.$$

Proof. Fix any k' -dimensional subspaces $W_R \subseteq \ker A$ and $W_L \subseteq \ker A^*$. Since $m > \Delta$, by the design property, there must exist indices $i, j \in [m]$ for which

$$\sigma_{\min}(Q_j|_{W_R}), \sigma_{\min}(Q_i|_{W_L}) \geq \eta.$$

We set

$$W'_R := \ker(Q_j|_{\ker(A)}) = \ker(Q_j) \cap \ker(A), \quad W'_L := \ker(Q_i|_{\ker(A^*)}) = \ker(Q_i) \cap \ker(A^*),$$

so that $W'_R + W_R = \ker(A)$ and $W'_L + W_L = \ker(A^*)$. Since $\dim \ker A = \dim \ker A^* = k$ and both $Q_j|_{\ker(A)}$ and $Q_i|_{\ker(A^*)}$ have rank k' , the rank-nullity theorem yields

$$\dim W'_R = \dim W'_L = k - k'.$$

Consider the restrictions

$$\tilde{A} := A|_{(W'_R)^\perp} : (W'_R)^\perp \rightarrow (W'_L)^\perp, \quad \tilde{Q}_j := Q_j|_{(W'_R)^\perp}, \quad \tilde{Q}_i := Q_i|_{(W'_L)^\perp}.$$

Choosing orthonormal bases of $(W'_R)^\perp$ and $(W'_L)^\perp$, we identify these maps with matrices of size $\tilde{n} \times \tilde{n}$ and $k' \times \tilde{n}$, respectively, where $\tilde{n} := n - k + k'$. Since $W'_R \subseteq \ker A$ and $W'_L \subseteq \ker A^*$,

$$\dim \ker \tilde{A} = k', \quad \sigma_{\min}^+(\tilde{A}) = \sigma_{\min}^+(A) \geq \gamma.$$

Moreover,

$$\ker \tilde{A} = \ker A \cap (W'_R)^\perp.$$

As $W'_R = \ker(Q_j|_{\ker A})$, the restriction of $Q_j|_{\ker A}$ to $\ker A \cap (W'_R)^\perp$ has precisely the nonzero singular values of $Q_j|_{\ker A}$. Therefore

$$\sigma_{\min}(\tilde{Q}_j|_{\ker \tilde{A}}) = \sigma_{\min}^+(Q_j|_{\ker A}) \geq \sigma_{\min}(Q_j|_{W_R}) \geq \eta,$$

where the middle inequality follows from the Courant-Fischer min-max characterization of singular values. Similarly,

$$\sigma_{\min}(\tilde{Q}_i|_{\ker \tilde{A}^*}) \geq \eta.$$

Hence [Lemma 3.5](#) gives

$$\sigma_{\min}(\tilde{A} + \tilde{Q}_i^* \tilde{Q}_j) \geq \frac{\eta^2 \gamma}{3}.$$

Finally, $A + Q_i^* Q_j$ vanishes on W'_R , and $(A + Q_i^* Q_j)^*$ vanishes on W'_L . Thus, its restriction from $(W'_R)^\perp$ to $(W'_L)^\perp$ is exactly $\tilde{A} + \tilde{Q}_i^* \tilde{Q}_j$, which is invertible. Consequently,

$$\ker(A + Q_i^* Q_j) = W'_R,$$

and thus

$$\dim \ker(A + Q_i^* Q_j) = k - k', \quad \sigma_{\min}^+(A + Q_i^* Q_j) \geq \frac{\eta^2 \gamma}{3}. \quad \square$$

Observe that the condition in [Lemma 3.6](#) that the subspace design should be of size $m > \Delta$ simply ensures that the constraint of being a subspace design is non-trivial.

In most of our applications we will only need the above kernel peeling lemma for Hermitian matrices A . We make the following remark about this case:

Remark 3.7 (Positive semidefinite perturbations). If A in [Lemma 3.6](#) is Hermitian, then the two indices can be chosen equal: there exists $i \in [m]$ such that

$$\dim \ker(A + Q_i^* Q_i) = k - k', \quad \sigma_{\min}^+(A + Q_i^* Q_i) \geq \frac{\eta^2 \gamma}{3}.$$

Indeed, since $\ker(A) = \ker(A^*)$, we may choose $W_R = W_L$ in the proof of [Lemma 3.6](#). The condition $m > \Delta$ guarantees an index i with $\sigma_{\min}(Q_i|_{W_R}) \geq \eta$, and the remainder of the proof applies unchanged. Thus it suffices to consider the m positive semi-definite perturbations $Q_i^* Q_i$ rather than all m^2 products $Q_i^* Q_j$.

For general A , the same conclusion holds if $m > 2\Delta$: at most Δ indices fail the required bound on each of W_R and W_L , so some index satisfies both bounds.

3.3 Explicit non-robust construction based on Wronskian evaluation maps

We first describe a simplified variant over $\mathbb{C}$ of the explicit subspace design due to Guruswami and Kopparty [GK16] that fits into our framework. While their construction is based on the folded Wronskian over finite fields, it was extended to arbitrary fields and their parameter analysis was improved by Brakensiek, Chen, Dhar and Zhang [BCDZ26]. We give a construction based on the standard Wronskian and mostly follow the analysis of the latter work. We observe that the construction can be performed space-efficiently. Since our analysis is similar to what is already known from [BCDZ26], we defer proofs to the appendix.

For $x = (x_0, \dots, x_{n-1}) \in \mathbb{C}^n$, let

$$f_x(z) := x_0 + x_1 z + \dots + x_{n-1} z^{n-1}.$$

For $1 \leq k \leq n$ and $a \in \mathbb{C}$, define the linear map

$$Q_{k,a} : \mathbb{C}^n \longrightarrow \mathbb{C}^k, \quad Q_{k,a}(x) = (f_x(a), f'_x(a), \dots, f_x^{(k-1)}(a))^T$$

and note that $\text{rank}(Q_{k,a}) = k$.

Theorem 3.8 (Wronskian subspace design). *For every $1 \leq k < n$, and every set $T_k \subseteq \mathbb{R}$ of $m > k(n-k)$ points, the family*

$$\mathcal{H}_{n,k} := \{H_a : a \in T_k\}, \quad \text{where} \quad H_a := \ker(Q_{k,a}),$$

is a $(k, k(n-k))$ -strong subspace design. Moreover, if $T_k = \{0, 1, \dots, m-1\}$, then the entries of all $Q_{k,a}$ are of polynomial bit complexity and computable in space $O(\log nm)$.

We include a self-contained proof of the above theorem in [Appendix A](#).

3.4 Subspace designs need not be robust

Recall again that every robust subspace design induces an ordinary subspace design by taking the kernels of its maps. In this subsection, we show that the converse fails in a quantitative sense: An ordinary strong subspace design may have arbitrarily poor robustness. In fact our example is a special case of the Wronskian design from [Theorem 3.8](#). This result qualitatively matches an argument in [LQW⁺25, Theorem 1.8] showing that dimension expanders can be arbitrarily poor quantum expanders.

Theorem 3.9. *For every $n \geq 2$ and every $\eta = \eta(n) \in (0, 1]$, there exist maps*

$$Q_1, \dots, Q_n : \mathbb{C}^n \rightarrow \mathbb{C}$$

with $\|Q_i\|_{\text{op}} = 1$ such that $\{\ker(Q_1), \dots, \ker(Q_n)\}$ is a strong $(1, n-1)$ -subspace design, while $\{Q_1, \dots, Q_n\}$ is not a $(1, \eta, n-1)$ -weak robust subspace design.

Proof. For $a \in \mathbb{C}$, consider

$$Q_{1,a}(x_0, \dots, x_{n-1}) = x_0 + x_1 a + \dots + x_{n-1} a^{n-1}, \quad \widehat{Q}_a := \frac{Q_{1,a}}{\|Q_{1,a}\|_{\text{op}}}.$$

Choose $\varepsilon > 0$ such that

$$(n-1)\varepsilon < \eta \quad \text{and let} \quad T := \{0, \varepsilon, 2\varepsilon, \dots, (n-1)\varepsilon\}.$$

While it follows already from [Theorem 3.8](#), we show that $\{\ker(\widehat{Q}_a) : a \in T\}$ is a $(1, n-1)$ strong subspace design. To this end, let $W = \text{span}\{y\}$ be any one-dimensional subspace of $\mathbb{C}^n$, where $y \neq 0$, and associate with y the nonzero polynomial

$$f_y(z) = y_0 + y_1 z + \cdots + y_{n-1} z^{n-1}.$$

Since

$$\ker(\widehat{Q}_a) \cap W \neq \{0\} \iff f_y(a) = 0,$$

and f_y has degree at most $n-1$, this can occur for at most $n-1$ points $a \in T$. For one-dimensional W , weak and strong subspace designs coincide, proving the claim.

On the other hand, let

$$W := \text{span}\{e_2\}, \quad e_2 = (0, 1, 0, \dots, 0)^T.$$

For every $a \in T$,

$$\sigma_{\min}(\widehat{Q}_a|_W) = |\widehat{Q}_a e_2| = \frac{|a|}{\|Q_{1,a}\|_{\text{op}}} \leq |a| \leq (n-1)\varepsilon < \eta,$$

where we used that $\|Q_{1,a}\|_{\text{op}} \geq 1$. Thus all n maps are bad on the same one-dimensional subspace W , and the family $\{\widehat{Q}_a : a \in T\}$ is not a $(1, \eta, n-1)$ -weak robust subspace design. $\square$

4 Probabilistic robust construction

In this section we detail our randomized construction of robust subspace designs. Our analysis borrows tools related to the Grassmannian and random matrix theory, which we introduce first.

4.1 The Grassmannian manifold of subspaces

Recall that we let $\text{Gr}(k, n)$ denote the complex Grassmannian of k -dimensional subspaces of $\mathbb{C}^n$. We equip $\text{Gr}(k, n)$ with the projection metric

$$d(W, V) := \|\Pi_W - \Pi_V\|_{\text{op}}.$$

We will use a net with respect to this metric. The following lemma follows from [\[Sza98\]](#), and is stated in e.g. [\[MWK⁺22\]](#).

Lemma 4.1 (Grassmannian net). *There exists an absolute constant $C_{\text{net}} > 0$ such that, for every $0 < \rho < 1$, the complex Grassmannian $\text{Gr}(k, n)$ admits a ρ -net $\{V_1, \dots, V_M\}$ with respect to the projection metric, of cardinality*

$$M \leq \left(\frac{C_{\text{net}}}{\rho} \right)^{2k(n-k)}.$$

That is, for every $W \in \text{Gr}(k, n)$, there exists $i \in [M]$ such that

$$d(V_i, W) \leq \rho.$$

Finally, we will also use that the singular values of a projected matrix are Lipschitz continuous on the Grassmannian manifold.

Lemma 4.2 (Lipschitz continuity). *Let $Q : \mathbb{C}^n \rightarrow \mathbb{C}^k$ and $W, V \in \text{Gr}(k, n)$. For every $j \in [k]$,*

$$|\sigma_j(Q|_W) - \sigma_j(Q|_V)| \leq \|Q\|_{\text{op}} \|\Pi_W - \Pi_V\|_{\text{op}}.$$

Proof. Since Q has a co-domain of dimension $k \leq n$ and $\dim W = \dim V = k$, we have $\sigma_j(Q|_W) = \sigma_j(Q\Pi_W)$ for every $j \leq k$, and similarly for V . Hence Weyl's inequality gives

$$|\sigma_j(Q|_W) - \sigma_j(Q|_V)| = |\sigma_j(Q\Pi_W) - \sigma_j(Q\Pi_V)| \leq \|Q(\Pi_W - \Pi_V)\|_{\text{op}},$$

which is at most

$$\|Q\|_{\text{op}}\|\Pi_W - \Pi_V\|_{\text{op}}.$$

□

4.2 Random Gaussian matrix theory

A standard complex Gaussian random variable, denoted $g \sim \mathcal{N}_{\mathbb{C}}(0, 1)$, has density $\pi^{-1}e^{-|z|^2}$ on $\mathbb{C}$. Equivalently, $g = (X + iY)/\sqrt{2}$, where X, Y are independent real standard Gaussians. The following lemma shows concentration of the operator norm of G .

Lemma 4.3 (Gaussian operator norm [AS17, Proposition 6.33]). *Let $G \in \mathbb{C}^{k \times n}$, with $k \leq n$, have independent entries distributed as $\mathcal{N}_{\mathbb{C}}(0, 1/n)$. Then*

$$\Pr[\|G\|_{\text{op}} > 3] \leq \frac{1}{2}e^{-n}.$$

Now consider a square, standard complex Gaussian matrix $G \in \mathbb{C}^{k \times k}$, and list its singular values

$$0 \leq \sigma_1(G) \leq \dots \leq \sigma_k(G).$$

For $\eta > 0$, define

$$b_\eta(G) := |\{j \in [k] : \sigma_j(G) \leq \eta\}|.$$

In [Appendix B](#) we prove the following bound on the concentration of $b_\eta(G)$.

Lemma 4.4 (Concentration of small singular values, [Claim B.2](#)). *Fix $0 < u \leq \frac{1}{2\sqrt{2}}$. If $G_1, \dots, G_m$ are independent copies of $G \in \mathbb{C}^{k \times k}$, then*

$$\Pr\left[\sum_{i=1}^m b_{u/\sqrt{k}}(G_i) \geq \frac{m}{4}\right] \leq e^{-u^2 m/3}.$$

4.3 The Gaussian construction

The following theorem combines the concentration bounds from last section with the Grassmannian net to argue that a family of $m = \text{poly}(n)$ independent Gaussian matrices $\{Q_1, \dots, Q_m\}$ form a robust subspace design with high probability.

For $Q : \mathbb{C}^n \rightarrow \mathbb{C}^k$, $W \in \text{Gr}(k, n)$ and $\eta \in (0, 1]$, we will use the notation

$$b_\eta(Q, W) := |\{j \in [k] : \sigma_j(Q|_W) \leq \eta\}|.$$

Equivalently, if $N \in \mathbb{C}^{n \times k}$ has orthonormal columns spanning W , then $b_\eta(Q, W)$ is the number of singular values of QN at most η .

Theorem 4.5 (Gaussian strong robust design). *There are absolute constants $C, c > 0$ with the following property. Fix $1 \leq k < n$ and*

$$m \geq Ck(n - k) \log(nk).$$

Choose $Q_1, \dots, Q_m \in \mathbb{C}^{k \times n}$ independently, with entries $\mathcal{N}_{\mathbb{C}}(0, 1/(9n))$. Then, with probability at least

$$1 - e^{-cm} - \frac{m}{2}e^{-n},$$

the family $\{Q_1, \dots, Q_m\}$ is a $(k, \frac{1}{18\sqrt{nk}}, m/4)$ -strong robust subspace design.

Proof. Fix constants $\eta = \frac{1}{9\sqrt{nk}}$ and $\rho = \frac{\eta}{2}$. Fix $W \in \text{Gr}(k, n)$ and let $N \in \mathbb{C}^{n \times k}$ have orthonormal columns spanning W . The key fact that we will use is that by unitary invariance of the complex Gaussian distribution,

$$Q_i N = \frac{1}{3\sqrt{n}} G_i,$$

for $G_i \in \mathbb{C}^{k \times k}$ with independent standard complex Gaussian entries. Since

$$b_\eta(Q_i, W) = b_{3\eta\sqrt{n}}(G_i) = b_{1/(3\sqrt{k})}(G_i),$$

by our choice of η , we can apply [Lemma 4.4](#) with $u = 1/3$ to get

$$\Pr \left[\sum_{i=1}^m b_\eta(Q_i, W) \geq \frac{m}{4} \right] \leq e^{-m/27}.$$

Now let $\mathcal{N}_\rho$ be a ρ -net of $\text{Gr}(k, n)$ as in [Lemma 4.1](#). A union bound gives

$$\Pr \left[\exists W \in \mathcal{N}_\rho : \sum_{i=1}^m b_\eta(Q_i, W) \geq \frac{m}{4} \right] \leq \left(\frac{C_{\text{net}}}{\rho} \right)^{2k(n-k)} e^{-m/27}.$$

By our choice of $\rho = \Theta(1/\sqrt{nk})$,

$$\log \left(\frac{C_{\text{net}}}{\rho} \right) = O(\log(nk)),$$

and so there exist constants C, c such that

$$\left(\frac{C_{\text{net}}}{\rho} \right)^{2k(n-k)} e^{-m/27} \leq e^{-cm}$$

provided $m \geq Ck(n-k)\log(nk)$.

By [Lemma 4.3](#) applied to the operator $3Q_i$ (whose entries are distributed as $\mathcal{N}_{\mathbb{C}}(0, 1/n)$), and a union bound over the Q_i 's, it holds that

$$\Pr \left[\exists i \in [m] : \|Q_i\|_{\text{op}} > 1 \right] \leq \frac{m}{2} e^{-n}.$$

Assume that neither failure event occurs. Let $W \in \text{Gr}(k, n)$ be arbitrary, and choose $V \in \mathcal{N}_\rho$ with $d(W, V) \leq \rho$. For every i and every singular-value index j , [Lemma 4.2](#) gives

$$|\sigma_j(Q_i|_W) - \sigma_j(Q_i|_V)| \leq \rho = \frac{\eta}{2}.$$

Therefore every singular value of $Q_i|_W$ at most $\eta/2$ corresponds to a singular value of $Q_i|_V$ at most η , and hence

$$b_{\eta/2}(Q_i, W) \leq b_\eta(Q_i, V).$$

By our assumption, summing over i yields

$$\sum_{i=1}^m b_{\eta/2}(Q_i, W) \leq \sum_{i=1}^m b_\eta(Q_i, V) \leq \frac{m}{4}.$$

This holds simultaneously for every W , making the Q_i 's a $(k, \eta/2, m/4)$ robust subspace design. $\square$

4.4 Computability in randomized logspace

We now show how to obtain a finite-precision version of the Gaussian construction in [Theorem 4.5](#) that is computable in randomized logspace. This will be used in all of our complexity-theoretic applications.

To this end, we first show that a randomized algorithm can output a finite-precision scalar that approximates a continuous Gaussian scalar. One subtlety in formalizing such an approximation is that the output of a randomized algorithm has a discrete distribution, whereas a Gaussian random variable has a continuous distribution. We handle this using the notion of a *coupling*. A coupling of two random variables X and Y is a pair of random variables (X', Y') defined on a common probability space such that X' has the same distribution as X and Y' has the same distribution as Y . As is customary, we will identify X' with X and Y' with Y once such a coupling has been fixed.

Lemma 4.6 (Finite-precision Gaussian sampling). *Let $\varepsilon, \delta \in (0, 1/2]$. There is a randomized algorithm which uses $O(\log(1/\varepsilon) + \log \log(1/\delta))$ workspace, runs in time $\text{poly}(1/\varepsilon, \log(1/\delta))$, and outputs a complex rational $\tilde{g} \in \mathbb{Q}(i)$ with the following properties:*

- (i) *it has an $O(\log(1/\varepsilon) + \log \log(1/\delta))$ -bit dyadic representation,*
- (ii) *its modulus is bounded by $|\tilde{g}| = O(\sqrt{\log(1/\delta)})$, and*
- (iii) *there is a coupling of $\tilde{g}$ with $g \sim \mathcal{N}_{\mathbb{C}}(0, 1)$ such that*

$$\Pr[|\tilde{g} - g| > \varepsilon] \leq \delta.$$

Proof. We first consider the case of a real Gaussian scalar. Let $Z_1, \dots, Z_N$ be independent Bernoulli random variables taking values 0 and 1 with equal probability, and set $S_N := Z_1 + \dots + Z_N$. Furthermore, we define

$$\tilde{X}_N := \frac{2S_N - N}{\sqrt{N}}.$$

By Tusnády's lemma [[MZ12](#), Proposition T], there is a coupling of $\tilde{X}_N$ with a standard real Gaussian $X \sim \mathcal{N}_{\mathbb{R}}(0, 1)$ such that

$$|X - \tilde{X}_N| \leq \frac{2}{\sqrt{N}} + \frac{X^2}{4\sqrt{N}}.$$

For $B := \left\lceil \sqrt{2 \log(4/\delta)} \right\rceil$, we find by the standard Gaussian tail bound,

$$\Pr[|X| > B] \leq 2e^{-B^2/2} \leq \frac{\delta}{2}.$$

We now choose

$$N = 2^{2k-1} = O(B^4/\varepsilon^2) = O(\log^2(1/\delta)/\varepsilon^2)$$

for some integer $k \geq 1$ large enough so that, on the event $|X| \leq B$, $|X - \tilde{X}_N| \leq \varepsilon$. Consequently,

$$\Pr[|X - \tilde{X}_N| > \varepsilon] \leq \frac{\delta}{2}.$$

We now pass to a complex Gaussian. Take two independent copies $\tilde{X}_N$ and $\tilde{Y}_N$ of the above construction, coupled respectively with independent standard real Gaussians $X, Y \sim \mathcal{N}_{\mathbb{R}}(0, 1)$, and define $g := (X + iY)/\sqrt{2} \sim \mathcal{N}_{\mathbb{C}}(0, 1)$. Furthermore, set

$$\tilde{g} := (\tilde{X}_N + i\tilde{Y}_N)/\sqrt{2},$$

whenever both $\tilde{X}_N$ and $\tilde{Y}_N$ are bounded by $B + \varepsilon$ in modulus, and set $\tilde{g} := 0$ otherwise. Since $N = 2^{2k-1}$, we have

$$\frac{\tilde{X}_N}{\sqrt{2}} = \frac{2S_N - N}{\sqrt{2N}} = \frac{2S_N - N}{2^k},$$

and likewise for $\tilde{Y}_N/\sqrt{2}$. Hence $\tilde{g}$ is a dyadic complex rational. Moreover, since $2S_N - N$ has $O(\log N)$ bits and the denominator is $2^k \leq N$, its dyadic representation uses $O(\log N) = O(\log(1/\varepsilon) + \log \log(1/\delta))$ bits. Further,

$$|\tilde{g}| \leq B + \varepsilon = O(\sqrt{\log(1/\delta)}).$$

On the event $|X| \leq B$ and $|Y| \leq B$, we have that $|\tilde{X}_N| \leq B + \varepsilon$ and $|\tilde{Y}_N| \leq B + \varepsilon$. Therefore

$$|g - \tilde{g}|^2 \leq \frac{|X - \tilde{X}_N|^2}{2} + \frac{|Y - \tilde{Y}_N|^2}{2} \leq \varepsilon^2.$$

Thus, by a union bound,

$$\Pr[|g - \tilde{g}| > \varepsilon] \leq \delta.$$

It remains to verify the computational bounds. An algorithm can generate the N Bernoulli variables used for the real and imaginary part sequentially and only needs to maintain their running sum and a counter. From the final value of the sum, it can directly determine whether the $B + \varepsilon$ truncation condition is satisfied and accordingly output either the resulting dyadic value or 0. It may therefore use $O(\log N) = O(\log(1/\varepsilon) + \log \log(1/\delta))$ workspace and run in time $\text{poly}(1/\varepsilon, \log(1/\delta))$. $\square$

The above lemma shows in particular that for inverse-polynomial accuracy and inverse-exponential failure probability, we only require $O(\log n)$ workspace to approximately sample a Gaussian scalar. Applying the sampler independently entrywise gives the following finite-precision version of the Gaussian construction from [Theorem 4.5](#).

Theorem 4.7 (Finite-precision strong robust subspace designs). *Let $C, c > 0$ be the constants from [Theorem 4.5](#). For every $1 \leq k < n$ and*

$$m \geq Ck(n - k) \log(nk),$$

there is a randomized algorithm that outputs $\tilde{Q}_1, \dots, \tilde{Q}_m \in \mathbb{Q}(i)^{k \times n}$ such that, with probability at least

$$1 - e^{-cm} - \frac{m}{2}e^{-n} - e^{-(m+n)},$$

the family $\{\tilde{Q}_1, \dots, \tilde{Q}_m\}$ is a $\left(k, \frac{1}{144\sqrt{nk}}, \frac{m}{4}\right)$ strong robust subspace design.

The algorithm uses $O(\log(nm))$ workspace and runs in $\text{poly}(n, m)$ time. Moreover, all entries of the matrices $\tilde{Q}_i$ are complex rationals representable using $O(\log(nm))$ bits.

Proof. Set

$$\eta := \frac{1}{18\sqrt{nk}}, \quad \tilde{\eta} := \frac{\eta}{8} = \frac{1}{144\sqrt{nk}},$$

and choose

$$d := \lceil 6\sqrt{n} \rceil.$$

Independently for every $\ell \in [m]$, $i \in [k]$, and $j \in [n]$, apply the finite-precision scalar sampler from [Lemma 4.6](#) to obtain a dyadic complex rational $\tilde{g}_{\ell,i,j}$ with accuracy and failure probability

$$\varepsilon := \frac{d}{144nk}, \quad \delta := \frac{2^{-2(m+n)}}{mkn}.$$

We may couple these variables with independent $g_{\ell,i,j} \sim \mathcal{N}_{\mathbb{C}}(0, 1)$ such that

$$\Pr[|\tilde{g}_{\ell,i,j} - g_{\ell,i,j}| > \varepsilon] \leq \delta$$

for every entry.

Let $G_\ell, \tilde{G}_\ell \in \mathbb{C}^{k \times n}$ denote the matrices with entries $g_{\ell,i,j}$ and $\tilde{g}_{\ell,i,j}$, respectively, and define

$$Q_\ell := \frac{G_\ell}{d}, \quad \tilde{Q}_\ell := \frac{\tilde{G}_\ell}{d}.$$

Let $\mathcal{E}$ be the event that $|\tilde{g}_{\ell,i,j} - g_{\ell,i,j}| \leq \varepsilon$ simultaneously for all mkn entries. By a union bound,

$$\Pr[\mathcal{E}^c] \leq mkn \cdot \delta \leq 2^{-2(m+n)} \leq e^{-(m+n)}.$$

On $\mathcal{E}$, for every $\ell \in [m]$,

$$\|\tilde{Q}_\ell - Q_\ell\|_{\text{op}} \leq \|\tilde{Q}_\ell - Q_\ell\|_{\text{F}} \leq \frac{\sqrt{kn}\varepsilon}{d} \leq \frac{1}{144\sqrt{nk}} = \tilde{\eta}.$$

We next show that the $\tilde{Q}_\ell$ satisfy the robust-design property. By [Theorem 4.5](#), with probability at least $1 - e^{-cm} - \frac{m}{2}e^{-n}$, the matrices $\frac{G_\ell}{3\sqrt{n}}$ form a $(k, \eta, m/4)$ -strong robust subspace design. Since

$$Q_\ell = \frac{3\sqrt{n}}{d} \frac{G_\ell}{3\sqrt{n}}$$

and $\frac{1}{4} \leq \frac{3\sqrt{n}}{d} \leq \frac{1}{2}$, the matrices Q_ℓ form a $(k, \frac{\eta}{4}, \frac{m}{4})$ strong robust subspace design. Now fix an arbitrary k -dimensional subspace $W \in \text{Gr}(k, n)$. By Weyl's inequality, on the event $\mathcal{E}$ we have, for every $\ell \in [m]$ and $r \in [k]$,

$$\left| \sigma_r(\tilde{Q}_\ell|_W) - \sigma_r(Q_\ell|_W) \right| \leq \|\tilde{Q}_\ell - Q_\ell\|_{\text{op}} \leq \tilde{\eta}.$$

Consequently, we find that if $\sigma_r(\tilde{Q}_\ell|_W) \leq \tilde{\eta}$, then also $\sigma_r(Q_\ell|_W) \leq 2\tilde{\eta} = \frac{\eta}{4}$. Hence

$$b_{\tilde{\eta}}(\tilde{Q}_\ell, W) \leq b_{\eta/4}(Q_\ell, W).$$

Summing over ℓ gives

$$\sum_{\ell=1}^m b_{\tilde{\eta}}(\tilde{Q}_\ell, W) \leq \sum_{\ell=1}^m b_{\eta/4}(Q_\ell, W) \leq \frac{m}{4}.$$

Moreover, in the success event of the Gaussian construction and on event $\mathcal{E}$, $\|\tilde{Q}_\ell\|_{\text{op}} \leq \frac{1}{2} + \tilde{\eta} \leq 1$ since $\|Q_\ell\|_{\text{op}} \leq \frac{1}{2}$. Thus, the family $\{\tilde{Q}_1, \dots, \tilde{Q}_m\}$ is a $(k, \tilde{\eta}, \frac{m}{4})$ strong robust subspace design, with probability at least

$$1 - e^{-cm} - \frac{m}{2}e^{-n} - e^{-(m+n)}.$$

It remains to verify the computational bounds. To see this, we simply verify that

$$1/\varepsilon = O(k\sqrt{n}) = \text{poly}(n), \quad \log(1/\delta) = O(m + n + \log(mkn)).$$

Hence, by [Lemma 4.6](#), each entry can be generated in $\text{poly}(n, m)$ time using $O(\log(1/\varepsilon) + \log \log(1/\delta)) = O(\log(nm))$ workspace and, has an exact $O(\log(nm))$ -bit rational representation. The entries can be generated sequentially, using only additional $O(\log(nm))$ -bit counters. $\square$

5 Explicit robust constructions

In the following section we describe explicit constructions of robust subspace designs. First, in [Section 5.1](#) we describe a very simple construction based on coordinate projections that yields a design of size $\binom{n}{k}$. Then, in [Section 5.2](#) we combine this construction with an idea based on sketching to reduce the size to $n^4 k^{O(k)}$.

5.1 Explicit construction based on coordinate projections

In this section, we present a simple construction of a weak robust subspace design. The construction has a good inverse polynomial robustness parameter η . However, it is of large size $m = \binom{n}{k}$. This unfortunately makes it impractical for our later complexity theoretic applications, which require polynomial-sized designs.

For $I \subseteq [n]$ with $|I| = k$, let $P_I : \mathbb{C}^n \rightarrow \mathbb{C}^k$ denote the coordinate projection onto I . The following theorem is a restatement of [Theorem 1.4](#). It uses a lemma about maximum-volume submatrices from [\[GOS⁺10\]](#), who developed it in the context of matrix approximation theory.

Theorem 5.1. *The family*

$$\mathcal{P}_k = \{P_I : I \subseteq [n], |I| = k\}$$

is an $(k, \eta, \binom{n}{k} - 1)$ -weak robust subspace design with $\eta = 1/(2\sqrt{nk})$.

Proof. It suffices to show that for any fixed k -dimensional subspace $W \subseteq \mathbb{C}^n$, there exists $I \subseteq [n]$ with $|I| = k$ such that $\sigma_{\min}(P_I|_W) \geq 1/\sqrt{nk} > \eta$. With $N \in \mathbb{C}^{n \times k}$ having orthonormal columns spanning W , we claim that this is satisfied by I maximizing

$$|\det(P_I N)|.$$

Following [\[GOS⁺10, Lemma 1\]](#), the matrix $P_I N$ is called a *dominant submatrix* of N and its entries satisfy

$$|N(P_I N)^{-1}(i, j)| \leq 1.$$

It remains to note that

$$\sigma_{\min}(P_I|_W) = \sigma_{\min}(P_I N) = 1/\|(P_I N)^{-1}\|_{\text{op}}$$

and that

$$\|(P_I N)^{-1}\|_{\text{op}} = \|N(P_I N)^{-1}\|_{\text{op}} \leq \|N(P_I N)^{-1}\|_F \leq \sqrt{nk}. \quad \square$$

5.2 Improved explicit construction using sketching

We first state the main result of this section.

Theorem 5.2 (Explicit robust family). *For every $1 \leq k \leq n$ there is an explicitly computable family $\mathcal{Q}_{n,k} \subseteq \mathbb{Q}^{k \times n}$ of size $m = O(n^4 k^{O(k)})$ which is a (k, η, Δ) -weak robust subspace design, with parameters*

$$\eta = \Omega\left(\frac{1}{k^{3/2}\sqrt{n}}\right), \quad \Delta < m.$$

Moreover, every entry is rational of polynomial bit length in n and enumerable in space $O(\log n + k \log k)$.

The proof uses a composition of two families. The first family maps into $O(k^2)$ dimensions rather than into $\mathbb{C}^k$. Effectively, it is a sketching family that robustly embeds the original n dimensions into $O(k^2)$ dimensions. We then compose this family with our explicit coordinate-projection design to reduce the output dimension to k , yielding a valid robust subspace design.

5.2.1 Robust sketching family

Our improved explicit construction goes through an intermediate dimension reduction phase based on sketching. To describe this step, we use the following terminology. It is effectively a relaxation of a weak robust subspace design where the maps are allowed to have rank $d > k$.

Definition 5.3 (Weak robust sketching family). Let $d \geq k$. A family of linear maps

$$A_1, \dots, A_m : \mathbb{C}^n \rightarrow \mathbb{C}^d$$

is a (k, η, Δ) -weak robust sketching family if $\Delta < m$, $\|A_i\|_{\text{op}} \leq 1$ for every i and, for every k -dimensional subspace $W \subseteq \mathbb{C}^n$,

$$|\{i \in [m] : \sigma_{\min}(A_i|_W) \leq \eta\}| \leq \Delta.$$

When $d = k$, this is exactly a (k, η, Δ) -weak robust subspace design.

We next show that an explicit, polynomial-sized family of sketching matrices yields a weak robust sketching family (hence the name). To establish this, we adapt some ideas from [CW09, Section 2.1, Eq. (1)], and [Sar06] in our setting. Our main technical tool is based on the *fourth moment method*.

Definition 5.4 (Tug-of-war matrix [Sar06]). A random matrix $S \in \mathbb{R}^{r \times n}$ is called a t -tug-of-war matrix if its entries are t -wise independent and, for every $a \in [r]$ and $i \in [n]$,

$$\Pr\left(S_{a,i} = \pm \frac{1}{\sqrt{r}}\right) = \frac{1}{2}.$$

Sarlós proves the corresponding estimate below over the reals [Sar06]. The same fourth-moment calculation gives the following complex version; we include the short argument for completeness.

Lemma 5.5 (Expected product of tug-of-war matrices). *Let $A, B \in \mathbb{C}^{n \times k}$ and let $S \in \mathbb{R}^{r \times n}$ be a 4-tug-of-war matrix. Then*

$$\mathbb{E}[\|B^* S^* S A - B^* A\|_F^2] \leq \frac{2}{r} \|A\|_F^2 \|B\|_F^2.$$

Proof. For $x, y \in \mathbb{C}^n$, write

$$Z := \langle Sx, Sy \rangle - \langle x, y \rangle.$$

If the a -th row of $\sqrt{r} S$ is $\varepsilon_a = (\varepsilon_{a1}, \dots, \varepsilon_{an})$, then

$$Z = \frac{1}{r} \sum_{a=1}^r \sum_{i \neq j} \varepsilon_{ai} \varepsilon_{aj} \overline{x_i} y_j.$$

By 4-wise independence, the cross terms corresponding to distinct rows vanish in the second moment. For one row, grouping the terms indexed by $\{i, j\}$ and using $|u + v|^2 \leq 2(|u|^2 + |v|^2)$ gives

$$\mathbb{E}|Z|^2 \leq \frac{2}{r} \|x\|^2 \|y\|^2.$$

Apply this estimate to every pair of columns of A and B and sum the resulting bounds over the entries of $B^* S^* S A - B^* A$. $\square$

Here, we will just need to consider the case $A = B = N \in \mathbb{C}^{n \times k}$ for some N satisfying $N^*N = I_k$. From this bound, we can deduce a lower bound on the minimum singular value with high probability.

Lemma 5.6. *Let $N \in \mathbb{C}^{n \times k}$ have orthonormal columns, let $S \in \mathbb{R}^{r \times n}$ be a 4-tug-of-war random matrix, and let $0 \leq \beta < 1$. Then*

$$\Pr[\sigma_{\min}(SN) \leq \sqrt{\beta}] \leq \frac{2k^2}{(1-\beta)^2 r},$$

and $\|S\|_{\text{op}} \leq \sqrt{n}$.

Proof. The second bound follows from $\|S\|_{\text{op}}^2 \leq \|S\|_F^2 = rn \cdot \frac{1}{r} = n$.

For the first bound, set

$$G := N^* S^* S N = (SN)^*(SN) \in \mathbb{C}^{k \times k}.$$

By Lemma 5.5 with $A = B = N$,

$$\mathbb{E}\|G - I_k\|_F^2 \leq \frac{2\|N\|_F^4}{r} = \frac{2k^2}{r}.$$

Hence, for every $a > 0$, Markov's inequality gives

$$\Pr[\|G - I_k\|_{\text{op}}^2 \geq a] \leq \Pr[\|G - I_k\|_F^2 \geq a] \leq \frac{2k^2}{ar}. \quad (3)$$

On the other hand, for every $x \in \mathbb{C}^k$, $x^* G x \geq (1 - \|G - I_k\|_{\text{op}}) \|x\|^2$, and therefore

$$\sigma_{\min}(SN)^2 = \lambda_{\min}(G) \geq 1 - \|G - I_k\|_{\text{op}}.$$

Consequently, by Equation (3),

$$\Pr[\sigma_{\min}(SN)^2 \leq \beta] \leq \Pr[\|G - I_k\|_{\text{op}} \geq 1 - \beta] = \Pr[\|G - I_k\|_{\text{op}}^2 \geq (1 - \beta)^2] \leq \frac{2k^2}{(1 - \beta)^2 r}. \quad \square$$

Having established the robust embedding properties of a 4-tug-of-war matrix ensemble, it remains to notice that there exists a logspace-enumerable and polynomial-sized instantiation of this ensemble. It follows from a standard construction of Alon–Babai–Itai [ABI86, Proposition 6.3] giving a logspace-enumerable uniform sample space $\mathcal{T} = \{t \in \{\pm 1\}^M\}$ of size $|\mathcal{T}| \leq 16M^4$ with vectors whose coordinates are 4-wise independent unbiased signs. More precisely, let $\ell = \lceil \log_2 M \rceil$ and $q = 2^\ell$. Taking $d = 4$ and this choice of q in [ABI86, Proposition 6.3] yields 4-wise independent random variables uniformly distributed over $\mathbb{F}_q$ (equivalently, over $[q]$ after fixing an identification). To obtain unbiased signs, we apply the standard field trace

$$\text{Tr}_{\mathbb{F}_q/\mathbb{F}_2} : \mathbb{F}_q \rightarrow \mathbb{F}_2, \quad \text{Tr}(z) = z + z^2 + \dots + z^{2^{\ell-1}},$$

and map $b \in \mathbb{F}_2$ to $(-1)^b$; see, e.g., [DPT24]. Since the trace is a surjective $\mathbb{F}_2$ -linear map, this preserves 4-wise independence and gives unbiased signs.

At this point, we have all the ingredients to construct and analyze the sketching family.

Theorem 5.7 (Robust sketching family). *For every $1 \leq k \leq n$, let $r := 16k^2$. There is an indexed family*

$$S_1, \dots, S_m \in \mathbb{Q}^{r \times n}, \quad m \leq 16(rn)^4,$$

which is a $(k, \eta, m/2)$ -weak robust sketching family, with $\eta = \Omega(1/\sqrt{n})$. The family is enumerable in logspace and its coefficients can be computed in logspace.

Proof. By [ABI86], we can generate in logspace an indexed family of sign matrices

$$E_1, \dots, E_m \in \{\pm 1\}^{r \times n}, \quad m = O((rn)^4),$$

such that the rn coordinates of a uniformly chosen E_i are 4-wise independent unbiased signs. Recalling that $r = 16k^2$, and we can then define $S_i^{(0)} := \frac{1}{4k} E_i$ so that a uniformly chosen $S_i^{(0)}$ is a 4-tug-of-war matrix.

Let d be the smallest power of two satisfying $\sqrt{n} \leq d < 2\sqrt{n}$. Define $S_i := \frac{1}{d} S_i^{(0)}$. Every entry of S_i is rational of polynomial bit length, and

$$\|S_i\|_{\text{op}} \leq \frac{\sqrt{n}}{d} \leq 1.$$

Now fix a k -dimensional subspace $W \subseteq \mathbb{C}^n$ and let $N \in \mathbb{C}^{n \times k}$ have orthonormal columns spanning W . Applying Lemma 5.6 with $\beta = 1/2$ and $r = 16k^2$,

$$\Pr \left[\sigma_{\min}(S^{(0)}N) \leq \frac{1}{\sqrt{2}} \right] \leq \frac{2k^2}{(1/2)^2 16k^2} = \frac{1}{2}.$$

Therefore at most half of the indices satisfy

$$\sigma_{\min}(S_i N) \leq \frac{1}{\sqrt{2}d}.$$

Since $d < 2\sqrt{n}$, we may take $\eta := \frac{1}{\sqrt{2}d} = \Omega(1/\sqrt{n})$, establishing that the family is a $(k, \eta, m/2)$ -weak robust sketching family. $\square$

5.3 Composing families

For the rank-completion application, the final maps must have output dimension k (and hence rank at most k), whereas the sketching family has output dimension $r = \Theta(k^2)$. We now combine the sketching family with the coordinate construction in order to reduce the output dimension from r to k .

Theorem 5.8 (Composition). *Let $A_1, \dots, A_{m_A} \in \mathbb{C}^{n_{AB} \times n}$ be an (m_A) -sized (k, η_A, Δ_A) -weak robust sketching family, and let $B_1, \dots, B_{m_B} \in \mathbb{C}^{k \times n_{AB}}$ be an (m_B) -sized (k, η_B, Δ_B) -weak robust subspace design, with $\Delta_A < m_A$ and $\Delta_B < m_B$.*

Then

$$\{B_j A_i : i \in [m_A], j \in [m_B]\}$$

is a $(k, \eta_{AB}, \Delta_{AB})$ -weak robust subspace design, where $m_{AB} = m_A m_B$ and $\eta_{AB} = \eta_A \eta_B$. Moreover,

$$\Delta_{AB} = m_A \Delta_B + m_B \Delta_A - \Delta_A \Delta_B < m_{AB}.$$

If the two families can be enumerated and their entries computed in logspace, then the composed family has the same property.

Proof. The size follows directly, and the operator norms are bounded by 1 because the operator norm is submultiplicative.

Fix a k -dimensional subspace $W \subseteq \mathbb{C}^n$. There are at least $m_A - \Delta_A$ indices i such that

$$\sigma_{\min}(A_i|_W) > \eta_A.$$

Fix such an i . Then $A_i|_W$ is injective, so $A_i W$ is a k -dimensional subspace of $\mathbb{C}^{n_{AB}}$. By the design property of the B_j 's, at least $m_B - \Delta_B$ indices j satisfy

$$\sigma_{\min}(B_j|_{A_i W}) > \eta_B.$$

For every such pair (i, j) and every non-zero $x \in W$,

$$\|B_j A_i x\| \geq \sigma_{\min}(B_j|_{A_i W}) \|A_i x\| \geq \sigma_{\min}(B_j|_{A_i W}) \sigma_{\min}(A_i|_W) \|x\| > \eta_B \eta_A \|x\|.$$

Hence there are at least

$$(m_A - \Delta_A)(m_B - \Delta_B)$$

good pairs. Equivalently, the number of bad pairs is at most

$$m_A m_B - (m_A - \Delta_A)(m_B - \Delta_B) = m_A \Delta_B + m_B \Delta_A - \Delta_A \Delta_B.$$

This proves the claimed parameters.

For the final assertion, the pairs (i, j) can be enumerated in logspace, and each entry of $B_j A_i$ can be computed entrywise in logspace from the corresponding entries of A_i and B_j . $\square$

We can now prove our main theorem.

Proof of Theorem 5.2. We compose the two constructions above. First, Theorem 5.7 gives a robust sketching family $A_i : \mathbb{C}^n \rightarrow \mathbb{C}^{n_{AB}}$ with

$$n_{AB} = r = 16k^2, \quad m_A = O((rn)^4), \quad \eta_A = \Omega(1/\sqrt{n}), \quad \Delta_A = m_A/2.$$

Second, Theorem 1.4, applied in ambient dimension n_{AB} , gives a coordinate subspace design $B_j : \mathbb{C}^{n_{AB}} \rightarrow \mathbb{C}^k$ with

$$m_B = \binom{n_{AB}}{k}, \quad \eta_B = \frac{1}{2\sqrt{n_{AB}k}}, \quad \Delta_B = m_B - 1.$$

By Theorem 5.8, their compositions form a weak robust subspace design with

$$m = m_A m_B = O\left((k^2 n)^4 \binom{16k^2}{k}\right) = O\left(n^4 k^{O(k)}\right),$$

where we used $\binom{16k^2}{k} \leq (16ek)^k$. Moreover,

$$\eta = \eta_A \eta_B = \Omega\left(\frac{1}{\sqrt{n}} \cdot \frac{1}{\sqrt{16k^3}}\right) = \Omega\left(\frac{1}{k^{3/2}\sqrt{n}}\right).$$

Finally, $\Delta = m_A \Delta_B + m_B \Delta_A - \Delta_A \Delta_B < m_A m_B = m$.

All matrices have rational entries of polynomial bit length, and the construction is enumerable in space $O(\log m) = O(\log n + k \log k)$ by the corresponding properties of the two individual families and the composition theorem. $\square$

6 Isolating a unique small quantum witness and other applications

In this section we show our main application of robust subspace designs and kernel peeling, to quantum witness isolation, as well as new linear algebra containment results. Before getting into this, let us recall some basic complexity theory notation:

We use L , RL , $coRL$, and BPL for deterministic logspace and its randomized one-sided- and two-sided-error variants. These machines have read-only access to the n -bit input, and an additional work tape where they may only use $O(\log n)$ tape cells. Randomized machines run in polynomial time, and error probabilities are at most $1/3$. A family of circuits or matrices is *logspace uniform* if a deterministic logspace algorithm can output its description on a write-only uni-directional tape.

In an oracle class such as $RL^{\mathcal{O}}$, the machine has access to a fixed problem in $\mathcal{O}$ and may write a polynomial-length query on a write-only oracle tape that is not counted as workspace. If the oracle problem has a promise, answers to queries outside that promise may be arbitrary. All claimed correctness and error guarantees must hold regardless of those answers.

6.1 Quantum space-bounded complexity classes

Let $s(n) = \Omega(\log n)$ be a space bound, and let $c, k : \mathbb{N} \rightarrow [0, 1]$ satisfy $c(n) > k(n)$ for every n . The functions c and k are the completeness and soundness thresholds, respectively. We assume that s , c , and k are computable in space $O(s(n))$. All unitary circuits use a fixed finite universal gate set with entries in $\mathbb{Q}(i)$.

Following [FR21], we use space-uniform families of unitary quantum circuits. Specifically, a family $\{Q_x\}$ is space- s uniform if a deterministic Turing machine, on input $x \in \{0, 1\}^n$, runs in space $O(s(n))$ and outputs the description of Q_x . Each circuit acts on $O(s(n))$ qubits and consists of at most $2^{O(s(n))}$ gates.

Definition 6.1 (Quantum space-bounded computation). A promise problem $L = (L_{\text{yes}}, L_{\text{no}})$ belongs to $\text{QUSPACE}_{c,k}(s(n))$ if there exists a space- s uniform family of unitary quantum circuits $\{Q_x\}$ such that

- if $x \in L_{\text{yes}}$, then $\Pr[Q_x \text{ accepts}] \geq c(n)$,
- if $x \in L_{\text{no}}$, then $\Pr[Q_x \text{ accepts}] \leq k(n)$.

Here Q_x acts on $a_n = O(s(n))$ qubits initialized in $|0\rangle^{\otimes a_n}$ and accepts if a final computational-basis measurement of the first qubit yields outcome 1.

For the witness-based classes below, the circuit Q_x acts on a_n auxiliary qubits and w_n witness qubits, where $a_n + w_n = O(s(n))$. Given a witness $|\psi\rangle \in \mathbb{C}^{2^{w_n}}$, the circuit is initialized to $|0\rangle^{\otimes a_n} |\psi\rangle$, and it accepts if the first of the auxiliary qubits is measured to be in state $|1\rangle$ at the end of the computation. Throughout, all witness states are understood to be normalized.

Definition 6.2 (Quantum space-bounded Merlin–Arthur protocols). A promise problem $L = (L_{\text{yes}}, L_{\text{no}})$ belongs to $\text{QUMASPACE}_{c,k}(s(n))$ if there exists a space- s uniform family of unitary verification circuits $\{Q_x\}$ such that

- if $x \in L_{\text{yes}}$, then there exists a witness $|\psi\rangle \in \mathbb{C}^{2^{w_n}}$ with $\Pr[Q_x \text{ accepts } |\psi\rangle] \geq c(n)$,
- if $x \in L_{\text{no}}$, then every witness $|\psi\rangle \in \mathbb{C}^{2^{w_n}}$ satisfies $\Pr[Q_x \text{ accepts } |\psi\rangle] \leq k(n)$.

We next introduce space-bounded analogues of the classes FewQMA and UQMA introduced in [JKK⁺11]. These classes mirror the classical restrictions underlying FewP (introduced in [All86]) and UP (introduced in [Val76]), and their logspace counterparts FewL and UL (introduced in [BJLR91]): the number of accepting computation paths is replaced by the dimension of an accepting witness subspace W_x , which is then restricted to either polynomial dimension (FewQMA) or dimension 1 (UQMA).

In the quantum setting, we additionally require an acceptance gap between the accepting subspace and its orthogonal complement. To distinguish this requirement from soundness on no-instances, let $\ell : \mathbb{N} \rightarrow [0, 1]$ be computable in space $O(s(n))$, with $\ell(n) < c(n)$. This parameter bounds the acceptance probability of witnesses orthogonal to the accepting subspace on yes-instances.

Definition 6.3 (Few quantum witnesses). A promise problem $L = (L_{\text{yes}}, L_{\text{no}})$ belongs to $\text{FewQUMASPACE}_{c,\ell,k}(s(n))$ if there exist a polynomial q and a space- s uniform family of unitary verification circuits $\{Q_x\}$ such that

- if $x \in L_{\text{yes}}$, then there exists a subspace $W_x \subseteq \mathbb{C}^{2^{w_n}}$ with $1 \leq \dim W_x \leq q(n)$ such that

$$\Pr[Q_x \text{ accepts } |\psi\rangle] \geq c(n) \quad \text{for every } |\psi\rangle \in W_x,$$

and

$$\Pr[Q_x \text{ accepts } |\phi\rangle] \leq \ell(n) \quad \text{for every } |\phi\rangle \in W_x^\perp,$$

- if $x \in L_{\text{no}}$, then every witness $|\psi\rangle \in \mathbb{C}^{2^{w_n}}$ satisfies $\Pr[Q_x \text{ accepts } |\psi\rangle] \leq k(n)$.

Definition 6.4 (A unique quantum witness). A promise problem $L = (L_{\text{yes}}, L_{\text{no}})$ belongs to $\text{UQUMASPACE}_{c,\ell,k}(s(n))$ if there exists a space- s uniform family of unitary verification circuits $\{Q_x\}$ such that

- if $x \in L_{\text{yes}}$, then there exists a one-dimensional subspace $W_x \subseteq \mathbb{C}^{2^{w_n}}$ such that

$$\Pr[Q_x \text{ accepts } |\psi\rangle] \geq c(n) \quad \text{for every } |\psi\rangle \in W_x,$$

and

$$\Pr[Q_x \text{ accepts } |\phi\rangle] \leq \ell(n) \quad \text{for every } |\phi\rangle \in W_x^\perp,$$

- if $x \in L_{\text{no}}$, then every witness $|\psi\rangle \in \mathbb{C}^{2^{w_n}}$ satisfies $\Pr[Q_x \text{ accepts } |\psi\rangle] \leq k(n)$.

We will focus on the special cases:

$$\begin{aligned} \text{BQUL} &:= \text{QUSPACE}_{2/3, 1/3}(\log n), \\ \text{QUMAL} &:= \text{QUMASPACE}_{2/3, 1/3}(\log n), \\ \text{QUMAL}_1 &:= \text{QUMASPACE}_{1, 1/3}(\log n), \\ \text{FewQUMAL}_1 &:= \text{FewQUMASPACE}_{1, 1/3, 1/3}(\log n), \\ \text{UQUMAL}_1 &:= \text{UQUMASPACE}_{1, 1/3, 1/3}(\log n). \end{aligned}$$

Here we note that, for logarithmic space, the polynomial upper bound on $\dim W_x$ is automatic, since the witness register has polynomial dimension. The only additional restriction in FewQUMAL_1 compared to QUMAL_1 is therefore the acceptance bound on $W_x^\perp$ on yes-instances.

Allowing general quantum channels, including intermediate measurements, gives the corresponding classes without the unitary subscript U. Fefferman and Remscrem [FR21, Lemma 18] showed that

$$\text{BQL} = \text{QMAL} = \text{BQUL} = \text{QUMAL}.$$

Consequently, the bounded-error versions of our Few and Unique classes also coincide with BQ_UL . The key idea in [FR21] behind removing a logarithmic-size quantum witness is to first amplify the verifier using space-efficient error reduction (see e.g. [FKL⁺16]) and then replace the witness by the maximally mixed state. Since the witness space has polynomial dimension, this leaves an inverse-polynomial acceptance gap, which can again be amplified to obtain a bounded-error algorithm. This procedure does not, however, preserve perfect completeness. Our focus is therefore on the perfect-completeness witness classes, for which we have

$$\text{UQ}_\text{UMAL}_1 \subseteq \text{FewQ}_\text{UMAL}_1 \subseteq \text{Q}_\text{UMAL}_1 \subseteq \text{Q}_\text{UMAL}.$$

Note that the bounded-error equivalence between space-bounded unitary circuits and general quantum channels does not establish $\text{Q}_\text{UMAL}_1 = \text{QMAL}_1$.

One motivation for studying Q_UMAL_1 is that it completely characterizes the complexity of testing the singularity of well-conditioned matrices [FR21]. This makes it a natural quantum counterpart to $\text{C}_=\text{L}$, the classical counting class characterized by testing whether a general (potentially ill-conditioned) matrix is singular (see e.g. [ABO99]). We show that $\text{FewQ}_\text{UMAL}_1$ and UQ_UMAL_1 likewise admit complete problems obtained by imposing additional spectral-gap and kernel-dimension promises. We will use this characterization to prove our main application, Theorem 1.1. To state these characterizations, we introduce the following promise problems:

κ -conditioned-SINGULAR	<i>Input:</i> $A \in \mathbb{C}^{n \times n}$ with $\ A\ _\text{op} \leq 1$. <i>Yes:</i> $\dim \ker(A) \geq 1$. <i>No:</i> $\dim \ker(A) = 0$ and $\sigma_\text{min}(A) \geq 1/\kappa(n)$.
κ -conditioned-FewSINGULAR	<i>Input:</i> $A \in \mathbb{C}^{n \times n}$ with $\ A\ _\text{op} \leq 1$ and $\sigma_\text{min}^+(A) \geq 1/\kappa(n)$. <i>Yes:</i> $\dim \ker(A) \geq 1$. <i>No:</i> $\dim \ker(A) = 0$.
κ -conditioned-UniqueSINGULAR	<i>Input:</i> $A \in \mathbb{C}^{n \times n}$ with $\ A\ _\text{op} \leq 1$ and $\sigma_\text{min}^+(A) \geq 1/\kappa(n)$. <i>Yes:</i> $\dim \ker(A) = 1$. <i>No:</i> $\dim \ker(A) = 0$.

We write pc-SINGULAR , pc-FewSINGULAR and pc-UniqueSINGULAR for the above problems with polynomial conditioning parameter, i.e. for the case $\kappa(n) = \text{poly}(n)$.

Lemma 6.5 (Completeness of poly-conditioned singularity). *Under deterministic logspace many-one reductions,*

- (i) pc-SINGULAR is Q_UMAL_1 -complete (proven in [FR21, Lemma 35]),
- (ii) pc-FewSINGULAR is $\text{FewQ}_\text{UMAL}_1$ -complete, and
- (iii) pc-UniqueSINGULAR is UQ_UMAL_1 -complete.

The hardness proofs for parts (ii) and (iii) refine the circuit-to-Hamiltonian argument used for part (i) in [FR21, Lemma 35]. The additional observations are that the construction preserves the dimension of the perfectly accepting witness subspace and that an acceptance gap on its orthogonal complement yields a spectral gap also on yes-instances. The corresponding containment results follow from the kernel-verification procedure also invoked in [FR21, Lemma 35]. We defer the proofs to Section C.2.

6.2 Isolating a unique small quantum witness

We can now finally prove our main application of robust subspace designs from the introduction:

Theorem 1.1. $\text{FewQUMAL}_1 \subseteq \text{RL}^{\text{UQUMAL}_1}$.

Proof. We prove that pc-FewSINGULAR can be decided by a randomized logspace machine with access to an oracle in UQUMAL_1 . The result then follows since pc-FewSINGULAR is FewQUMAL_1 -complete by Lemma 6.5.

The oracle problem: We introduce an auxiliary promise problem $\text{pc-UniqueSINGULAR}'$. Its input consists of a positive semidefinite matrix $H \in \mathbb{C}^{n \times n}$ and a matrix $Q \in \mathbb{C}^{k' \times n}$, where $0 \leq k' \leq n$. Writing $M = H + Q^*Q$, the input promise is that for all entries i, j

$$|Q(i, j)| \leq 1, \quad \|H\|_{\text{op}} \leq 1, \quad \sigma_{\min}^+(M) \geq 1/\text{poly}(n).$$

The yes-instances satisfy $\dim \ker(M) = 1$, and the no-instances satisfy $\dim \ker(M) = 0$.

This problem belongs to UQUMAL_1 . Indeed, given (H, Q) on the input tape, a deterministic logspace machine can compute the entries of

$$\widehat{M} := \frac{H + Q^*Q}{1 + nk'}.$$

Then $\|\widehat{M}\|_{\text{op}} \leq 1$ since $\|Q^*Q\|_{\text{op}} \leq nk'$, $\ker(\widehat{M}) = \ker(M)$, and $\sigma_{\min}^+(\widehat{M}) \geq 1/\text{poly}(n)$, yielding an immediate reduction to pc-UniqueSINGULAR , which belongs to UQUMAL_1 by Lemma 6.5. The purpose of supplying H and Q separately is that the oracle can reread the entries of Q to obtain easy access to Q^*Q , whereas the randomized reduction only needs to stream the entries of Q once.

The randomized reduction: Let $A \in \mathbb{C}^{n \times n}$ be a promised input to pc-FewSINGULAR , and put

$$H = A^*A, \quad k = \dim \ker(A).$$

Then

$$H \succeq 0, \quad \|H\|_{\text{op}} \leq 1, \quad \ker(H) = \ker(A), \quad \lambda_{\min}^+(H) = \sigma_{\min}^+(A)^2 \geq 1/\text{poly}(n).$$

For each $k' \in \{1, \dots, n-1\}$, we use Theorem 4.7 to generate a strong robust subspace design in randomized logspace of the form

$$\mathcal{Q}^{(k')} = \{Q_1^{(k')}, \dots, Q_{m_{k'}}^{(k')}\} \subseteq \mathbb{C}^{k' \times n},$$

where $m_{k'} = \text{poly}(n)$ and the robustness parameter $\eta_{k'}$ is inverse-polynomial. The failure bound in Theorem 4.7 ensures this happens with probability at least $2/3$ for all sufficiently large n . In particular every entry of each $Q_i^{(k')}$ is of modulus at most 1 in that case. Smaller dimensions can be handled by padding A with an identity block. For $k' = 0$, take the family consisting of just the zero map.

For each generated matrix $Q_i^{(k')}$, we write $(H, Q_i^{(k')})$ to the oracle tape, except we replace entries of modulus greater than 1 by 0. We then query $\text{pc-UniqueSINGULAR}'$ with $(H, Q_i^{(k')})$ as input. We accept if any query returns yes, and reject otherwise. It is easy to see that the reduction uses logarithmic workspace and polynomial time: The entries of $H = A^*A$ can be computed in deterministic logspace from the input, and each $Q_i^{(k')}$ is streamed directly to the oracle tape as it is generated. Beyond the workspace needed for these procedures, only loop indices need to be stored. Hence, the total space usage is only $O(\log n)$.

Soundness: Suppose that A is a no-instance. Then H is positive definite and

$$\lambda_{\min}(H) = \sigma_{\min}(A)^2 \geq 1/\text{poly}(n).$$

For every generated matrix Q , regardless of whether the robust-design construction succeeds, $H + Q^*Q \succeq H$. Hence $H + Q^*Q$ is invertible and

$$\lambda_{\min}(H + Q^*Q) \geq \lambda_{\min}(H) \geq 1/\text{poly}(n).$$

Every query is therefore a promised no-instance of $\text{pc-UniqueSINGULAR}'$, and the algorithm rejects with probability 1.

Completeness: Suppose that A is a yes-instance, so $k \geq 1$. If $k = 1$, the query with $Q = 0$ is a promised yes-instance, and the algorithm accepts.

If $k \geq 2$, consider the family generated for $k' = k - 1$. With probability at least $2/3$, this family has the required robust-design property. On this event, the kernel peeling lemma, [Lemma 3.6](#), specifically its Hermitian version [Remark 3.7](#), gives some $Q_i^{(k-1)}$ such that

$$\dim \ker(H + (Q_i^{(k-1)})^* Q_i^{(k-1)}) = 1$$

and

$$\lambda_{\min}^+(H + (Q_i^{(k-1)})^* Q_i^{(k-1)}) \geq 1/\text{poly}(n).$$

This query is a promised yes-instance of $\text{pc-UniqueSINGULAR}'$, so the algorithm accepts. Other queries may lie outside the oracle promise, but their answers cannot prevent acceptance. Thus the algorithm accepts with probability at least $2/3$. $\square$

Since the only place where we use randomness in the reduction above is the construction of the robust subspace designs, we immediately find the following:

Corollary 6.6. *If there exists a logspace-uniform weak robust subspace design of polynomial size with inverse polynomial robustness parameter for every ambient subspace dimension $1 \leq k < n$, then $\text{FewQUMAL}_1 \subseteq \text{L}^{\text{UQUMAL}_1}$.*

Instead of computing the robust subspace designs used in [Theorem 1.1](#), we can also supply it as an advice string to obtain:

Corollary 6.7. $\text{FewQUMAL}_1 \subseteq \text{L}^{\text{UQUMAL}_1}/\text{poly}$.

Here we use standard non-uniform logspace notation. A promise problem $L = (L_{\text{yes}}, L_{\text{no}})$ is in L/poly if there exist advice strings $a(n) \in \{0, 1\}^{\text{poly}(n)}$, and a deterministic logspace machine M such that, for every $x \in L_{\text{yes}} \cup L_{\text{no}}$, $x \in L$ if and only if $M(x, a(|x|))$ accepts. Similarly, $\text{L}^{\text{UQUMAL}_1}/\text{poly}$, or $\text{L}^{\text{QUMAL}_1}/\text{poly}$ denote deterministic logspace with polynomial-size advice and oracle access to a promise problem in UQUMAL_1 or QUMAL_1 , respectively.

6.3 Other applications

In this subsection we show that robust and non-robust subspace designs and our kernel peeling lemma help us to solve linear algebra problems such as nullity testing and linear system feasibility in logspace. The non-robust subspace designs based on the Wronskian construction from [Theorem 3.8](#) allow us to put nullity testing in C=L and linear system feasibility in $\text{L}^{\text{C=L}}$, recovering the bounds from [\[ABO99\]](#). Robust subspace design allow us to put well-conditioned versions of nullity testing

and feasibility in $\text{coRL}^{\text{QUMAL}_1}$ and $\text{BPL}^{\text{QUMAL}_1}$, respectively. Before getting into the proofs, let us first formalize the problem statements:

NULL_k	<i>Input:</i> $A \in \mathbb{C}^{n \times n}$. <i>Yes:</i> $\dim \ker(A) \geq k$. <i>No:</i> $\dim \ker(A) < k$.
FSLE	<i>Input:</i> $A \in \mathbb{C}^{n \times n}$ and $b \in \mathbb{C}^n$. <i>Yes:</i> $\exists x \in \mathbb{C}^n : Ax = b$. <i>No:</i> $\forall x \in \mathbb{C}^n : Ax \neq b$.
κ -conditioned- NULL_k	<i>Input:</i> $A \in \mathbb{C}^{n \times n}$ with $\ A\ _{\text{op}} \leq 1$. <i>Yes:</i> $\dim \ker(A) \geq k$. <i>No:</i> $\dim \ker(A) < k$ and $\sigma_k(A) \geq 1/\kappa(n)$.
(κ, ε) -conditioned-FSLE	<i>Input:</i> $A \in \mathbb{C}^{n \times n}$ and $b \in \mathbb{C}^n$ with $\ A\ _{\text{op}}, \ b\ \leq 1$ and $\sigma_{\min}^+(A) \geq 1/\kappa(n)$. <i>Yes:</i> $\exists x \in \mathbb{C}^n : Ax = b$. <i>No:</i> $\forall x \in \mathbb{C}^n : \ Ax - b\ \geq \varepsilon(n)$.

We write pc-NULL_k and pc-FSLE for κ -conditioned- NULL_k and (κ, ε) -conditioned-FSLE, respectively, with $\kappa(n), \varepsilon(n)^{-1} = \text{poly}(n)$. Furthermore, we write NULL without an index if k is given as part of the input. We define κ -conditioned- NULL , and pc-NULL in the natural way.

Remark 6.8. The above problems NULL and FSLE are formulated for matrices with integer entries from $\mathbb{Z}$ in [ABO99]. They remark that their computational complexities remain unchanged when one allows rational entries from $\mathbb{Q}$ instead. We make the further observation that their complexity further remains unchanged if one allows Gaussian rationals $\mathbb{Q}(i)$, i.e. complex numbers of the form $(a + ib)/c$ for integers a, b, c . The reason is that we can use the following realification of a complex matrix $M = A + iB$ with $A, B \in \mathbb{Q}^{n \times n}$:

$$M = A + iB \quad \mapsto \quad \mathcal{R}(M) := \begin{pmatrix} A & -B \\ B & A \end{pmatrix}.$$

Then

$$\text{rank}_{\mathbb{R}}(\mathcal{R}(M)) = 2 \text{rank}_{\mathbb{C}}(M),$$

and, for any complex vector $b = s + it \in \mathbb{Q}(i)^n$,

$$Mx = b \quad \text{if and only if} \quad \mathcal{R}(M) \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} s \\ t \end{pmatrix},$$

where $u = \Re(x)$ and $v = \Im(x)$.

Throughout, we always assume that the entries of the input matrices are complex rationals of polynomial bit complexity.

6.3.1 Classical upper bounds for nullity testing and feasibility

We use ordinary subspace designs and kernel peeling to give an alternative reduction from NULL to SINGULAR . This recovers the classical upper bounds of Allender, Beals, and Ogihara [ABO99] for

these problems, whose original proof uses a sequence of reductions including Berkowitz's algorithm [Ber84].

Recall that a language L belongs to $\mathbf{C=L}$ if there exists a nondeterministic logspace machine M such that, for every input x , $x \in L$ if and only if M has equally many accepting and rejecting computation paths on x . We use the following standard facts.

Fact 6.9 ([AO94, MV97, ABO99]). **SINGULAR** is $\mathbf{C=L}$ -complete under deterministic logspace many-one reductions, and $\mathbf{C=L}$ is closed under deterministic logspace conjunctive truth-table reductions.

Here a conjunctive truth-table reduction from a language (or promise problem) A to B outputs a polynomial-length list of queries to B so that an input is in A if and only if every query is a yes-instance. We write $A \preceq_{\text{ctt}}^L B$.

Lemma 6.10. $\text{NULL} \preceq_{\text{ctt}}^L \text{SINGULAR}$.

Proof. Let (A, k) be an instance of **NULL**, where $A \in \mathbb{C}^{n \times n}$ is the matrix which we wish to decide whether $\dim \ker(A) \geq k$. For each $1 \leq k' < k$, take the Wronskian maps

$$Q_1^{(k')}, \dots, Q_{m_{k'}}^{(k')} \in \mathbb{C}^{k' \times n}$$

from Theorem 3.8, evaluated at $m_{k'} = k'(n - k') + 1$ distinct points $T_{k'} = \{0, \dots, m_{k'} - 1\}$. For $k' = 0$, take the zero map. Output singularity queries for each matrix

$$A + (Q_i^{(k')})^* Q_j^{(k')}, \quad 0 \leq k' < k, \quad i, j \in [m_{k'}].$$

There are polynomially many queries, and their entries are computable in deterministic logspace by the logspace uniformity of the Wronskian construction and since matrix multiplication is in deterministic logspace.

If $\dim \ker(A) \geq k$, every perturbation has rank at most $k' < k$, and hence

$$\text{rank}\left(A + (Q_i^{(k')})^* Q_j^{(k')}\right) \leq \text{rank}(A) + k' \leq n - k + k' < n.$$

Thus every query is singular. Conversely, if $1 \leq d := \dim \ker(A) < k$, the kernel peeling argument in Lemma 3.6, applied with $k' = d$, yields indices i, j for which

$$\dim \ker\left(A + (Q_i^{(d)})^* Q_j^{(d)}\right) = 0.$$

For $d = 0$, the query A itself is nonsingular. Thus at least one query is nonsingular, proving the reduction. $\square$

Consequently, $\text{NULL} \in \mathbf{C=L}$, recovering [ABO99, Proposition 2.5]. This also gives an alternative route to $\text{FSLE} \in \mathbf{L}^{\mathbf{C=L}}$ through [ABO99, Corollary 2.15]: for a linear system $Ax = b$,

$$Ax = b \text{ is feasible} \iff \text{rank}(A) = \text{rank}([A \ b]).$$

Both ranks can be determined in deterministic logspace using nullity queries.

6.3.2 Quantum upper bounds for nullity testing and feasibility

We first apply randomized robust subspace designs to conditioned nullity testing. The reduction has one-sided error: on a yes-instance all oracle queries are singular, while on a no-instance a successful design produces an invertible query.

Theorem 6.11. $\text{pc-NULL} \in \text{coRL}^{\text{QUMAL}_1}$.

Proof. Similar to the proof in [Theorem 1.1](#), we consider an auxiliary problem **pc-SINGULAR'**. Its input consists of a positive semi-definite matrix $H \succeq 0$ and a matrix $Q \in \mathbb{C}^{k' \times n}$, where $0 \leq k' < n$, all entries of Q are bounded by $|Q(i, j)| \leq 1$ and $\|H\|_{\text{op}} \leq 1$. Writing $M = H + Q^*Q$, the yes-instances have $\dim \ker M \geq 1$, and the no-instances have $\lambda_{\min}(M) \geq 1/\text{poly}(n)$. This problem belongs to QUMAL_1 : given (H, Q) on a read-only input tape, the entries of

$$\widehat{M} := \frac{H + Q^*Q}{1 + nk'}$$

are computable in deterministic logspace. The matrix $\widehat{M}$ is a contraction, has the same kernel as M , and is inverse-polynomially conditioned on no-instances.

Let (A, k) be a promised input to **pc-NULL**. Put $H = AA^*$. Then $H \succeq 0$, $\|H\|_{\text{op}} \leq 1$, and $\dim \ker H = \dim \ker A$. If $k = 1$, make the single oracle query $(H, 0)$. Otherwise set $k' = k - 1$ and use [Theorem 4.7](#) to sample a robust subspace design

$$\mathcal{Q}^{(k')} = \{Q_1^{(k')}, \dots, Q_m^{(k')}\} \subseteq \mathbb{C}^{k' \times n},$$

of polynomial size and robustness $\eta = \Omega(1/\sqrt{nk'})$ with probability at least $2/3$. Query **pc-SINGULAR'** on $(H, Q_i^{(k')})$ for each i , except replace entries of modulus larger than 1 by 0. Accept if all query answers are yes. Since entries of H are computable in deterministic logspace from A , and the $Q_i^{(k')}$ s are streamed to the oracle tape, the total space usage of the procedure is $O(\log n)$.

If $\dim \ker A \geq k$, every query is singular, since

$$\text{rank}(H + (Q_i^{(k')})^* Q_i^{(k')}) \leq \text{rank } H + k' \leq n - k + (k - 1) < n.$$

This also covers $k = 1$ with $Q = 0$. Hence the algorithm accepts yes-instances with certainty, regardless of the sampled subspace design.

If (A, k) is a no-instance, its promise gives $\lambda_k(H) = \sigma_k(A)^2 \geq 1/\text{poly}(n)$. For $k = 1$, $(H, 0)$ is a promised no-query. For $k \geq 2$, let H_1 contain the first k' terms in a spectral decomposition of H , and put $H_2 = H - H_1$. Then $H_1 \succeq 0$, $\dim \ker H_2 = k'$, and $\lambda_{\min}^+(H_2) = \lambda_k(H)$. On the successful design event, the Hermitian case of [Lemma 3.6](#) and [Remark 3.7](#) yields an index i for which

$$\lambda_{\min}(H_2 + (Q_i^{(k')})^* Q_i^{(k')}) \geq \frac{\eta^2 \lambda_k(H)}{3} \geq \frac{1}{\text{poly}(n)}.$$

Adding $H_1 \succeq 0$ preserves this lower bound. Thus one query is a promised no-instance, and the algorithm rejects. Answers to other queries outside the promise cannot change this outcome. The rejection probability is at least $2/3$. $\square$

We next transfer this result to well-conditioned linear-system feasibility. The following rank-parity construction adapts the classical reduction of [[ABO99](#), Corollary 2.15]. We add the analysis of the spectral gap for our setting.

Lemma 6.12. *Let $\kappa \geq 1$ and $\varepsilon \in (0, 1]$. Given an instance (A, b) of (κ, ε) -conditioned-FSLE, one can construct in deterministic logspace a Hermitian contraction $H \in \mathbb{C}^{2n \times 2n}$ such that*

$$Ax = b \text{ is feasible} \iff \text{rank } H \text{ is even}, \quad \lambda_{\min}^+(H) \geq \frac{\varepsilon^2}{18\kappa^2}.$$

Proof. Put $T = [A \ b]$ and

$$H := \frac{1}{2} \begin{pmatrix} AA^* & 0 \\ 0 & TT^* \end{pmatrix} = \frac{1}{2} \begin{pmatrix} AA^* & 0 \\ 0 & AA^* + bb^* \end{pmatrix}.$$

Its entries are computable exactly in deterministic logspace, and $\|H\|_{\text{op}} \leq 1$ because $\|A\|_{\text{op}}, \|b\| \leq 1$. If $\text{rank } A = r$, then $\text{rank } T = r$ when $Ax = b$ is feasible and $\text{rank } T = r + 1$ otherwise. Hence $\text{rank } H = r + \text{rank } T$ has the claimed parity.

It remains to bound the smallest non-zero eigenvalue. Since $\sigma_{\min}^+(A) \geq 1/\kappa$, it suffices to prove

$$\sigma_{\min}^+(T) \geq \frac{\varepsilon}{3\kappa},$$

since $\lambda_{\min}^+(H) = \frac{1}{2} \min\{(\sigma_{\min}^+(A))^2, (\sigma_{\min}^+(T))^2\}$.

If $Ax = b$ is feasible, then $b \in \text{im } A$. The matrices $TT^* = AA^* + bb^*$ and AA^* have the same kernel, and $TT^* \succeq AA^*$. Therefore $\sigma_{\min}^+(T) \geq \sigma_{\min}^+(A) \geq 1/\kappa$.

Otherwise, if $Ax = b$ is infeasible, decompose $b = b_{\parallel} + b_{\perp}$ with $b_{\parallel} \in \text{im } A$ and $b_{\perp} \perp \text{im } A$. The input promise gives $\|b_{\perp}\| \geq \varepsilon$. For $y \in \text{im } T = \text{im } A \oplus \text{span}\{b_{\perp}\}$, set

$$t := \frac{\langle b_{\perp}, y \rangle}{\|b_{\perp}\|^2}, \quad x := A^+(y - tb),$$

where A^+ is the Moore-Penrose pseudoinverse, which maps every vector $z \in \text{im}(A)$ to its minimum norm preimage x such that $Ax = z$. Then $T(x, t) = Ax + tb = y$. Moreover,

$$|t| \leq \frac{\|y\|}{\varepsilon}, \quad \|x\| \leq \kappa \|y - tb\| \leq \kappa \left(1 + \frac{1}{\varepsilon}\right) \|y\|,$$

where we used that $\|A^+\|_{\text{op}} \leq \kappa$. Thus $\|(x, t)\| \leq 3\kappa \|y\|/\varepsilon$. The minimum-norm preimage of y has no larger norm, so $\|T^+\|_{\text{op}} \leq 3\kappa/\varepsilon$, proving the claimed bound on $\sigma_{\min}^+(T)$ and hence on H . $\square$

For every threshold $k \in \{1, \dots, 2n\}$, the query “ $\dim \ker H \geq k$ ” is a promised instance of $\tilde{\kappa}$ -conditioned-NULL $_k$ with $\tilde{\kappa} = 18\kappa^2/\varepsilon^2$: if the answer is no, its k th smallest singular value is at least $\lambda_{\min}^+(H)$. These queries determine $\text{rank } H$ and therefore the feasibility of $Ax = b$.

Theorem 6.13. $\text{pc-FSLE} \in \text{BPL}^{\text{QUMAL}_1}$.

Proof. Amplify the algorithm of Theorem 6.11 to error smaller than $1/(10n)$ on each threshold query. A union bound then lets us recover $\text{rank } H$ with bounded error. $\square$

7 Acknowledgements

The authors gratefully acknowledge stimulating discussions with Quinten Tupker on the topic of logspace complexity, as well as with Isabella Ziccardi about Gaussian sampling. This work has received support under the program “Investissement d’Avenir” launched by the French Government and implemented by ANR, with the reference “ANR-22-CMAS-0001, QuantEdu-France”. SA was supported in part by the European QuantERA project QOPT (ERA-NET Cofund 2022-25), the French PEPR integrated projects EPiQ (ANR-22-PETQ0007) and HQI (ANR-22-PNCQ-0002), and the French ANR project QUOPS (ANR-22-CE47-0003-01).

AI disclosure: The main direction and conceptual contributions (notion of robust subspace designs, kernel peeling and the complexity-theoretic applications) were conceived by the authors. GPT 5.6 Sol came up with the Grassmannian net argument for the randomized construction and with an improved version of our explicit robust construction (improving the robustness parameter from inverse exponential to inverse polynomial in [Theorem 1.4](#)). GPT 6.0 Astra came up with the idea for improving the size from $\binom{n}{k}$ to $\text{poly}(n)k^{O(k)}$ through sketching in [Theorem 5.2](#). GPT 5.6 Sol, 6.0 Astra and 6.1 Sol were also used to improve the presentation of parts of the manuscript and to help identify errors. The authors take full responsibility for the content and correctness of the manuscript.

References

- [ABI86] Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of Algorithms*, 7(4):567–583, 1986.
- [ABO99] Eric Allender, Robert Beals, and Mitsunori Ogihara. The complexity of matrix rank and feasible systems of linear equations. *Computational Complexity*, 8(2):99–126, 1999.
- [ABOBS22] Dorit Aharonov, Michael Ben-Or, Fernando G.S.L. Brandão, and Or Sattath. The pursuit of uniqueness: Extending Valiant-Vazirani theorem to the probabilistic and quantum settings. *Quantum*, 6:668, March 2022.
- [AGZ10] Greg W Anderson, Alice Guionnet, and Ofer Zeitouni. *An introduction to random matrices*, volume 118 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 2010.
- [AHHN26] Anurag Anshu, Jonas Haferkamp, Yeongwoo Hwang, and Quynh T. Nguyen. On the Complexity of Unique Quantum Witnesses and Quantum Approximate Counting. In Shubhangi Saraf, editor, *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*, volume 362 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:20, Dagstuhl, Germany, 2026. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [All86] Eric W. Allender. The complexity of sparse sets in P. In Alan L. Selman, editor, *Structure in Complexity Theory*, pages 1–11, Berlin, Heidelberg, 1986. Springer Berlin Heidelberg.
- [AO94] Eric Allender and Mitsunori Ogihara. Relationships among PL, #L, and the determinant. In *Proceedings of IEEE 9th Annual Conference on Structure in Complexity Theory*, pages 267–278, 1994.
- [App26] Benny Applebaum. From nondeterministic to modular branching programs, uniformly and in logspace. Technical Report TR26–131, Electronic Colloquium on Computational Complexity, 2026.
- [AS17] Guillaume Aubrun and Stanisław J Szarek. *Alice and Bob meet Banach*, volume 223. American Mathematical Society, 2017.
- [BASTS08] Avraham Ben-Aroya, Oded Schwartz, and Amnon Ta-Shma. Quantum expanders: Motivation and constructions. In *2008 23rd Annual IEEE Conference on Computational Complexity*, pages 292–303. IEEE, 2008.

- [BCDZ26] Joshua Brakensiek, Yeyuan Chen, Manik Dhar, and Zihan Zhang. From random to explicit via subspace designs with applications to local properties and matroids. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, STOC 2026. Association for Computing Machinery, 2026.
- [Ber84] Stuart J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Information Processing Letters*, 18(3):147–150, 1984.
- [Bha97] Rajendra Bhatia. *Matrix Analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer, 1997.
- [BJLR91] Gerhard Buntrock, Birgit Jenner, Klaus-Jörn Lange, and Peter Rossmanith. Unambiguity and fewness for logarithmic space. In L. Budach, editor, *Fundamentals of Computation Theory*, pages 168–179, Berlin, Heidelberg, 1991. Springer Berlin Heidelberg.
- [CDL01] Andrew Chiu, George Davida, and Bruce Litow. Division in logspace-uniform NC¹. *RAIRO. Theoretical Informatics and Applications*, 35(3):259–275, 2001.
- [CGG⁺26] Abhranil Chatterjee, Sumanta Ghosh, Rohit Gurjar, Roshan Raj, and Thomas Thierauf. Bipartite matching is in NC. *Electronic Colloquium on Computational Complexity (ECCC)*, TR26-100, 2026.
- [Coo85] Stephen A. Cook. A taxonomy of problems with fast parallel algorithms. *Information and Control*, 64(1–3):2–22, January 1985.
- [CW09] Kenneth L. Clarkson and David P. Woodruff. Numerical linear algebra in the streaming model. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 205–214. ACM, 2009.
- [DF06] Patrick Desrosiers and Peter J. Forrester. Asymptotic correlations for Gaussian and Wishart matrices with external source. *International Mathematics Research Notices*, 2006:27395, 2006.
- [DPT24] Dean Doron, Edward Pyne, and Roei Tell. Opening up the distinguisher: A hardness to randomness approach for $BPL = L$ that uses properties of BPL. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*, pages 2039–2049. ACM, 2024. See Claim 3.17 in the ECCC version, Report TR23-208.
- [FG15] Michael A. Forbes and Venkatesan Guruswami. Dimension expanders via rank condensers. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2015)*, pages 800–814. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2015.
- [FKL⁺16] Bill Fefferman, Hirotada Kobayashi, Cedric Yen-Yu Lin, Tomoyuki Morimae, and Harumichi Nishimura. Space-Efficient Error Reduction for Unitary Quantum Computations. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, volume 55 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 14:1–14:14, Dagstuhl, Germany, 2016. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

- [FL18] Bill Fefferman and Cedric Yen-Yu Lin. A complete characterization of unitary quantum space. In Anna R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*, volume 94 of *LIPICs*, pages 4:1–4:21, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [For10] Peter J. Forrester. *Log-Gases and Random Matrices*, volume 34 of *London Mathematical Society Monographs*. Princeton University Press, Princeton, NJ, 2010.
- [FR21] Bill Fefferman and Zachary Remscrim. Eliminating intermediate measurements in space-bounded quantum computation. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, pages 1343–1356. Association for Computing Machinery, 2021.
- [GK16] Venkatesan Guruswami and Swastik Kopparty. Explicit subspace designs. *Combinatorica*, 36(2):161–185, 2016.
- [GOS⁺10] Sergei A Goreinov, Ivan V Oseledets, Dmitry V Savostyanov, Eugene E Tyrtysnikov, and Nikolay L Zamarashkin. How to find a good submatrix. In *Matrix Methods: Theory, Algorithms And Applications: Dedicated to the Memory of Gene Golub*, pages 247–256. World Scientific, 2010.
- [GX13] Venkatesan Guruswami and Chaoping Xing. List decoding Reed–Solomon, algebraic-geometric, and Gabidulin subcodes up to the singleton bound. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC)*, pages 843–852. ACM, 2013.
- [HKPV06] J. Ben Hough, Manjunath Krishnapur, Yuval Peres, and Bálint Virág. Determinantal processes and independence. *Probability Surveys*, 3:206–229, 2006.
- [JKK⁺11] Rahul Jain, Iordanis Kerenidis, Greg Kuperberg, Miklos Santha, Or Sattath, and Shengyu Zhang. On the power of a unique quantum witness, 2011.
- [JKK⁺12] Rahul Jain, Iordanis Kerenidis, Greg Kuperberg, Miklos Santha, Or Sattath, and Shengyu Zhang. On the power of a unique quantum witness. *Theory of Computing*, 8(17):375–400, 2012.
- [KSV02] Alexei Yu. Kitaev, Alexander Shen, and Mikhail N. Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate Studies in Mathematics*. American Mathematical Society, 2002.
- [LQW⁺25] Yinan Li, Youming Qiao, Avi Wigderson, Yuval Wigderson, and Chuanqi Zhang. On linear-algebraic notions of expansion. *Theory of Computing*, 21(1):1–28, 2025.
- [LZ08] Alexander Lubotzky and Efim Zelmanov. Dimension expanders. *Journal of Algebra*, 319(2):730–738, 2008.
- [MV97] Meena Mahajan and V. Vinay. Determinant: Combinatorics, algorithms, and complexity. *Chicago Journal of Theoretical Computer Science*, 1997(5), 1997.
- [MWK⁺22] Oren Mangoubi, Yikai Wu, Satyen Kale, Abhradeep Thakurta, and Nisheeth K Vishnoi. Private matrix approximation and geometry of unitary orbits. In *Conference on Learning Theory*, pages 3547–3588. PMLR, 2022.

- [MZ12] David M. Mason and Harrison H. Zhou. Quantile coupling inequalities and their applications. *Probability Surveys*, 9:439–479, 2012.
- [NIS25] NIST Digital Library of Mathematical Functions. Chapter 18: Orthogonal polynomials, section 18.14: Inequalities. <https://dlmf.nist.gov/18.14>, 2025. Eq. (18.14.8), Release 1.2.4 of 2025-03-15.
- [Sar06] Tamás Sarlós. Improved approximation algorithms for large matrices via random projections. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 143–152. IEEE Computer Society, 2006.
- [Sza98] Stanisław J. Szarek. Metric entropy of homogeneous spaces. *Banach Center Publications*, 43:395–410, 1998.
- [TS13] Amnon Ta-Shma. Inverting well conditioned matrices in quantum logspace. In *Proceedings of the Annual ACM Symposium on Theory of Computing*, pages 881–890, 06 2013.
- [Val76] Leslie G. Valiant. Relative complexity of checking and evaluating. *Information Processing Letters*, 5(1):20–23, 1976.
- [Ver18] Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, 2018.
- [VV86] Leslie G. Valiant and Vijay V. Vazirani. NP is as easy as detecting unique solutions. *Theor. Comput. Sci.*, 47(3):85–93, 1986.

A Explicit construction based on Wronskian evaluation maps

Lemma A.1 (Wronskian degree bound). *Let $N \in \mathbb{C}^{n \times k}$ have rank k . Then*

$$W_N(z) := \det(Q_{k,z}N)$$

is a nonzero polynomial of degree at most $k(n - k)$.

Proof. Let $f_1, \dots, f_k \in \mathbb{C}[z]_{<n}$ be the polynomials corresponding to the columns of N . Since N has rank k , the polynomials $f_1, \dots, f_k$ are linearly independent. By an invertible change of basis, we may assume that

$$\deg(f_1) > \deg(f_2) > \dots > \deg(f_k).$$

Indeed, replacing N by NR for some invertible matrix $R \in \mathbb{C}^{k \times k}$ changes W_N only by the nonzero scalar $\det(R)$. Writing $d_j := \deg(f_j)$, we thus have

$$d_j \leq n - j.$$

Now

$$Q_{k,z}N = \begin{pmatrix} f_1(z) & \dots & f_k(z) \\ f_1'(z) & \dots & f_k'(z) \\ \vdots & & \vdots \\ f_1^{(k-1)}(z) & \dots & f_k^{(k-1)}(z) \end{pmatrix},$$

so W_N is the Wronskian of $f_1, \dots, f_k$. Over a field of characteristic zero, the Wronskian of linearly independent polynomials is nonzero. Finally, as $\deg(f_j^{(\ell)}) \leq d_j - \ell$, we find that

$$\deg W_N \leq \sum_{j=1}^k d_j - \frac{k(k-1)}{2} \leq \sum_{j=1}^k (n-j) - \frac{k(k-1)}{2} = k(n-k). \quad \square$$

With the above lemma we can now prove [Theorem 3.8](#).

Theorem 3.8 (Wronskian subspace design). *For every $1 \leq k < n$, and every set $T_k \subseteq \mathbb{R}$ of $m > k(n-k)$ points, the family*

$$\mathcal{H}_{n,k} := \{H_a : a \in T_k\}, \quad \text{where} \quad H_a := \ker(Q_{k,a}),$$

is a $(k, k(n-k))$ -strong subspace design. Moreover, if $T_k = \{0, 1, \dots, m-1\}$, then the entries of all $Q_{k,a}$ are of polynomial bit complexity and computable in space $O(\log nm)$.

Proof. Fix $1 \leq k < n$, a set $T_k \subseteq \mathbb{R}$ of m distinct points and the family $\mathcal{H}_{n,k}$ as in the statement. Let $W \subseteq \mathbb{C}^n$ be a k -dimensional subspace, and let $N \in \mathbb{C}^{n \times k}$ have columns forming a basis of W . By [Lemma A.1](#), $W_N(z) = \det(Q_{k,z}N)$ is a nonzero polynomial of degree

$$\deg W_N \leq k(n-k).$$

For $a \in T_k$, let

$$k_a := \dim(H_a \cap W) = \dim \ker(Q_{k,a}N).$$

Since we have $\text{rank}(Q_{k,a}N) = k - k_a$, there exist invertible matrices $P, Q \in \mathbb{C}^{k \times k}$ such that

$$P(Q_{k,a}N)Q = \begin{pmatrix} I_{k-k_a} & 0 \\ 0 & 0 \end{pmatrix},$$

i.e. every entry in the last k_a rows of $P(Q_{k,z}N)Q$ vanishes at $z = a$, and is therefore divisible by $(z - a)$. Factoring one copy of $(z - a)$ from each of these rows yields

$$(z - a)^{k_a} \mid \det P(Q_{k,z}N)Q.$$

On the other hand,

$$\det P(Q_{k,z}N)Q = \det(P)W_N(z)\det(Q),$$

where $\det(P), \det(Q) \neq 0$. Consequently,

$$(z - a)^{k_a} \mid W_N(z).$$

Putting all together, we find

$$\sum_{a \in T_k} \dim(H_a \cap W) = \sum_{a \in T_k} k_a \leq \sum_{a \in T_k} \text{ord}_a(W_N(z)) \leq \deg W_N(z) \leq k(n-k),$$

where $\text{ord}_a(W_N(z))$ denotes the multiplicity of the root a in $W_N(z)$, i.e. the maximum number of times $k \geq 0$ such that $(z - a)^k \mid W_N(z)$. Thus, $\mathcal{H}_{n,k}$ is a $(k, k(n-k))$ -strong subspace design.

The computability in space $O(\log nm)$ of the entries of the $Q_{k,a}$ given $T_k = \{0, \dots, m-1\}$ follows directly since integer arithmetic and matrix multiplication over polynomial-bit numbers can be performed in space logarithmic in the input size (see e.g. [\[CDL01\]](#)). $\square$

B Proofs of Gaussian matrix concentration

Let $G \in \mathbb{C}^{k \times k}$ be a square, standard complex Gaussian matrix; that is, its entries are independent complex Gaussian random variables with density $\frac{1}{\pi} e^{-|z|^2}$. We devote this appendix to characterizing the smaller singular values of G . Our discussion will largely follow [For10, HKPV06].

B.1 Background: Laguerre unitary ensemble and determinantal point processes

If G is a square, standard complex Gaussian matrix, then the random matrix GG^* is drawn from the *complex Wishart ensemble*. Let $x_1, \dots, x_k$ be the (unordered) eigenvalues of the positive semidefinite matrix GG^* . Equivalently, $x_1, \dots, x_k$ are the (unordered) squared singular values of G . By [For10, Proposition 3.2.2], the random variable $\mathcal{X} = (x_1, \dots, x_k)$ follows the square Laguerre unitary ensemble (LUE) with joint density $p(x_1, \dots, x_k)$ proportional to

$$p(x_1, \dots, x_k) \propto \prod_{1 \leq i < j \leq k} (x_i - x_j)^2 \prod_{i=1}^k e^{-x_i}, \quad x_1, \dots, x_k \geq 0.$$

The LUE $\mathcal{X} = (x_1, \dots, x_k)$ is a point process on the real half line $[0, \infty)$. For any subset $D \subset [0, \infty)$, we denote by $\mathcal{X}(D)$ the number of points in $\mathcal{X}$ that fall in D , i.e.,

$$\mathcal{X}(D) = |\{j \in [k] : x_j \in D\}|.$$

Following [HKPV06], for any $1 \leq \ell \leq k$ we define the *joint intensities* (or correlation functions) of $\mathcal{X}$ as functions $\rho_\ell : [0, \infty)^\ell \rightarrow [0, \infty)$ such that for any family of mutually disjoint subsets $D_1, \dots, D_\ell \subset [0, \infty)$ we have

$$\mathbb{E} \left[\prod_{j=1}^\ell \mathcal{X}(D_j) \right] = \int_{\prod_{j=1}^\ell D_j} \rho_\ell(y_1, \dots, y_\ell) dy_1 \dots dy_\ell.$$

In particular, when $\ell = k$ then $\rho_\ell = k!p$ while for $\ell = 1$ the intensity $\rho_1(x)$ simply corresponds to the eigenvalue density.

The key fact is that the LUE is a *determinantal point process* (DPP). For general background, see [For10]. This means that there exists a kernel $K : [0, \infty)^2 \rightarrow \mathbb{R}$ such that

$$\rho_\ell(x_1, \dots, x_\ell) = \det(K(x_i, x_j)_{1 \leq i, j \leq \ell}).$$

For the LUE, the kernel is

$$K(x, y) = e^{-(x+y)/2} \sum_{j=0}^{k-1} L_j(x) L_j(y),$$

with $L_j(x)$ the j -th Laguerre polynomial, normalized so that

$$\int_0^\infty e^{-x} L_j(x) L_\ell(x) dx = \delta_{j\ell}.$$

To explain the above facts more specifically, we refer to Section 2 of [DF06]. The equation 2.3 is the key for us which states

$$K(x, y) = \sqrt{w(x)w(y)} \sum_{j=0}^{k-1} \frac{p_j(x)p_j(y)}{\|p_j\|^2}.$$

For us the weight functions are $w(x) = e^{-x}$ and $w(y) = e^{-y}$. Furthermore, as explained in Equation 2.7 of [DF06],

$$p_j(x) = (-1)^j j! L_j(x) \quad p_j(y) = (-1)^j j! L_j(y).$$

Putting these values, we derive the required expression for the kernel entry $K(x, y)$ ².

By the standard inequality $e^{-x/2} |L_k(x)| \leq L_k(0) \leq 1$ for $x \geq 0$ [NIS25], we get that

$$\rho_1(x) = K(x, x) \leq k.$$

It follows that for every $0 < u \leq 1$ we have

$$\mathbb{E}[\mathcal{X}([0, u^2/k])] = \int_0^{u^2/k} \rho_1(x) dx \leq \int_0^{u^2/k} k dx = u^2. \quad (4)$$

B.2 Singular value concentration

To prove a concentration bound over families of Gaussian matrices, we will use the following theorem. It is essentially [AGZ10, Corollary 4.2.24], or alternatively follows by applying [HKPV06, Theorem 7] to the restricted DPP $\mathcal{X}|_D$ with kernel $K|_D(x, y) = K(x, y)$ if $(x, y) \in D$ and $K|_D(x, y) = 0$ otherwise.

Theorem B.1. *Suppose $\mathcal{X}$ is a DPP on $[0, \infty)$ with a trace-class kernel K . Fix a measurable subset $D \subset [0, \infty)$ and consider the restricted kernel $K|_D$ with eigenvalues $\{\lambda_j(K|_D) \in [0, 1]\}$. We have that*

$$\mathcal{X}(D) \stackrel{d}{=} \sum_j \text{Bernoulli}(\lambda_j).$$

I.e., $\mathcal{X}(D)$ is distributed as a sum of independent $\text{Bernoulli}(\lambda_j)$ random variables.

We will apply this theorem to the LUE kernel. To get a bound on the eigenvalues $\lambda_j(K|_D)$, we note that $K|_D$ is positive semi-definite and that

$$\sum_{j=1}^k \lambda_j(K|_D) = \text{Tr}(K|_D) = \int_D K(x, x) dx \leq k|D|,$$

with $|D|$ the Lebesgue measure of D . This implies that $\mathcal{X}(D)$ will be distributed as a sum of k Bernoulli with sum of parameters bounded by $k|D|$. We use this to prove the following claim, which implies Lemma 4.4.

Claim B.2. *Fix $0 < u \leq 1$. If $\mathcal{X}_1, \dots, \mathcal{X}_m$ are independent copies of the LUE $\mathcal{X}$, then*

$$\Pr \left[\sum_{i=1}^m \mathcal{X}_i([0, u^2/k]) \geq 2u^2 m \right] \leq e^{-u^2 m/3}.$$

Proof. Let $S = \sum_{i=1}^m \mathcal{X}_i([0, u^2/k])$. By the preceding argument, we get that S behaves as a sum of mk independent Bernoulli random variables with sum of parameter bounded by $u^2 m$, and so mean $\mu = \mathbb{E}S \leq u^2 m$.

Let $1 + \delta = \frac{2u^2 m}{\mu}$. By the multiplicative Chernoff bound for independent Bernoulli variables, we know that for all $\delta \geq 0$ it holds that

$$\Pr[S \geq (1 + \delta)\mu] \leq e^{-\delta^2 \mu / (2 + \delta)}.$$

²For the square matrix case the term α used in [DF06] is 0.

Let $x := \frac{\mu}{u^2 m} \in (0, 1]$. Then,

$$\delta = \frac{2-x}{x} \quad \text{and} \quad x + \delta x = 2, \text{ hence } \frac{\delta^2 \mu}{2 + \delta} = u^2 m \frac{(2-x)^2}{2+x}.$$

The claim follows by noticing that, for $0 \leq x \leq 1$,

$$\frac{(2-x)^2}{2+x} \geq \frac{1}{3}.$$

□

C Proofs for complexity-theoretic applications

C.1 Constant pc-NULL reduces to pc-SINGULAR

Here we provide a logspace many-one reduction from pc-NULL_k for $k \in O(1)$ to pc-SINGULAR . The underlying idea is due to Jain, Kerenidis, Kuperberg, Santha, Sattath and Zhang [JKK⁺12]. For a background on antisymmetric subspaces, we refer to [Bha97, Section 1.5].

Theorem C.1. *Let $k = O(1)$. Then*

$$\kappa\text{-conditioned-NULL}_k \preceq_m^L (k\kappa^2)\text{-conditioned-SINGULAR}.$$

In particular,

$$\text{pc-NULL}_k \preceq_m^L \text{pc-SINGULAR}.$$

Proof. Let $A \in \mathbb{C}^{n \times n}$ be an input matrix to the problem pc-NULL_k with $\|A\|_{\text{op}} \leq 1$, and set

$$H := A^* A.$$

Then H is positive semidefinite, $\ker(H) = \ker(A)$, $\|H\|_{\text{op}} \leq 1$ and the eigenvalues of H are the squared singular values of A . Write $0 \leq \lambda_1 \leq \dots \leq \lambda_n$ for these eigenvalues, with corresponding orthonormal eigenvectors $\psi_1, \dots, \psi_n$.

We use the k -th antisymmetric tensor power $\bigwedge^k \mathbb{C}^n$. This is a $\binom{n}{k}$ -dimensional space with orthonormal basis

$$e_{i_1} \wedge \dots \wedge e_{i_k}, \quad 1 \leq i_1 < \dots < i_k \leq n.$$

If we index the standard basis by the k -subsets of $[n]$, then we can identify $\bigwedge^k \mathbb{C}^n$ with $\mathbb{C}^{\binom{n}{k}}$. Alternatively, we can identify $\bigwedge^k \mathbb{C}^n$ with the antisymmetric subspace of $(\mathbb{C}^n)^{\otimes k}$ by letting

$$v_1 \wedge \dots \wedge v_k := \frac{1}{\sqrt{k!}} \sum_{\sigma \in S_k} \text{sgn}(\sigma) v_{\sigma(1)} \otimes \dots \otimes v_{\sigma(k)}.$$

Now consider the operator on $(\mathbb{C}^n)^{\otimes k}$

$$H^{(k)} := \sum_{r=1}^k I^{\otimes(r-1)} \otimes H \otimes I^{\otimes(k-r)}.$$

Since $H^{(k)}$ commutes with every permutation of the tensor factors, it maps antisymmetric tensors to antisymmetric tensors. Hence $\bigwedge^k \mathbb{C}^n$ as the antisymmetric subspace is invariant under $H^{(k)}$.

Let $K \in \mathbb{C}^{\binom{n}{k} \times \binom{n}{k}}$ be the matrix representation of

$$\frac{1}{k} H^{(k)} \big|_{\bigwedge^k \mathbb{C}^n}$$

with respect to the basis $\{e_{i_1} \wedge \cdots \wedge e_{i_k}\}_{1 \leq i_1 < \cdots < i_k \leq n}$. Since $\|H\|_{\text{op}} \leq 1$, we have $\|H^{(k)}\|_{\text{op}} \leq k$, and therefore $\|K\|_{\text{op}} \leq 1$.

The vectors

$$\psi_{i_1} \wedge \cdots \wedge \psi_{i_k}, \quad 1 \leq i_1 < \cdots < i_k \leq n,$$

form an orthonormal basis of $\bigwedge^k \mathbb{C}^n$ and satisfy

$$H^{(k)}(\psi_{i_1} \wedge \cdots \wedge \psi_{i_k}) = (\lambda_{i_1} + \cdots + \lambda_{i_k})(\psi_{i_1} \wedge \cdots \wedge \psi_{i_k}).$$

Consequently, the spectrum of K is

$$\left\{ \frac{\lambda_{i_1} + \cdots + \lambda_{i_k}}{k} : 1 \leq i_1 < \cdots < i_k \leq n \right\}.$$

Suppose first that $\dim \ker(A) \geq k$. Then $\lambda_1 = \cdots = \lambda_k = 0$, so 0 is an eigenvalue of K and hence K is singular.

Conversely, suppose that $\dim \ker(A) < k$ and $\sigma_k(A) \geq 1/\kappa$. Since the positive eigenvalues of $H = A^*A$ are the squares of the positive singular values of A , $\lambda_k \geq 1/\kappa^2$. Therefore

$$\lambda_{\min}(K) = \frac{\lambda_1 + \cdots + \lambda_k}{k} \geq \frac{\lambda_k}{k} \geq \frac{1}{k\kappa^2}.$$

Thus K is a valid no-instance of $(k\kappa^2)$ -conditioned-SINGULAR.

It remains to verify that K can be output in logarithmic space. Since $k = O(1)$,

$$\dim K = \binom{n}{k} = \text{poly}(n).$$

Index the rows and columns of K by k -subsets

$$I = \{i_1 < \cdots < i_k\}, \quad J = \{j_1 < \cdots < j_k\}$$

of $[n]$. To determine the entries of K , note that for arbitrary $v_1, \dots, v_k \in \mathbb{C}^n$,

$$H^{(k)}(v_1 \wedge \cdots \wedge v_k) = \sum_{r=1}^k v_1 \wedge \cdots \wedge H v_r \wedge \cdots \wedge v_k.$$

Hence

$$H^{(k)}(e_{j_1} \wedge \cdots \wedge e_{j_k}) = \sum_{r=1}^k \sum_{p=1}^n H(p, j_r) e_{j_1} \wedge \cdots \wedge e_p \wedge \cdots \wedge e_{j_k}.$$

Thus an application of $H^{(k)}$ can replace at most one element of J , and therefore

$$K(I, J) = 0$$

whenever I and J differ in more than one element.

If $I = J$, only the terms with $p = i_r$ contribute, and hence

$$K(I, I) = \frac{1}{k} \sum_{r=1}^k H(i_r, i_r).$$

If I and J differ in exactly one element, say

$$I \setminus J = \{i_a\}, \quad J \setminus I = \{j_b\},$$

then the unique contributing term replaces e_{j_b} by e_{i_a} . Reordering the resulting wedge into increasing order contributes the sign $(-1)^{a+b}$, so

$$K(I, J) = \frac{(-1)^{a+b}}{k} H(i_a, j_b).$$

Finally,

$$H(i, j) = \sum_{\ell=1}^n \overline{A(\ell, i)} A(\ell, j).$$

Since $k = O(1)$, the k -subsets I, J can be stored using $O(\log n)$ space, and every entry of H , and hence every entry of K , can be computed using logarithmic workspace and standard logspace arithmetic. Therefore $A \mapsto K$ is a logspace many-one reduction. $\square$

C.2 Completeness of well-conditioned singularity

In this subsection we prove the natural well-conditioned singularity completeness results for QUMAL_1 , FewQUMAL_1 and UQUMAL_1 . Let us recall the theorem:

Lemma 6.5 (Completeness of poly-conditioned singularity). *Under deterministic logspace many-one reductions,*

- (i) pc-SINGULAR is QUMAL_1 -complete (proven in [FR21, Lemma 35]),
- (ii) pc-FewSINGULAR is FewQUMAL_1 -complete, and
- (iii) pc-UniqueSINGULAR is UQUMAL_1 -complete.

Proof. We establish hardness and containment for all three claims simultaneously.

The circuit Hamiltonian. Let x have length n , and let $Q_x = U_T \cdots U_1$ be the corresponding verifier, acting on $a + w = O(\log n)$ qubits, with $T = \text{poly}(n)$. Set $V_0 = I$, $V_t = U_t \cdots U_1$. Let $\Pi_{\text{acc}} = |1\rangle\langle 1|_1 \otimes I$ project onto the first auxiliary qubit being in state $|1\rangle$, i.e. the case where the circuit accepts, and put $\Pi_{\text{rej}} = I - \Pi_{\text{acc}} = |0\rangle\langle 0|_1 \otimes I$. The acceptance operator is

$$E_x = (\langle 0^a | \otimes I) V_T^* \Pi_{\text{acc}} V_T (|0^a\rangle \otimes I).$$

Thus $\langle \psi | E_x | \psi \rangle$ is the acceptance probability of the circuit given the witness state $|\psi\rangle$. In particular, the perfectly accepting witness subspace is

$$W_x = \ker(I - E_x) = \ker(\Pi_{\text{rej}} V_T (|0^a\rangle \otimes I)).$$

We use the clock construction underlying [FR21, Lemma 35], originating in [KSV02]. On $\mathbb{C}^{2^{a+w}} \otimes \mathbb{C}^{T+1}$, define

$$\begin{aligned} H_{\text{prop}} &= \frac{1}{2} \sum_{t=1}^T \left(I \otimes (|t\rangle\langle t| + |t-1\rangle\langle t-1|) - U_t \otimes |t\rangle\langle t-1| - U_t^* \otimes |t-1\rangle\langle t| \right), \\ H_{\text{in}} &= \left(\sum_{j=1}^a |1\rangle\langle 1|_j \otimes I_{\text{rest}} \right) \otimes |0\rangle\langle 0|, \\ H_{\text{out}} &= \Pi_{\text{rej}} \otimes |T\rangle\langle T|, \\ H_x &= H_{\text{prop}} + H_{\text{in}} + H_{\text{out}}. \end{aligned}$$

Here $|1\rangle\langle 1|_j$ acts on the j th auxiliary qubit. Each summand in H_x is positive semi-definite, and one easily finds that,

$$\|H_{\text{prop}}\|_{\text{op}} \leq 2, \quad \|H_{\text{in}}\|_{\text{op}} \leq a, \quad \|H_{\text{out}}\|_{\text{op}} \leq 1, \quad \text{and thus} \quad \|H_x\|_{\text{op}} \leq a + 3.$$

As a hard instance, we output

$$A_x = \frac{H_x}{a + 3}.$$

This is a positive semidefinite $D \times D$ matrix, with $D = 2^{a+w}(T+1)$ and operator norm $\|A_x\|_{\text{op}} \leq 1$. Its entries are computable in deterministic logspace, since they involve only clock indices, auxiliary-bit tests, and entries of a single gate U_t .

Preservation of witness dimension: Define the history-state isometry

$$F_x |\psi\rangle = \frac{1}{\sqrt{T+1}} \sum_{t=0}^T V_t |0^a\rangle |\psi\rangle \otimes |t\rangle.$$

For a vector $\sum_{t=0}^T |y_t\rangle \otimes |t\rangle$, zero propagation energy is equivalent to $|y_t\rangle = U_t |y_{t-1}\rangle$ for every t . Zero initialization energy requires $|y_0\rangle = |0^a\rangle |\psi\rangle$ for some $|\psi\rangle \in \mathbb{C}^{2^w}$, and zero output energy requires $\Pi_{\text{rej}} |y_T\rangle = 0$. Since all three terms are positive semi-definite, this proves

$$\ker(A_x) = \ker(H_x) = F_x W_x, \quad \dim \ker(A_x) = \dim W_x.$$

Thus the construction preserves the dimension of the perfectly accepting witness subspace exactly.

The spectral gap analysis: Assume that, for some $0 < \gamma \leq 1$,

$$\Pr[Q_x \text{ accepts } |\psi\rangle] \leq 1 - \gamma \quad \text{for every normalized } |\psi\rangle \in W_x^\perp.$$

We claim that

$$\lambda_{\min}^+(A_x) \geq \frac{\gamma}{24(a+3)(T+1)^2}. \quad (5)$$

Since $A_x = H_x/(a+3)$, we have $\lambda_{\min}^+(A_x) = \lambda_{\min}^+(H_x)/(a+3)$. It therefore suffices to prove the corresponding lower bound for H_x , that is $\lambda_{\min}^+(H_x) \geq \frac{\gamma}{24(T+1)^2}$.

To this end, fix an arbitrary vector

$$|y\rangle = \sum_{t=0}^T |y_t\rangle \otimes |t\rangle \perp \ker(H_x), \quad |\psi_0\rangle := (\langle 0^a| \otimes I) |y_0\rangle.$$

The vectors $|y_t\rangle$ and $|\psi_0\rangle$ need not be normalized. Writing Π_j for the projector onto auxiliary qubit j being 1, the energy is

$$\eta := \langle y | H_x | y \rangle = \frac{1}{2} \sum_{t=1}^T \| |y_t\rangle - U_t |y_{t-1}\rangle \|^2 + \sum_{j=1}^a \|\Pi_j |y_0\rangle\|^2 + \|\Pi_{\text{rej}} |y_T\rangle\|^2. \quad (6)$$

Claim C.2. *We have*

- (i) $\|V_t^* |y_t\rangle - |y_0\rangle\|^2 \leq 2t\eta$ for every $0 \leq t \leq T$,
- (ii) $\| |y_0\rangle - |0^a\rangle |\psi_0\rangle \|^2 \leq \eta$, and
- (iii) $\|\Pi_{\text{rej}} |y_T\rangle\|^2 \leq \eta$.

Proof. Parts (ii) and (iii) follow from Equation (6), using $\sum_j \Pi_j \succeq I - |0^a\rangle\langle 0^a| \otimes I$ for part (ii). For part (i), the case $t = 0$ is immediate. For $t \geq 1$, telescoping and Cauchy-Schwarz give

$$\|V_t^* |y_t\rangle - |y_0\rangle\|^2 = \left\| \sum_{r=1}^t V_r^* (|y_r\rangle - U_r |y_{r-1}\rangle) \right\|^2 \leq t \sum_{r=1}^t \| |y_r\rangle - U_r |y_{r-1}\rangle \|^2 \leq 2t\eta. \quad \square$$

We next bound the rejection amplitude of $|\psi_0\rangle$. Decompose it as

$$\Pi_{\text{rej}} V_T |0^a\rangle |\psi_0\rangle = \Pi_{\text{rej}} V_T (|0^a\rangle |\psi_0\rangle - |y_0\rangle) + \Pi_{\text{rej}} V_T (|y_0\rangle - V_T^* |y_T\rangle) + \Pi_{\text{rej}} |y_T\rangle.$$

Since Π_{rej} is a contraction and V_T is unitary, Cauchy-Schwarz and Claim C.2 imply

$$\begin{aligned} \|\Pi_{\text{rej}} V_T |0^a\rangle |\psi_0\rangle\|^2 &\leq 3 \left(\| |0^a\rangle |\psi_0\rangle - |y_0\rangle \|^2 + \| |y_0\rangle - V_T^* |y_T\rangle \|^2 + \|\Pi_{\text{rej}} |y_T\rangle\|^2 \right) \\ &\leq 3(\eta + 2T\eta + \eta) = 6(T+1)\eta. \end{aligned}$$

The rejection map $\Pi_{\text{rej}} V_T (|0^a\rangle \otimes I)$ vanishes on W_x . Applying the assumed rejection gap to $\Pi_{W_x^\perp} |\psi_0\rangle$, therefore gives

$$\gamma \left\| \Pi_{W_x^\perp} |\psi_0\rangle \right\|^2 \leq \|\Pi_{\text{rej}} V_T |0^a\rangle |\psi_0\rangle\|^2 \leq 6(T+1)\eta. \quad (7)$$

Now define the corresponding zero-energy history

$$|h\rangle := \sum_{t=0}^T V_t |0^a\rangle \Pi_{W_x} |\psi_0\rangle \otimes |t\rangle = \sqrt{T+1} F_x \Pi_{W_x} |\psi_0\rangle.$$

Since $\ker(H_x) = F_x W_x$, $|h\rangle \in \ker(H_x)$. Using clock orthogonality, unitarity, and Cauchy-Schwarz, we obtain

$$\begin{aligned} \| |y\rangle - |h\rangle \|^2 &= \sum_{t=0}^T \| V_t^* |y_t\rangle - |0^a\rangle \Pi_{W_x} |\psi_0\rangle \|^2 \\ &\leq 3 \sum_{t=0}^T \left(\| V_t^* |y_t\rangle - |y_0\rangle \|^2 + \| |y_0\rangle - |0^a\rangle |\psi_0\rangle \|^2 + \left\| \Pi_{W_x^\perp} |\psi_0\rangle \right\|^2 \right) \\ &\leq 3 \sum_{t=0}^T \left(2t + 1 + \frac{6(T+1)}{\gamma} \right) \eta \leq \frac{24(T+1)^2}{\gamma} \eta, \end{aligned}$$

where the last line uses Claim C.2 and Equation (7) and $\gamma \leq 1$. Finally, $|y\rangle \perp |h\rangle$, so

$$\| |y\rangle \|^2 \leq \| |y\rangle - |h\rangle \|^2 \leq \frac{24(T+1)^2}{\gamma} \langle y | H_x | y \rangle.$$

Taking the minimum Rayleigh quotient over nonzero vectors orthogonal to $\ker(H_x)$ yields

$$\lambda_{\min}^+(H_x) = \min_{0 \neq |y\rangle \perp \ker(H_x)} \frac{\langle y | H_x | y \rangle}{\| |y\rangle \|^2} \geq \frac{\gamma}{24(T+1)^2}.$$

Dividing by $a+3$ proves Equation (5).

On no-instances of all three classes, every witness is accepted with probability at most $1/3$ and $W_x = \{0\}$, so the assumption holds with $\gamma = 2/3$. The same is true on yes-instances of FewQUMAL₁ and UQUMAL₁, by their acceptance gap on $W_x^\perp$. In these cases,

$$\lambda_{\min}^+(A_x) \geq \frac{1}{36(a+3)(T+1)^2} \geq \frac{1}{\text{poly}(n)}.$$

For ordinary QUMAL_1 , the assumption need not hold on yes-instances, only the existence of a zero eigenvalue is required. This establishes the three hardness claims.

Containment: We use the one-sided kernel-verification procedure invoked in [FR21, Lemma 35] (see also [FL18, Lemma 19 and Proposition 32]). For a Hermitian contraction H and $\delta \geq 1/\text{poly}(n)$, there is a deterministic logspace-uniform unitary verifier, using $O(\log n)$ qubits and $T = \text{poly}(n)$ gates, with the following properties:

- (i) every normalized vector in $\ker(H)$ is accepted with probability 1, and
- (ii) every normalized vector in the span of the eigenvectors of H with eigenvalues satisfying $|\lambda| \geq \delta$ is accepted with probability at most $1/3$.

When A is the input matrix, apply this procedure to $H = A^*A$ and $\delta = \kappa(n)^{-2}$. Since $\ker(H) = \ker(A)$ and the eigenvalues of H are the squared singular values of A , this gives perfect completeness and soundness for standard pc-SINGULAR (as is already described in [FR21, Lemma 35]). For pc-FewSINGULAR and pc-UniqueSINGULAR , the positive singular value gap additionally ensures acceptance probability at most $1/3$ on $\ker(A)^\perp$ on yes-instances. Their kernel dimensions therefore give the required perfectly accepting subspaces, proving containment in FewQUMAL_1 and UQUMAL_1 , respectively. $\square$