Strictly Scalable Exterior Decoders for Quantum Lists: Exact Full-Spark Widths and Fixed-Probe Weyl Bayes Curves

Lluis Eriksson

August 24, 2026

Abstract

A quantum list measurement succeeds when its classical output contains the prepared label. General zero-error feasibility is already characterized by the learning width, equivalently the factor width of the Gram matrix. We use a physical-space form of that known criterion to close an exact realization-sensitive branch and then go beyond feasibility. For every full-spark ensemble of N pure-state rays spanning $\mathbb{C}^r$ that is *strictly scalable*—its rays admit nonzero tight representatives—the minimum list size for zero error is

$$\ell_{\min} = N - r + 1.$$

The converse is forced by full spark: a smaller-list effect must annihilate r spanning states and hence vanish. Attainment is constructive. Hodge duals of all $(r-1)$-fold frame wedges form a tight frame, and their rank-one projectors give a POVM whose output is the complementary list. The factor-width criterion becomes an annihilator cone in the physical Hilbert space; standard conic Carathéodory then gives a realization with at most r^2 outcomes. Below threshold, a smallest-eigenvalue functional gives a strict Bayes-error floor for every full-spark ensemble and remains positive under explicit operator-norm perturbations. Applied to a flat Schmidt-rank-r probe of the complete d-dimensional Weyl-channel ensemble, the theorem yields the exact fixed-probe frontier $\ell_{\min} = d - r + 1$. This strictly exceeds both the support-matroid threshold $\lceil d/r \rceil$ and the global summed-projector obstruction on infinite families. An arithmetic-support probe with the same Schmidt rank has a different exact threshold. More strongly, when $r \mid d$, a dimension converse and that arithmetic construction close the optimization over *all* pure Schmidt-rank-r probes: $\min_\psi \ell_{\min}(\psi) = d/r$. For the consecutive rank-two probe we also determine the complete one-shot Bayes curve,

$$P^\star_{\mathrm{succ}}(d, \ell) = \frac{\ell + \sin(\pi\ell/d)/\sin(\pi/d)}{d}, \qquad 1 \le \ell \le d.$$

The nondivisible probe optimum, adaptive channel testers, and asymptotic capacity remain outside scope.

1 Introduction

Quantum state exclusion asks a measurement to certify hypotheses that did not produce the observed system. Its list formulation returns a set of candidate labels and succeeds when the true label lies in that set. General semidefinite formulations and optimality conditions are known [1]; antidistinguishability has exact algebraic and Gram-matrix criteria in several settings [2, 3, 4], and finite group orbits admit exact one-state-exclusion solutions [5]. Recent work also studies asymptotic zero-error list capacity for pure-state classical–quantum channels [6]. Two-out-of-four state elimination has both an exact finite analysis and an optical realization [7, 8]; the multiple-exclusion task itself is therefore not new here.

Most directly, Johnston, Lovitz, Russo, and Sikora define the same finite candidate-list task as k-learnability and prove that its minimum perfect list size, the *learning width*, equals the factor width of the state Gram matrix [9]. Their result subsumes the general feasibility question. We do not claim a new criterion for arbitrary ensembles. Our contribution is an exact evaluation of that width on the strictly scalable full-spark branch, a weighted exterior-power factorization that gives the POVM explicitly, and exact fixed- and optimized-probe Weyl consequences including a complete rank-two Bayes curve. We also record a simple perturbatively stable spectral subthreshold bound.

The present question is deliberately different and finite: given one copy of one of N prescribed pure states, what is the smallest output list that can contain the label with certainty? This question separates two layers. Linear dependence controls how many states a nonzero effect can annihilate, while positivity decides whether sufficiently many such effects can resolve the identity. A support matroid captures the first layer but not the second [10]. We identify a broad branch where both layers close exactly.

Let $g_1, \dots, g_N \in \mathbb{C}^r$ be unit state vectors. We call their ray ensemble strictly scalable when positive numbers s_i exist such that $f_i = s_i g_i$ form a tight frame [11]. This is broader than requiring the physical state vectors themselves to form a unit-norm tight frame, while remaining invariant under the choice of nonzero ray representatives. Tightness of the scaled representatives supplies a positive identity resolution. Full spark—every r rays are independent—maximizes resistance to erasure and is a standard frame-theory notion [12, 13]. Neither property alone determines our threshold. Together they do:

$$\boxed{\ell_{\min} = N - r + 1.}$$

The construction uses exterior algebra. For every $(r-1)$-subset E of labels, the Hodge dual w_E of $\bigwedge_{i \in E} f_i$ is orthogonal exactly to those $r-1$ frame vectors. Exterior powers of the tight frame operator show that the projectors $|w_E\rangle\langle w_E|$ resolve the identity after one common normalization. The measurement reports $[N] \setminus E$.

The cross-product tightness identity itself is classical in real frame geometry [14]; full-spark harmonic frames are likewise classical [12]. We do not claim either ingredient as new. The contribution is their exact synthesis into a quantum list decoder, the matching full-spark converse, a physical-space realization of the known factor-width criterion, and the resulting Weyl frontiers.

Our main results are:

1. The known learning-width/factor-width criterion has an equivalent physical-space form: the identity lies in a cone generated by projectors onto vectors with sufficiently many frame zeros (Theorem 2.1). This form also yields the standard r^2 conic-Carathéodory realization.
2. Every strictly scalable full-spark ensemble of N pure-state rays in $\mathbb{C}^r$ has exact perfect list threshold $N - r + 1$, attained by an explicit weighted Hodge POVM (Theorem 3.2).
3. For arbitrary nonnegative priors and density operators, the simple spectral quantity γ_ℓ gives the universal error floor $P_{\mathrm{err}} \geq r\gamma_\ell$. It is strictly positive below the full-spark threshold when every prior is positive, and is stable under operator-norm perturbations (Theorem 4.1).
4. For the fixed flat consecutive-support Schmidt-rank-r probe of all d^2 Weyl channels, perfect decoding is possible exactly when $\ell \geq d - r + 1$ (Theorem 6.1).
5. For $r = 2$, every subthreshold optimum is explicit: $P^\star_{\mathrm{succ}}(d, \ell) = [\ell + \sin(\pi\ell/d)/\sin(\pi/d)]/d$ (Theorem 6.3).
6. If $r \mid d$, optimizing over every pure Schmidt-rank-r probe gives the exact global value $\min_\psi \ell_{\min}(\psi) = d/r$ (Corollary 6.5).

The hypotheses and resource claims are intentionally narrow. Full spark and strict scalability are separately necessary for the stated universal theorem, as counterexamples show. The global Weyl optimum is closed only in the divisor branch $r \mid d$; the nondivisible branch, adaptive uses, and indefinite causal order are not optimized here.

2 Learning width, factor width, and annihilator cones

Let $[N] = \{1, \dots, N\}$ and let $g_i \in \mathbb{C}^r$ be unit vectors spanning $\mathbb{C}^r$. Write $\rho_i = |g_i\rangle\langle g_i|$. A list POVM of size at most ℓ is a family

$$\{M_L \succeq 0 : L \subseteq [N],\ |L| \le \ell\}, \qquad \sum_L M_L = \mathbb{I}_r.$$

Zero effects and unused lists are allowed. This is the weak exclusion convention; a strong convention requiring every listed subset to occur with a nonzero effect is a different problem [15]. It is perfect if

$$\operatorname{Tr}(\rho_i M_L) = 0 \qquad \text{whenever } i \notin L. \tag{1}$$

For any prior with full support, this is equivalent to unit average success. Lists of smaller cardinality may be padded, so "at most ℓ" and "exactly ℓ" have the same feasibility threshold.

The *learning width* is the least such feasible ℓ. If $G = [\langle g_i, g_j\rangle]_{i,j=1}^N$ is the Gram matrix, its factor width is the least ℓ for which

$$G = \sum_a |u_a\rangle\langle u_a|, \qquad |\operatorname{supp} u_a| \le \ell. \tag{2}$$

Johnston et al. prove that learning width equals factor width [9]. The next theorem records the equivalent physical-space form needed for the exterior construction; the equivalence itself is not claimed new.

For $v \neq 0$, define its zero set with respect to the frame

$$Z(v) = \{i \in [N] : \langle g_i, v\rangle = 0\}$$

and the admissible annihilator set

$$\mathcal{V}_\ell = \{v \in \mathbb{C}^r \setminus \{0\} : |Z(v)| \ge N - \ell\}.$$

Theorem 2.1 (Factor width in physical-space form). *A perfect list decoder of size at most ℓ exists if and only if*

$$\mathbb{I}_r \in \operatorname{cone}\{|v\rangle\langle v| : v \in \mathcal{V}_\ell\}. \tag{3}$$

Equivalently, the Gram matrix has factor width at most ℓ. If it exists, one can choose a perfect decoder with at most r^2 nonzero outcomes.

Proof. Suppose first that $\mathbb{I}_r = \sum_{a=1}^m c_a |v_a\rangle\langle v_a|$ with $c_a > 0$ and $v_a \in \mathcal{V}_\ell$. Set $M_a = c_a|v_a\rangle\langle v_a|$ and report $L_a = [N] \setminus Z(v_a)$. Then $|L_a| \le \ell$, the effects sum to the identity, and $i \notin L_a$ implies $\langle g_i, v_a\rangle = 0$. Hence Eq. (1) holds.

Conversely, let $\{M_L\}$ be perfect. Positivity gives

$$\langle g_i, M_L g_i\rangle = 0 \quad \Longrightarrow \quad M_L^{1/2} g_i = 0.$$

Thus every eigenvector with positive eigenvalue in a spectral decomposition of M_L is orthogonal to all g_i with $i \notin L$, hence belongs to $\mathcal{V}_\ell$. Summing all spectral decompositions proves Eq. (3).

Finally, Hermitian $r \times r$ matrices form a real vector space of dimension r^2. Conic Caratheodory reduces any representation of $\mathbb{I}_r$ in Eq. (3) to at most r^2 generators. The first part of the proof turns those generators into a perfect decoder.

For completeness, let $V : \mathbb{C}^N \to \mathbb{C}^r$ have columns g_i, so $G = V^*V$. A cone decomposition $\mathbb{I}_r = \sum_a c_a|v_a\rangle\langle v_a|$ maps to $u_a = \sqrt{c_a}V^*v_a$, giving Eq. (2); the support of u_a is the reported list. Conversely, every factor u_a in Eq. (2) lies in $\operatorname{ran} G = \operatorname{ran} V^*$, so write $u_a = V^*v_a$. Then $V^*(\sum_a |v_a\rangle\langle v_a| - \mathbb{I}_r)V = 0$. Since V is surjective, the middle operator vanishes. This is the coordinate translation of the known factor-width criterion. $\square$

Theorem 2.1 is the realization-sensitive, physical-space form of the learning-width criterion. The zero sets are flats of the represented support matroid, but the inclusion of $\mathbb{I}_r$ in their projector cone depends on angles and positivity. For $\ell = N - 1$ it is closely related to known rank-one descriptions of antidistinguishing POVMs [2]; its role here is to connect sparse Gram factors to physical annihilators and to compress the exterior decoder below.

3 The exact scalable full-spark threshold

A nonzero frame $F = (f_i)_{i=1}^N \subset \mathbb{C}^r$ is tight if

$$\sum_{i=1}^{N} |f_i\rangle\langle f_i| = \alpha \mathbb{I}_r, \qquad \alpha = \frac{1}{r}\sum_{i=1}^{N} \|f_i\|^2 > 0. \tag{4}$$

It is full spark if every r frame vectors are linearly independent. For the associated physical states we normalize only at the end,

$$g_i = \frac{f_i}{\|f_i\|}, \qquad \rho_i = |g_i\rangle\langle g_i|. \tag{5}$$

Thus the hypothesis concerns the rays: equivalently, the unit vectors g_i form a strictly scalable full-spark ensemble, with positive scaling weights $s_i = \|f_i\|$. Individual normalization need not preserve tightness; the representatives in Eq. (4) are the data used by the exterior construction.

For $E = \{i_1, \dots, i_{r-1}\} \subset [N]$, let

$$f_E = f_{i_1} \wedge \cdots \wedge f_{i_{r-1}} \in \bigwedge^{r-1} \mathbb{C}^r.$$

Fix an antiunitary Hodge identification $\star : \bigwedge^{r-1} \mathbb{C}^r \to \mathbb{C}^r$ and put $w_E = \star f_E$. Equivalently, w_E is the cofactor vector characterized up to the fixed Hodge convention by

$$\langle f_i, w_E\rangle = 0 \quad (i \in E), \qquad \|w_E\|^2 = \det[\langle f_i, f_j\rangle]_{i,j\in E}. \tag{6}$$

Lemma 3.1 (Exterior tightness)**.** *If F obeys Eq. (4), then*

$$\sum_{E\in\binom{[N]}{r-1}} |w_E\rangle\langle w_E| = \alpha^{r-1}\mathbb{I}_r. \tag{7}$$

Proof. Let $T : \mathbb{C}^N \to \mathbb{C}^r$ be the synthesis operator $Te_i = f_i$. Then $TT^* = \alpha\mathbb{I}_r$. Apply the $(r-1)$st exterior-power functor:

$$\Big(\bigwedge^{r-1} T\Big)\Big(\bigwedge^{r-1} T\Big)^* = \bigwedge^{r-1}(TT^*) = \alpha^{r-1}\mathbb{I}_{\wedge^{r-1}\mathbb{C}^r}.$$

The columns of $\bigwedge^{r-1} T$ are the f_E. Conjugation by the antiunitary Hodge map preserves rank-one sums and sends f_E to w_E, which gives Eq. (7). $\square$

The identity is a complex version of the classical cross-product tight-frame construction [14]. We use it as a POVM normalization, not as a new frame-theory claim.

Theorem 3.2 (Strictly scalable full-spark list threshold)**.** *Let $F = (f_i)_{i=1}^N$ be a nonzero tight full-spark frame spanning $\mathbb{C}^r$, with $N > r$, and let ρ_i be the normalized ray states in Eq. (5). The minimum perfect list size is*

$$\ell_{\min} = N - r + 1. \tag{8}$$

Equivalently, the Gram matrix of the normalized rays has factor width $N - r + 1$. At the threshold, an explicit perfect POVM is

$$M_E = \alpha^{-(r-1)}|w_E\rangle\langle w_E|, \qquad L_E = [N] \setminus E, \qquad |E| = r - 1. \tag{9}$$

Moreover, a perfect threshold decoder exists using at most r^2 of these complementary lists.

Proof. For the converse, suppose $\ell \le N - r$. Every allowed list L omits at least r labels. Perfectness implies $M_L g_i = 0$, equivalently $M_L f_i = 0$, for all omitted i. Any r omitted frame vectors span $\mathbb{C}^r$ by full spark, so $M_L = 0$. This would hold for every outcome, contradicting $\sum_L M_L = \mathbb{I}_r$. Therefore $\ell \ge N - r + 1$.

For attainment, Lemma 3.1 shows that the effects in Eq. (9) sum to $\mathbb{I}_r$. Equation (6) and $g_i \parallel f_i$ show that M_E annihilates every physical state with label in E, exactly the labels omitted from the reported list. Hence the decoder is perfect with list size $N - r + 1$.

Full spark makes each $(r-1)$-fold wedge nonzero and prohibits a vector from having r frame zeros. Consequently every generator of the threshold cone in Theorem 2.1 is proportional to some w_E. Conic Caratheodory therefore selects at most r^2 of the displayed effects, with possibly different positive weights, while retaining the identity resolution. □

Proposition 3.3 (Constructive r^2 compression)**.** *The symmetric Hodge POVM in Eq. (9) can be compressed to at most r^2 nonzero outcomes by finitely many real-nullspace computations and nonnegative weight updates. The retained outcomes report the same lists and remain a perfect decoder.*

Proof. Write the current identity resolution as $\mathbb{I}_r = \sum_{j=1}^m c_j P_j$, where $c_j > 0$ and every P_j is one of the Hodge projectors. Regard Hermitian matrices as vectors in the real space $\mathrm{Herm}(r)$ of dimension r^2. If $m > r^2$, compute a nonzero real vector β satisfying $\sum_j \beta_j P_j = 0$. Its coefficients have both signs: a nonzero positive combination of nonzero positive semidefinite matrices cannot vanish. After changing the sign of β if necessary, set

$$t = \min_{\beta_j > 0} \frac{c_j}{\beta_j}, \qquad c_j' = c_j - t\beta_j.$$

Then every $c_j' \geq 0$, at least one becomes zero, and $\sum_j c_j' P_j = \mathbb{I}_r$. Delete zero coefficients and repeat. The procedure terminates with at most r^2 effects. Only weights change, so every retained effect keeps its original annihilated set and reported list. □

Remark 3.4 (What each hypothesis does)**.** *Full spark is used only in the converse and to identify the maximal zero sets. Tightness is used only to place the identity in the annihilator cone through Eq. (7). Section 7 shows that neither hypothesis can be deleted from the universal statement.*

Remark 3.5 (Weighted physical form)**.** *If unit state representatives g_i obey $\sum_i \lambda_i |g_i\rangle\langle g_i| = \mathbb{I}_r$ with every $\lambda_i > 0$, take $f_i = \sqrt{\lambda_i} g_i$ and $\alpha = 1$. Then the threshold effects are the Hodge projectors of the weighted wedges $\bigwedge_{i \in E} \sqrt{\lambda_i} g_i$. Hence the theorem is genuinely projective: unequal representative norms are not physical probabilities, but a strictly positive scalability certificate used to synthesize the POVM. The unit-norm tight theorem is the special case $\lambda_i = r/N$.*

The number $\binom{N}{r-1}$ of effects in Eq. (9) is a symmetric description, not a hardware lower bound. Proposition 3.3 turns the standard conic-Caratheodory bound into an explicit finite compression algorithm. It requires nullspaces of real matrices with r^2 rows; it is not asserted to minimize the number of outcomes or to be numerically well conditioned without certified linear algebra.

4 A spectral subthreshold certificate

The zero-error threshold is discontinuous as a yes/no statement, but its failure admits a quantitative and perturbatively stable certificate. Let $p_i \geq 0$, $\sum_i p_i = 1$, and let ρ_i be arbitrary density operators on $\mathbb{C}^r$. For every allowed list L define the omitted-prior operator

$$Q_L = \sum_{i \notin L} p_i \rho_i, \qquad \gamma_\ell = \min_{|L| \leq \ell} \lambda_{\min}(Q_L).$$

Theorem 4.1 (Spectral list-error bound)**.** *For every list POVM of size at most ℓ,*

$$P_{\mathrm{err}} \geq r\gamma_\ell. \tag{10}$$

If the states are pure, have full-support priors, and their rays are full spark, then $\gamma_\ell > 0$ whenever $\ell \leq N - r$. Moreover, if $\|\tilde{\rho}_i - \rho_i\|_{\mathrm{op}} \leq \varepsilon_i$, then

$$\widetilde{P}_{\mathrm{err}} \geq r \max\left\{0, \min_{|L| \leq \ell} \left[\lambda_{\min}(Q_L) - \sum_{i \notin L} p_i \varepsilon_i\right]\right\} \geq r \max\left\{0, \gamma_\ell - \sum_i p_i \varepsilon_i\right\}. \tag{11}$$

Table 1: Support threshold versus the exact perfect-list threshold for a representative unit-norm tight subfamily.

N	r	support/global threshold $\lceil N/r \rceil$	exact $N-r+1$
4	2	2	3
6	2	3	5
6	3	2	4
8	4	2	5
10	5	2	6

Proof. For a POVM $\{M_L\}$, the probability of omitting the true label is

$$P_{\text{err}} = \sum_L \operatorname{Tr}(M_L Q_L).$$

Since $Q_L \succeq \gamma_\ell \mathbb{I}_r$ for every allowed list,

$$P_{\text{err}} \geq \gamma_\ell \sum_L \operatorname{Tr} M_L = r\gamma_\ell.$$

If $\ell \leq N-r$, every L omits at least r labels. For a full-spark pure ensemble those omitted rays span $\mathbb{C}^r$; because their priors are positive, Q_L is positive definite. There are finitely many lists, so their smallest eigenvalues have a strictly positive minimum.

Finally, $\|\widetilde{Q}_L - Q_L\|_{\text{op}} \leq \sum_{i \notin L} p_i \varepsilon_i$. Weyl's eigenvalue perturbation inequality gives the first lower bound in Eq. (11); replacing every omitted-list sum by $\sum_i p_i \varepsilon_i$ gives the second. □

This is the scalar smallest-eigenvalue witness obtained from the standard quantum decision SDP [16, 17]; no novelty is claimed for that dual mechanism. Its usefulness here is that it depends on the actual angles and priors, not only on the support matroid. It therefore supplies quantitative information that the zero-error width discards. It is a certified lower bound, not generally the exact Bayes error; Section 6.1 gives a nontrivial family whose complete optimum can nevertheless be solved.

5 What support-only obstructions miss

For a full-spark N-frame in $\mathbb{C}^r$, the rank of a subfamily is $\min\{|A|, r\}$. The ℓ-fold support-matroid obstruction therefore vanishes once

$$\ell \geq \left\lceil \frac{N}{r} \right\rceil. \tag{12}$$

This is a necessary support threshold, not a measurement construction [10]. Theorem 3.2 gives the exact physical threshold $N-r+1$, which may be much larger.

For the unit-norm tight subfamily the common global projector test also becomes

$$\sum_i \rho_i = \frac{N}{r}\mathbb{I}_r \preceq \ell \mathbb{I}_r$$

at Eq. (12). Thus both coarse tests can be silent while perfect decoding remains impossible. We stress "global": stronger subset-wise projector criteria exist in exclusion theory [15], and no equivalence with them is claimed.

For example, keeping $r=2$ and increasing N produces an infinite strict family: the support obstruction disappears at $\lceil N/2 \rceil$, while the physical threshold is $N-1$.

6 A fixed-probe Weyl-channel frontier

We now turn the frame theorem into an exact communication statement. Let $X|x\rangle = |x+1 \bmod d\rangle$ and $Z|x\rangle = \omega^x|x\rangle$, where $\omega = e^{2\pi i/d}$. One of the d^2 equiprobable channels

$$\mathcal{U}_{ab}(\cdot) = X^a Z^b(\cdot)Z^{-b}X^{-a}, \qquad (a,b) \in \mathbb{Z}_d^2,$$

is used once. Dense coding is the rank-d endpoint [18]. Partially entangled deterministic dense coding and its probe dependence are classical [19, 20]. We fix the specific flat consecutive-support Schmidt-rank-r probe

$$|\Phi_r\rangle = \frac{1}{\sqrt{r}}\sum_{x=0}^{r-1}|x\rangle_A|x\rangle_R, \qquad 1 \le r \le d. \tag{13}$$

No optimization over probes is implicit. The companion record [10] gives a general unitary-error-basis list cap and reports the point $(d,r,\ell) = (4,2,2)$ for a consecutive-support probe. The theorem below strictly generalizes that fixed-probe comparison across all d and r; it does not re-claim the companion fixture.

Theorem 6.1 (Fixed-probe Weyl frontier)**.** *For the complete d-dimensional Weyl ensemble and the probe Eq. (13), perfect one-use list decoding is possible if and only if*

$$\ell \ge d - r + 1. \tag{14}$$

At equality, a decoder first reads the shift sector and then applies the Hodge POVM of Theorem 3.2 to the phase frame in that sector.

Proof. The output vector is

$$|\psi_{ab}\rangle = \frac{1}{\sqrt{r}}\sum_{x=0}^{r-1}\omega^{bx}|x+a\rangle_A|x\rangle_R. \tag{15}$$

Different a give mutually orthogonal subspaces. For fixed a, the d vectors are unitarily equivalent to

$$|\varphi_b\rangle = \frac{1}{\sqrt{r}}(1, \omega^b, \dots, \omega^{(r-1)b})^T \in \mathbb{C}^r.$$

They form a unit-norm tight frame with bound d/r. Every $r \times r$ minor of the synthesis matrix is a Vandermonde determinant on distinct dth roots of unity, so the frame is full spark [12]. Theorem 3.2 supplies a perfect within-sector list of size $d-r+1$, and the shift measurement identifies a without error. This proves sufficiency.

For necessity, let M_L be any effect of a perfect global list measurement. Replacing every effect by its pinching into the orthogonal shift sectors preserves positivity, completeness, and all probabilities of the states in Eq. (15); hence we may take every effect block diagonal. In each sector, L contains at most ℓ phase labels. If $\ell \le d-r$, the effect's block must annihilate at least r phase states; full spark forces that block to vanish. Every block vanishes, hence $M_L = 0$. Since all effects would vanish, no POVM exists. Thus Eq. (14) is necessary. □

6.1 The complete rank-two Bayes curve

For $r = 2$, the phase ensemble in each shift sector is the regular equatorial d-gon

$$|\varphi_b\rangle = \frac{|0\rangle + \omega^b|1\rangle}{\sqrt{2}}.$$

Lemma 6.2 (Largest sum of polygon vertices)**.** *For $0 \le k \le d$,*

$$\max_{S\subseteq\mathbb{Z}_d,\ |S|=k}\left|\sum_{b\in S}\omega^b\right| = \frac{\sin(\pi k/d)}{\sin(\pi/d)},$$

with the right-hand side interpreted as zero for $k = 0, d$. A consecutive arc attains the maximum. Endpoint projection ties may give more than one maximizing arc but do not change the value.

Proof. Let S maximize the modulus and, when its sum is nonzero, let u be its unit complex direction. Then

$$\left|\sum_{b\in S}\omega^b\right| = \sum_{b\in S}\Re(\overline{u}\omega^b).$$

Replacing S by the k vertices with largest projection on u cannot decrease this expression, while the modulus of their sum is at least that projection. The k largest projections on any direction form a consecutive arc of the regular polygon, with only the stated endpoint ambiguity. A geometric-series evaluation of such an arc gives the displayed sine ratio for $1 \le k < d$. When $k > d/2$, the same value also follows by complementing: the sum over the complement is the negative of the original sum because all d roots sum to zero. The cases $k = 0, d$ are immediate. $\square$

Minimum-error discrimination of geometrically uniform states and the associated square-root measurement are classical [21]. The following result solves the different list-valued objective for every list size.

Theorem 6.3 (Exact rank-two Weyl list curve)**.** *For the complete uniform d^2 Weyl ensemble and the consecutive-support probe $|\Phi_2\rangle$, the optimal one-use success probability with lists of size at most ℓ is*

$$P^\star_{\mathrm{succ}}(d,\ell) = \frac{1}{d}\left(\ell + \frac{\sin(\pi\ell/d)}{\sin(\pi/d)}\right), \qquad 1 \le \ell \le d. \tag{16}$$

In particular $P^\star_{\mathrm{succ}} < 1$ for $\ell \le d-2$ and equals one for $\ell = d-1, d$, recovering Theorem 6.1 at $r = 2$.

Proof. Pinch any global measurement into the orthogonal shift sectors as in the proof of Theorem 6.1. For a set $L \subseteq \mathbb{Z}_d$ of phase labels, $|L| \le \ell$, put

$$A_L = \sum_{b\in L}|\varphi_b\rangle\langle\varphi_b| = \frac{1}{2}\begin{pmatrix} |L| & \sum_{b\in L}\omega^{-b} \\ \sum_{b\in L}\omega^b & |L| \end{pmatrix}.$$

Hence

$$\lambda_{\max}(A_L) = \frac{1}{2}\left(|L| + \left|\sum_{b\in L}\omega^b\right|\right).$$

Lemma 6.2 shows that the largest possible root sum has magnitude

$$R_{d,\ell} = \frac{\sin(\pi\ell/d)}{\sin(\pi/d)}. \tag{17}$$

For the upper bound, write $L_a = \{b : (a,b) \in L\}$. If $|L_a| < \ell$, padding it to ℓ phase labels can only increase the positive operator A_{L_a}. Hence each pinched effect block $M_L^{(a)}$ obeys

$$\mathrm{Tr}\Big(M_L^{(a)}A_{L_a}\Big) \le \frac{\ell + R_{d,\ell}}{2}\,\mathrm{Tr}\,M_L^{(a)}.$$

Summing over lists and the d two-dimensional shift sectors gives

$$P_{\mathrm{succ}} \le \frac{1}{d^2}\frac{\ell + R_{d,\ell}}{2}\sum_{a,L}\mathrm{Tr}\,M_L^{(a)} = \frac{1}{d^2}\frac{\ell + R_{d,\ell}}{2}(2d),$$

which is Eq. (16).

If $\ell = d$, the single effect $\mathbb{I}_2$ labelled by the full list attains one, so suppose $\ell < d$. Choose a consecutive base list L_0 and a unit eigenvector v_0 of A_{L_0} at its largest eigenvalue. Translate both by the cyclic phase action: $L_t = L_0 + t$ and $v_t = \mathrm{diag}(1, \omega^t)v_0$. Every v_t is equatorial and $\sum_{t=0}^{d-1}|v_t\rangle\langle v_t| = (d/2)\mathbb{I}_2$, so $M_t = (2/d)|v_t\rangle\langle v_t|$ is a POVM. It attains the largest eigenvalue for every translated list. Applying this covariant list POVM after the projective shift measurement proves equality in every sector and hence globally. $\square$

The endpoints have a direct meaning:

$$r = 1: \ \ell_{\min} = d, \qquad r = d: \ \ell_{\min} = 1.$$

Each additional flat Schmidt coefficient removes exactly one label from the zero-error list. Rearranging Eq. (14) gives the exact fixed-family memory witness

$$r \geq d - \ell + 1.$$

Here r is the Schmidt rank of the specified probe, not a claim about a general quantum-memory cost.

Remark 6.4 (The support pattern matters)**.** *The frontier is not determined by Schmidt rank alone. If $r \mid d$ and $q = d/r$, the equally entangled arithmetic-support probe*

$$|\widetilde{\Phi}_r\rangle = \frac{1}{\sqrt{r}} \sum_{t=0}^{r-1} |qt\rangle_A |t\rangle_R$$

produces mutually orthogonal sectors labelled by a and $b \bmod r$; each output vector is shared by exactly the q labels with the same residue class. Measuring those sectors therefore gives the exact threshold $\ell_{\min} = q = d/r$. Necessity follows because the q corresponding labels produce identical output states. Thus, for example, $(d, r) = (4, 2)$ gives $\ell_{\min} = 2$ for this probe but 3 for Eq. (13). Theorem 6.1 is consequently a support-pattern theorem. The next corollary shows that the arithmetic pattern is globally optimal in the divisor branch.

Corollary 6.5 (Optimized Weyl divisor branch)**.** *Suppose $r \mid d$. Among all pure probes of Schmidt rank exactly r for the complete uniform d^2 Weyl ensemble,*

$$\min_{\psi:\, \mathrm{SR}(\psi) = r} \ell_{\min}(\psi) = \frac{d}{r}. \tag{18}$$

The arithmetic-support probe in Remark 6.4 is optimal. The lower bound specializes the support-dimension mechanism already used for complete unitary error bases in [10]; the new point here is its matching arithmetic-support construction for every divisor pair.

Proof. Fix any such probe and let $\mathcal{K}$ be the span of its d^2 pure output vectors. Since the channel output has dimension d and the probe's reference support has dimension r, $D = \dim \mathcal{K} \leq dr$. Compress any perfect list POVM to $\mathcal{K}$, so that $\sum_L M_L = \mathbb{I}_{\mathcal{K}}$. Writing ρ_{ab} for the output states and using $\sum_{(a,b)\in L} \rho_{ab} \preceq |L| \mathbb{I}_{\mathcal{K}}$ gives

$$\begin{aligned} 1 &= \frac{1}{d^2} \sum_L \operatorname{Tr}\left(M_L \sum_{(a,b)\in L} \rho_{ab} \right) \\ &\leq \frac{\ell}{d^2} \sum_L \operatorname{Tr} M_L = \frac{\ell D}{d^2} \leq \frac{\ell r}{d}. \end{aligned}$$

Thus every rank-r probe needs $\ell \geq d/r$. The arithmetic-support probe attains equality by Remark 6.4. $\square$

7 Necessity of strict scalability and full spark

7.1 Tight but not full spark

In $\mathbb{C}^2$, take

$$f_1 = f_2 = e_1, \qquad f_3 = f_4 = e_2.$$

This is a unit-norm tight frame but not full spark. The basis measurement reports $\{1, 2\}$ or $\{3, 4\}$, so $\ell = 2$ is perfect, below $N - r + 1 = 3$. Full spark cannot be removed from the converse.

7.2 Full spark but not strictly scalable

Take the three qubit vectors

$$f_1 = |0\rangle, \qquad f_2 = \frac{3|0\rangle + |1\rangle}{\sqrt{10}}, \qquad f_3 = \frac{3|0\rangle - |1\rangle}{\sqrt{10}}.$$

Every pair is independent, so the frame is full spark. All three Bloch vectors have strictly positive z component. A list-two effect must omit at least one label and is therefore supported on the line orthogonal to the omitted state. Its Bloch vector has strictly negative z component. No positive combination of such nonzero effects can sum to the identity, whose Bloch vector is zero. Hence list two is impossible although $N - r + 1 = 2$. The same hemisphere argument shows that no strictly positive rescaling of these rays is tight: every weighted Bloch sum still has positive z component. Strict scalability cannot be removed from attainment. This is the qubit cone criterion of [2] in the present list language.

7.3 Other boundaries

Zero-prior labels must be removed before defining N. A nonzero effect cannot annihilate a full-rank state, so replacing frame vectors by arbitrary mixed states is invalid. Near-tight frames do not inherit exact zero error by continuity: feasibility is a closed conic condition with exact orthogonality constraints. Theorem 4.1 instead supplies a quantitative floor under bounded state perturbations, while Theorem 6.3 gives an exact subthreshold curve only for the specified rank-two Weyl family.

8 Relation to prior work and scope

Table 2 states the contribution boundary. General exclusion SDPs, antidistinguishability, group-covariant one-state exclusion, exterior tightness, full-spark harmonic frames, and the equality between learning width and Gram-matrix factor width, as well as the notion of factor-width rank, all predate this work. In particular, factor-width rank records the minimum number of sparse rank-one terms [22]; the r^2 bound below is the standard conic-Carathéodory consequence in the real vector space of Hermitian operators on the r-dimensional support. In particular, the group-generated solution of [5] concerns an outcome that excludes its matching single label. Our threshold decoder instead associates an outcome with $r - 1$ simultaneously excluded labels and solves the smallest retained list over an arbitrary strictly scalable full-spark ray ensemble. The channel-capacity results of [6] are asymptotic in block length; our result is a one-shot, fixed-codebook feasibility law. The companion record [10] already proves a general unitary-error-basis list cap and the single $(d, r, \ell) = (4, 2, 2)$ subthreshold fixture. Our formulas strictly generalize that reported point; we do not claim it again as an independent contribution. The new Weyl content here is the exact all-(d, r) zero-error threshold for the specified consecutive-support probe, its Hodge decoder, and the exact optimization over all rank-r pure probes when $r \mid d$, together with the complete rank-two list-valued Bayes curve.

We make no "first ever" priority claim. The general feasibility criterion is explicitly credited to [9]. Targeted comparison with the primary sources above did not locate the scalable all-frame evaluation $N - r + 1$, the complementary weighted Hodge decoder, the consecutive-support Weyl frontier $d - r + 1$, the global divisor-branch optimum d/r, or the exact rank-two list curve. The novelty claim is restricted to those evaluations and constructions, not to factor width, exclusion, exterior algebra, covariant discrimination, or dense coding individually.

9 Reproducibility

The accompanying lightweight NumPy replay checks seven finite layers:

1. harmonic frames are tight and full spark for a grid of small (N, r);

Table 2: Contribution boundary. "Used" means explicitly credited rather than claimed as new.

Topic	Status here	New synthesis claimed here
Learning width = Gram factor width	Used	Exact evaluation $N - r + 1$ on the scalable full-spark branch
Exclusion SDP and optimality	Used	Exact scalable full-spark list threshold
Pure-state annihilator/factor cones and factor-width rank	Used; the r^2 compression is a standard Carathéodory corollary	Weighted Hodge factorization and exact full-spark width
Cross-products/exterior powers of tight frames	Used	Hodge projectors as a threshold POVM
Full-spark Fourier/Vandermonde frames	Used	Exact consecutive-support Weyl frontier
Dense coding and Weyl channels	Used	Fixed-probe frontier, global divisor optimum, and rank-two Bayes curve
Geometrically uniform state discrimination	Used	Exact list-valued regular-polygon curve

2. exterior Hodge effects reproduce $\alpha^{r-1}\mathbb{I}$, including a non-unit-norm tight full-spark fixture;
3. nullspace elimination compresses a finite Hodge POVM to at most r^2 effects while retaining the identity resolution;
4. consecutive- and arithmetic-support Weyl fixtures agree with the closed formulas, including the divisor-branch optimum;
5. the rank-two Weyl Bayes formula agrees with explicit covariant POVMs for a grid of dimensions and list sizes;
6. the spectral subthreshold floor and its perturbation inequality hold on finite harmonic-frame fixtures;
7. both hypothesis counterexamples behave as stated.

The replay is diagnostic evidence, not a proof of the all-dimensional theorems. Those follow from the exterior-power, full-spark, and Weyl arguments in Sections 3 and 6. No formal-proof-assistant certification, peer review, or independent scientific validation is claimed.

10 Conclusion

For quantum list decoding, linear independence alone is an obstruction and positivity is the construction problem. Strictly scalable full-spark ray ensembles make the two meet exactly: full spark forces $\ell \geq N - r + 1$, while the $(r-1)$st exterior power of tightness builds the matching POVM. The known factor-width criterion, in physical-space annihilator form, explains why this construction is physical; standard conic Carathéodory explains the r^2-outcome compression.

The Weyl application turns the abstract threshold into a sharp resource frontier for a fixed family of partially entangled probes. In the divisor branch, a dimension converse closes the global rank-r probe optimization at d/r. For rank two, the regular-polygon geometry closes the entire one-shot list-success curve, while a simple spectral witness quantifies failure below threshold for arbitrary full-spark ensembles and survives bounded state perturbations. The construction also supplies an infinite separation between support-only feasibility tests and actual zero-error measurements. Natural open problems are to close the nondivisible probe optimum, obtain exact subthreshold curves beyond rank two, and find other represented matroids whose physical-space factor cone admits a closed positive identity resolution.

AI assistance disclosure. OpenAI Codex assisted with theorem exploration, literature triage, proof stress testing, verifier development, and manuscript editing. The author selected and reviewed the claims, proofs, citations, and computational evidence and remains responsible for the work. AI output is not treated as independent scientific evidence.

References

[1] Somshubhro Bandyopadhyay, Rahul Jain, Jonathan Oppenheim, and Christopher Perry. Conclusive exclusion of quantum states. *Physical Review A*, 89:022336, 2014.

[2] Teiko Heinosaari and Oskari Kerppo. Antidistinguishability of pure quantum states. *Journal of Physics A: Mathematical and Theoretical*, 51:365303, 2018.

[3] Vincent Russo and Jamie Sikora. Inner products of pure states and their antidistinguishability. *Physical Review A*, 107:L030202, 2023.

[4] Nathaniel Johnston, Vincent Russo, and Jamie Sikora. Tight bounds for antidistinguishability and circulant sets of pure quantum states. *Quantum*, 9:1622, 2025.

[5] Arnau Diebra, Santiago Llorens, Emili Bagan, Gael Sentis, and Ramon Munoz-Tapia. Quantum state exclusion for group-generated ensembles of pure states. *Physical Review Research*, 8:L012001, 2026.

[6] Marco Dalai, Filippo Girardi, and Ludovico Lami. Zero-error list decoding for classical–quantum channels, 2026. version 2.

[7] Jonathan Crickmore, Ittoop V. Puthoor, Berke Ricketti, Sarah Croke, Mark Hillery, and Erika Andersson. Unambiguous quantum state elimination for qubit sequences. *Physical Review Research*, 2:013256, 2020.

[8] Jonathan W. Webb, Ittoop V. Puthoor, Joseph Ho, Jonathan Crickmore, Emma Blakely, Alessandro Fedrizzi, and Erika Andersson. Experimental demonstration of optimal unambiguous two-out-of-four quantum state elimination. *Physical Review Research*, 5:023094, 2023.

[9] Nathaniel Johnston, Benjamin Lovitz, Vincent Russo, and Jamie Sikora. The complexity of perfect quantum state classification, 2025.

[10] Lluis Eriksson. Bayesian matroid-union bounds for quantum list discrimination: Support congestion, process compression, and exact adaptive-parallel phases, 2026. ARR-2026-15SJ1ANHDN8D88Z1, Archive for Rigorous Research.

[11] Gitta Kutyniok, Kasso A. Okoudjou, Friedrich Philipp, and Elizabeth K. Tuley. Scalable frames. *Linear Algebra and its Applications*, 438(5):2225–2238, 2013.

[12] Boris Alexeev, Jameson Cahill, and Dustin G. Mixon. Full spark frames. *Journal of Fourier Analysis and Applications*, 18(6):1167–1194, 2012.

[13] Peter G. Casazza and Gitta Kutyniok, editors. *Finite Frames: Theory and Applications*. Birkhauser, 2013.

[14] Grigory Ivanov. Tight frames and related geometric problems. *Canadian Mathematical Bulletin*, 64(4):942–963, 2021.

[15] Benjamin Stratton, Chung-Yun Hsieh, and Paul Skrzypczyk. Operational interpretation of the choi rank through exclusion tasks. *Physical Review A*, 110:L050601, 2024.

[16] Alexander S. Holevo. Statistical decision theory for quantum systems. *Journal of Multivariate Analysis*, 3:337–394, 1973.

[17] Horace P. Yuen, Robert S. Kennedy, and Melvin Lax. Optimum testing of multiple hypotheses in quantum detection theory. *IEEE Transactions on Information Theory*, 21(2):125–134, 1975.

[18] Charles H. Bennett and Stephen J. Wiesner. Communication via one- and two-particle operators on einstein–podolsky–rosen states. *Physical Review Letters*, 69:2881–2884, 1992.

[19] Shay Mozes, Jonathan Oppenheim, and Benni Reznik. Deterministic dense coding with partially entangled states. *Physical Review A*, 71:012311, 2005.

[20] Shengjun Wu, Scott M. Cohen, Yuqing Sun, and Robert B. Griffiths. Deterministic and unambiguous dense coding. *Physical Review A*, 73:042311, 2006.

[21] Yonina C. Eldar and G. David Forney, Jr. On quantum detection and the square-root measurement. *IEEE Transactions on Information Theory*, 47(3):858–872, 2001.

[22] Nathaniel Johnston, Shirin Moein, and Sarah Plosker. The factor width rank of a matrix, 2025.