TASK-AWARE CERTIFIED BLOCK FRAMES FOR EXACT LATENT RECOVERY

Anonymous authors
Paper under double-blind review

ABSTRACT

Most block-erasure designs optimize reconstruction stability. With erasure locations known at decoding, every full-column-rank survivor frame recovers the same latent representation in the zero-noise case. Under bounded observation noise, however, survivor geometry determines the amplification of each class-pair score difference. We exploit this remaining design freedom to reduce decision uncertainty while preserving exact recovery of the full latent representation at zero noise. For a linear head and a total ℓ_2 observation-noise budget, we derive the worst-case pairwise score uncertainty and the largest radius that preserves the multiclass prediction over all survivor sets; an explicit perturbation reaches its boundary. The resulting certificate identifies the decision directions that the block design must protect. **T**ask-**A**ware **C**ertified **B**lock **F**rames (TACBF) combines closed-form task-subspace allocation with learned rank-constrained block geometry. Across held-out test sets from seven datasets, TACBF improves both grid-based and continuous-radius measures of the area under the certified-accuracy curve (AUC) relative to a fixed-factor control paired by training seed. Continuous-radius AUC improves for every paired seed on all seven datasets, and TACBF has the highest mean grid AUC among seven methods on every dataset. The method preserves exact zero-noise recovery, transmitted width, and trace without changing the inference procedure; dataset-level mean clean accuracy differs by at most 0.25 percentage points.

1 INTRODUCTION

Redundant block encodings preserve latent content after packet erasures, yet conventional spectral, coherence, and reconstruction criteria leave downstream decision geometry uncontrolled. A well-conditioned survivor Gram therefore provides no guarantee against amplification of a class-separating direction.

With known erasure locations and full-rank least-squares (LS) decoding, every admissible survivor frame has the same zero-noise reconstruction s and therefore the same clean prediction. Under bounded observation noise, the inverse-Gram quadratic determines class-pair sensitivity. This invariance permits task-aware frame design without sacrificing recovery of any latent coordinate.

These observations motivate TACBF, which converts the exact certificate into a feasible block-frame design.

We establish a unified certificate-to-design framework with three components:

- **Exact certification:** we derive exact pairwise uncertainty and the exact all-survivor radius, with an explicit perturbation attaining its boundary, then reduce sample-wise worst-survivor evaluation to reusable class-pair geometry.
- **Certificate-to-design theory:** we uniquely solve the fixed-trace relaxation, derive exact scaling and task–reconstruction identities, and prove the unavoidable rank-constrained gap.
- **Feasible realization:** D^* allocates trace analytically to task directions while B optimizes the residual rank-constrained geometry; controlled comparisons show that both analytic allocation and learned block geometry improve certification.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Fig. 1 contrasts exact zero-noise recovery across admissible survivor sets with the decision-relevant geometry induced by bounded observation noise.

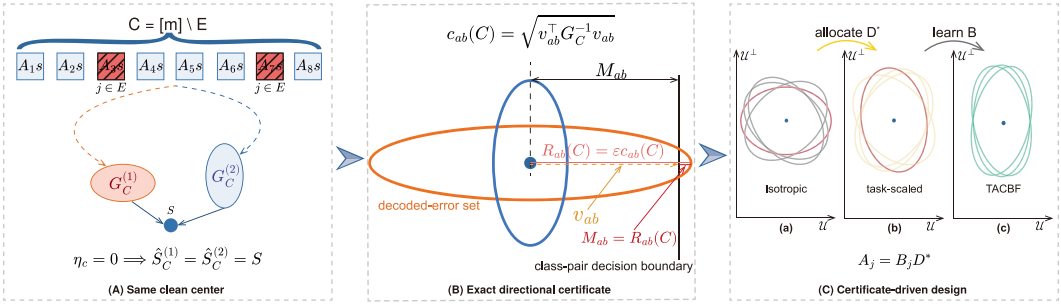


Figure 1: Certificate geometry and task-aware design. (a) Each admissible survivor set $C = [m] \setminus E$ yields the same noiseless LS reconstruction. (b) Bounded observation noise induces an LS reconstruction-error set whose support in the class-pair direction v_{ab} is $R_{ab}(C) = \epsilon c_{ab}(C)$. The set touches the class-pair decision boundary when $R_{ab}(C)$ equals the class-pair margin $M_{ab} = v_{ab}^\top s$. (c) D^* allocates the trace budget across $\mathcal{U}$ and $\mathcal{U}^\perp$; the learned Parseval factor B then shapes survivor geometry subject to the block-rank constraints.

2 Related Work

Frames under erasures. Classical frame and fusion-frame theories study reconstruction after coefficient, sampling, or subspace erasures (Goyal et al., 2001; Casazza & Kovačević, 2003; Larson & Scholze, 2015; Casazza et al., 2008; Casazza & Kutyniok, 2008; Köhldorfer et al., 2023). Erasure-optimal, Grassmannian, and Frame Codes constructions target spectra, distortion, coherence, or block-channel criteria (Holmes & Paulsen, 2004; Strohmer & Heath, 2003; Kutyniok et al., 2009; Jacoby & Zamir, 2024). Packet-oriented frame constructions further connect erasure robustness to graph structure and loss-insensitive linear transmission (Bodmann & Paulsen, 2005; Bodmann, 2007). Reconstruction-system analyses also characterize erasure-optimal protocols through frame potentials and spectral conditions (Massey, 2009). Among these approaches, Frame Codes is the closest block-channel formulation, but it targets reconstruction and channel criteria rather than decision robustness. Within the known-survivor block-erasure model, TACBF instead derives an exact decision-level objective for sample-wise prediction under bounded observation noise after LS decoding.

Task-aware communication. Task-oriented communication aligns coding and resource allocation with inference through information bottlenecks, task-based quantization, discriminative precoding, and rate–distortion–accuracy tradeoffs (Gündüz et al., 2023; Shao et al., 2022; Pezone et al., 2022; Shlezinger & Eldar, 2020; Cai et al., 2024; Kubiak & Hadfield, 2021; Kutay & Yener, 2024; Ma et al., 2025; Singh & Arif, 2026). Neural block-erasure semantic coding likewise optimizes reconstruction or channel distortion (Esfahanizadeh et al., 2026; Fayaz et al., 2026). TACBF unifies task-aware resource allocation with exact recovery of the full latent representation: it preserves every latent coordinate under fixed width and trace while certifying robustness to worst-case observation noise under every admissible erasure pattern. Task selectivity arises from both directional trace allocation in D^* and the distribution of those directions across blocks in B.

Certification and geometry. Smoothing and margin–Lipschitz methods certify robustness to input perturbations (Hein & Andriushchenko, 2017; Cohen et al., 2019; Tsuzuku et al., 2018). For each known survivor set, the LS error map yields an exact support-function certificate attained by an explicit perturbation. Grassmannian and alternating-projection methods separate subspaces (Dhillon et al., 2008), fusion-frame homotopy tightens frames on a manifold (Needham & Shonkwiler, 2023), and Parseval constraints control neural-layer geometry (Cisse et al., 2017). Here the Parseval condition parameterizes feasible block arrangements, whereas the certificate supplies the task-specific objective used to choose among them.

3 Problem Setup

For a positive integer L, write $[L] = \{1, \dots, L\}$. Let x have label $y \in [K]$, and let $s = f_\theta(x) \in \mathbb{R}^d$ be a unit-normalized latent representation. A block frame has operators $A_j \in \mathbb{R}^{r \times d}$, $j \in [m]$, and stacked operator $A \in \mathbb{R}^{mr \times d}$; its transmitted width is mr and its redundancy ratio is mr/d. After q erasures, $C \subseteq [m]$ indexes the $n = m - q$ survivors, A_C stacks them, and $G_C = A_C^\top A_C$ is the survivor Gram. Assume every admissible A_C has full column rank. The receiver observes

$$\widetilde{z}_C = A_C s + \eta_C, \qquad \|\eta_C\|_2 \le \epsilon, \qquad \widehat{s}_C = G_C^{-1} A_C^\top \widetilde{z}_C. \tag{1}$$

At $\eta_C = 0$, $\widehat{s}_C = s$: this is exact recovery of the full latent representation, including all task-orthogonal coordinates. We fix the data-independent trace budget $\mathrm{tr}(A^\top A) = \sum_j \|A_j\|_F^2 = P$. The frame survivor floor is

$$\alpha_q(A) = \min_{|C|=m-q} \lambda_{\min}(A_C^\top A_C). \tag{2}$$

Thus $\alpha_q(A) > 0$ is equivalent to exact recovery for every admissible survivor set; we impose the required survivor-floor threshold $\alpha_q(A) \ge \mu > 0$. Complete notation and the elementary feasibility condition $(m - q)r \ge d$ are given in Sec. A and Lemma 10.

The trace fixes the encoder's total squared gain: for $u \sim \mathrm{Unif}(\mathbb{S}^{d-1})$, $\mathbb{E}\|Au\|_2^2 = P/d$. The trace budget is distinct from the sample-wise codeword energy $\|As\|_2^2 = s^\top A^\top A s$: trace is a data-independent design resource, whereas codeword energy also depends on the latent distribution. The floor μ enforces full-space stability, while the certificate quantifies the task-direction uncertainty minimized by the design.

4 Exact Robustness Certificate

Let $W \in \mathbb{R}^{K \times d}$ be a bias-free linear classification head with rows $w_k^\top$. For classes $a \ne b$, define the pairwise decision direction $v_{ab} = w_a - w_b$.

Theorem 1 (Exact pairwise score uncertainty) *If A_C has full column rank and $\epsilon \ge 0$, then*

$$\sup_{\|\eta_C\|_2 \le \epsilon} \left| v_{ab}^\top (\widehat{s}_C - s) \right| = \epsilon \sqrt{v_{ab}^\top (A_C^\top A_C)^{-1} v_{ab}}. \tag{3}$$

For $v_{ab} \ne 0$, the bound is attained by aligning the observation noise with $A_C (A_C^\top A_C)^{-1} v_{ab}$.

The appendix gives the proof and an explicit perturbation attaining the bound (Sec. B.1). When at most q blocks may be erased, the worst coefficient is attained among sets with exactly q erasures. Define the score uncertainty $R_{ab}(C; \epsilon) := \epsilon c_{ab}(C)$, where $c_{ab}(C) = \sqrt{v_{ab}^\top G_C^{-1} v_{ab}}$, and its reusable worst-survivor value

$$\overline{c}_{ab,q} := \max_{|C|=m-q} c_{ab}(C), \qquad M_{ik}^0 := v_{y_i k}^\top s_i. \tag{4}$$

For a sample with strictly positive pairwise margins $M_{ik}^0 > 0$ for every $k \ne y_i$, the largest radius preserving its prediction for every competitor and survivor set is exactly

$$\epsilon_i^* = \min_{k \ne y_i} \frac{M_{ik}^0}{\overline{c}_{y_i k, q}}. \tag{5}$$

Only the clean margin varies across samples; $\overline{c}_{ab,q}$ is shared. Since $c_{ab} = c_{ba}$, this reduces evaluation from one combinatorial search per sample to one per unordered class pair. Reusing these coefficients enables class-pair branch-and-bound that returns certified radius intervals at any search budget and the exact worst-survivor coefficient at termination (Sec. B.2).

5 Task-Aware Block-Frame Design

Relaxing the block-rank constraints yields the optimal allocation; Parseval factorization then exposes the residual rank-constrained geometry optimized through B.

Fix W, define $\bar{w} := K^{-1}\sum_{k=1}^{K} w_k$, and let $\mathcal{U} = \mathrm{span}\{w_k - \bar{w}\}_{k=1}^{K}$ be its p-dimensional task subspace, with orthonormal basis U. We assume $1 \le p \le d$ throughout. Define the spectral task envelope and class-pair envelope

$$J(A) = \max_{|C|=n} \lambda_{\max}(U^\top G_C^{-1} U), \qquad J_{\mathrm{pair}}(A) = \max_{|C|=n} \max_{\widehat{v}\in\mathcal{D}} \widehat{v}^\top G_C^{-1} \widehat{v}, \tag{6}$$

where $\mathcal{D} = \{v_{ab}/\|v_{ab}\|_2 : a < b,\ v_{ab} \neq 0\}$. The spectral envelope supports the closed-form allocation, whereas the class-pair envelope is the normalized geometry used by the certificate.

Proposition 2 (Class-pair coverage bound) *Index $\mathcal{D}$ by $\widehat{v}_\ell$, set $z_\ell = U^\top \widehat{v}_\ell$, and define $\kappa_W = \min_{\|u\|_2=1} \max_\ell (u^\top z_\ell)^2 > 0$. Then*

$$\kappa_W J(A) \le J_{\mathrm{pair}}(A) \le J(A). \tag{7}$$

Because the class-pair directions span $\mathcal{U}$, κ_W is positive. Proposition 2 implies that the spectral minimizer is within a factor $1/\kappa_W$ of the class-pair optimum; computable lower bounds on κ_W quantify this factor (Sec. B.4).

5.1 Relaxed allocation

Write $H_j = A_j^\top A_j \succeq 0$, so $G_C = \sum_{j\in C} H_j$, and temporarily remove the block-rank constraint. Under the trace budget and $G_C \succeq \mu I$, the ideal allocation is unique.

Theorem 3 (Rank-relaxed task allocation) *For $1 \le p < d$, $0 < \mu < nP/(md)$, and $\mathcal{B} = nP/m - (d-p)\mu$,*

$$\min_{\substack{H_j\succeq 0,\ \sum_j \mathrm{tr}(H_j)=P \\ G_C\succeq \mu I\ \forall C}} \max_{|C|=n} \lambda_{\max}(U^\top G_C^{-1} U) = \frac{p}{\mathcal{B}}. \tag{8}$$

Equality uniquely determines every survivor Gram as

$$G_C = G_* := \frac{\mathcal{B}}{p} UU^\top + \mu(I - UU^\top). \tag{9}$$

When $1 \le q \le m-1$, this further forces $H_j = G_*/n$ for every block.

The relaxed optimum assigns the minimum allowed energy to $\mathcal{U}^\perp$ and the remaining trace uniformly within $\mathcal{U}$. Because G_*/n is full rank, rank-r blocks cannot attain the relaxation; B therefore controls the unavoidable gap. We write $J_{\mathrm{relax}} := p/\mathcal{B}$ for this rank-relaxed lower bound.

Relative to the isotropic survivor Gram at the same trace $\bar{t} = nP/m$, G_* reduces task uncertainty but increases full-space LS noise amplification. For this isotropic-to-G_* comparison, the exact exchange is

$$\Delta_{\mathrm{rec}} = \frac{\bar{t} - d\mu}{\mu}\, \Delta_{\mathrm{task}}, \tag{10}$$

with the component definitions and proof in Proposition 13.

A head-specific p-dimensional semidefinite program (SDP) gives the class-pair relaxation optimum

$$\gamma_{\mathrm{pair}}^{\mathrm{rel}} := \min_{\substack{T\succeq \mu I_p \\ \mathrm{tr}(T)=\mathcal{B}}} \max_\ell z_\ell^\top T^{-1} z_\ell \le p/\mathcal{B}. \tag{11}$$

Equality holds exactly when the normalized class-pair directions admit an isotropic weighting. We use the closed-form spectral solution as the training rule, while the class-pair relaxation provides an exact classifier-specific reference (Proposition 12). The directional inverse-information quantity $z^\top T^{-1} z$ is the classical c-optimal design criterion within optimal-design equivalence theory. Constructing information matrices from discrete subsets also connects this problem to sensor selection (Elfving, 1952; Kiefer, 1974; Joshi & Boyd, 2009).

5.2 FEASIBLE ANISOTROPIC LIFTING

Every full-column-rank frame has a unique right-polar factorization into a global symmetric positive-definite (SPD) allocation and a Parseval factor (Higham, 1986):

$$A = BD, \qquad D = (A^\top A)^{1/2} \succ 0, \qquad B = AD^{-1}, \qquad B^\top B = I, \tag{12}$$

where D allocates energy across latent directions, whereas B distributes those directions across blocks (Lemma 15). Its inherited blocks satisfy $A_j = B_j D$, $G_C = D(B_C^\top B_C)D$, and $\operatorname{tr}(A^\top A) = \operatorname{tr}(D^2)$. This factorization lets TACBF impose D^* while retaining the full family of rank-constrained Parseval factors B.

Let B_j be the blocks of a rank-constrained Parseval factor: its stacked matrix satisfies $B^\top B = I$. Assume its Parseval-factor survivor floor $\beta = \min_{|C|=n} \lambda_{\min}(B_C^\top B_C)$ is positive. For the task projector $\Pi = UU^\top$, choose

$$\tau = \mu/\beta, \quad \rho = \frac{P - (d-p)\tau}{p}, \quad D^* = \sqrt{\rho}\,\Pi + \sqrt{\tau}(I - \Pi), \quad A_j = B_j D^*. \tag{13}$$

Theorem 4 (Optimal anisotropic allocation under a scalar factor floor) *For $1 \le p \le d$ and $0 < \mu < \beta P/d$, A satisfies the trace budget and $\alpha_q(A) \ge \mu$. For every $v \in \mathcal{U}$ and survivor set C with $|C| = n$,*

$$v^\top (A_C^\top A_C)^{-1} v = \rho^{-1} v^\top (B_C^\top B_C)^{-1} v. \tag{14}$$

Given only the scalar bound $M \succeq \beta I$, an arbitrary SPD allocation has the worst-case spectral task envelope

$$\mathcal{R}_\beta(D) := \sup_{M \succeq \beta I} \lambda_{\max}\left(U^\top D^{-1} M^{-1} D^{-1} U\right) = \beta^{-1} \lambda_{\max}\left(U^\top D^{-2} U\right). \tag{15}$$

Among all SPD D satisfying $\operatorname{tr}(D^2) = P$ and $\beta D^2 \succeq \mu I$, D^* uniquely minimizes $\mathcal{R}_\beta(D)$ and attains $\mathcal{R}_\beta(D^*) = 1/(\beta\rho)$.

Thus a single scalar β suffices to obtain a uniform closed-form minimax allocation; the class-pair SDP in Eq. (11) gives a tighter classifier-specific reference by using the classifier's class-pair directions.

Relative to $D_{\text{iso}} = \sqrt{P/d}\, I$ with the same fixed factor, define

$$\Gamma_{\text{alloc}} := \frac{\rho}{P/d} = \frac{d}{p}\left(1 - \frac{(d-p)\mu}{\beta P}\right). \tag{16}$$

Holding the encoder, latent samples, B, and W fixed, the sample-wise robust radii and the full certified-accuracy curve scale exactly as

$$\epsilon_i^*(D^*) = \sqrt{\Gamma_{\text{alloc}}}\, \epsilon_i^*(D_{\text{iso}}), \qquad \text{CertAcc}_{\text{all}}^{D^*}(\epsilon) = \text{CertAcc}_{\text{all}}^{D_{\text{iso}}}\left(\epsilon/\sqrt{\Gamma_{\text{alloc}}}\right). \tag{17}$$

For every absolute-radius interval $[t_0, t_{\max}]$, continuous-radius AUC transforms exactly as

$$\int_{t_0}^{t_{\max}} \text{CertAcc}_{\text{all}}^{D^*}(\epsilon) d\epsilon = \sqrt{\Gamma_{\text{alloc}}} \int_{t_0/\sqrt{\Gamma_{\text{alloc}}}}^{t_{\max}/\sqrt{\Gamma_{\text{alloc}}}} \text{CertAcc}_{\text{all}}^{D_{\text{iso}}}(u) du. \tag{18}$$

The same factorization also quantifies the corresponding reconstruction tradeoff. For $F_C = B_C^\top B_C$, define $G_{C,D} := DF_C D$ for any allocation D, and let $U_\perp$ be an orthonormal basis of $\mathcal{U}^\perp$.

$$\begin{aligned} \operatorname{tr}(U^\top G_{C,D^*}^{-1} U) &= \rho^{-1} \operatorname{tr}(U^\top F_C^{-1} U), \\ \operatorname{tr}(U_\perp^\top G_{C,D^*}^{-1} U_\perp) &= \tau^{-1} \operatorname{tr}(U_\perp^\top F_C^{-1} U_\perp). \end{aligned} \tag{19}$$

For $p < d$, the task and orthogonal ratios to isotropic scaling are exactly $\Gamma_{\text{alloc}}^{-1} < 1$ and $\beta P/(d\mu) > 1$, with $1 < \Gamma_{\text{alloc}} < d/p$; for $p = d$, $\Gamma_{\text{alloc}} = 1$ and $U_\perp$ is empty. When $p = K - 1$ and the floor correction is small, $\Gamma_{\text{alloc}} \approx d/(K-1)$; maintaining a fixed gain as K grows therefore requires $d \propto K - 1$. For fixed (m, q), maintaining both mr/d and $(m-q)r \ge d$ requires $r \propto d$. Thus D^* provably contracts task-direction uncertainty at the cost of greater reconstruction error on $\mathcal{U}^\perp$. Further scaling and monotonicity properties are in Corollary 16; the Parseval factor-floor identity is Eq. (46).

5.3 RANK-CONSTRAINED GEOMETRY

Let $\mathcal{F}_r$ be the feasible family of positive-semidefinite block Grams with $\operatorname{rank}(H_j) \le r$, total trace P, and every $G_C \succeq \mu I$. The homogeneous relaxed contribution G_*/n is full rank, so rank-constrained blocks cannot realize it.

Theorem 5 (Strict rank-constrained gap) *Let $1 \le p < d$, $1 \le q \le m-1$, $n = m-q$, $0 < \mu < nP/(md)$, and $U \in \mathbb{R}^{d\times p}$ have orthonormal columns. If $r < d$ and $\mathcal{F}_r \neq \emptyset$, then*

$$\min_{(H_j)\in\mathcal{F}_r} \max_{C\subseteq[m],\,|C|=n} \lambda_{\max}\big(U^\top G_C^{-1} U\big) > \frac{p}{\mathcal{B}}. \tag{20}$$

Let $H_* = G_*/n$ have eigenvalues $\lambda_1(H_*) \ge \cdots \ge \lambda_d(H_*)$. Every feasible design also satisfies

$$\left(\frac{1}{\binom{m}{n}} \sum_{|C|=n} \|G_C - G_*\|_F^2\right)^{1/2} \ge \sqrt{\frac{nq}{m-1}} \left(\sum_{\ell=r+1}^{d} \lambda_\ell(H_*)^2\right)^{1/2}. \tag{21}$$

Equality would make every survivor Gram identical and, by exchanging one block, force every H_j to be full rank. Compactness makes the objective gap strict, while the spectral tail quantifies Gram separation. The proof and $r \ge p$ versus $r < p$ specializations are in Sec. B.6; the remaining rank-constrained design problem resides in B.

5.4 FACTORIZED OPTIMIZATION

The strict rank gap leaves a nontrivial survivor-geometry degree of freedom, which TACBF optimizes in B using a sampled approximation to the worst-survivor certificate loss. At the start of each epoch, TACBF recomputes $\mathcal{U}$ and β and updates D^* when $\beta > \mu d/P$; otherwise it retains the current value. Within each minibatch, reduced QR maps the raw block parameters onto the Parseval manifold (Edelman et al., 1998). Training uses noiseless LS reconstructions $\widehat{s}^0_{i,C}$ from $A_C s_i$; observation-noise sensitivity enters through the certificate term $R_{y_i k}$. For a minibatch of size N_{mb} and sampled survivor family $\mathcal{S}$, let $M_{ik}(C) = v_{y_i k}^\top \widehat{s}^0_{i,C}$. The training objective is

$$\frac{1}{N_{\mathrm{mb}}} \sum_i \max_{C\in\mathcal{S}} \mathrm{CE}_{i,C} + \frac{\lambda_{\mathrm{cert}}}{N_{\mathrm{mb}}} \sum_i \max_{\substack{C\in\mathcal{S}\\ k\neq y_i}} [R_{y_i k}(C;\epsilon_{\mathrm{train}}) + \gamma_{\mathrm{tr}} - M_{ik}(C)]_+ + \lambda_{\mathrm{floor}}\mathcal{L}_{\mathrm{floor}}. \tag{22}$$

The three terms train the task model, reduce exact certified-margin violations, and penalize survivor-floor violations, respectively. The fixed-factor control uses the same initialization, schedule, allocation update, and loss with B frozen; full details are in Sec. C.

After training, we recompute B, U, β, verify $\beta > \mu d/P$, and construct $A_j = B_j D^*$. Survivor sampling is confined to training; final certification exhaustively evaluates every admissible set with exact LS, while transmitted width and inference remain unchanged.

6 EXPERIMENTAL SETUP

Our primary end-to-end evaluation fixes $(m, q, r, P) = (8, 2, d/4, d)$. We additionally evaluate $q = 2$ models at $q = 0, \ldots, 3$, train $q = 3$ models on STL-10, and study survivor-search scalability at $m = 16$. The seven held-out benchmarks are CIFAR-10, STL-10, EuroSAT RGB, CIFAR-100, PCam, UCI HAR, and Banking77 (Krizhevsky, 2009; Coates et al., 2011; Helber et al., 2019; Veeling et al., 2018; Anguita et al., 2013; Casanueva et al., 2020). Six use $d = P = 128, r = 32$; CIFAR-100 uses $d = P = 512, r = 128$. At $q = 2$, all seven datasets use $m = 8$ and exhaustively evaluate all 28 survivor sets. Five paired seeds compare TACBF with a fixed-factor control matched in width, trace, allocation rule, certificate loss, schedule, initialization distribution, and LS decoder. For every dataset, the final latent $s = f_\theta(x)$ is ℓ_2-normalized before frame encoding and classification.

CIFAR-10, STL-10, and EuroSAT use a common absolute-ϵ grid. CIFAR-100, PCam, UCI HAR, and Banking77 use $\delta = \epsilon/\sqrt{\overline{E}_{\mathrm{control,train}}}$, where the denominator is the fixed-factor control's mean training-set codeword energy. Both radius rules and grids are chosen using training and development

Table 1: Five-seed paired comparison at $q = 2$; certified accuracy uses $\epsilon = 1.5$ or $\delta = 0.42$. Differences are percentage points where applicable. * marks an unadjusted paired 95% CI excluding zero; full intervals and Holm-adjusted continuous-radius AUC tests are in Tab. 8 and Tab. 11.

Setting	Dataset (K, d)	Grid AUC ↑		Certified accuracy ↑ (%) at $\epsilon = 1.5$ or $\delta = 0.42$	
		TACBF/fixed-factor	Δ	TACBF/fixed-factor (%)	Δ (pp)
Absolute ϵ	CIFAR-10 $(10, 128)$	**1.3606**/1.3516	**+0.0090***	**87.02**/84.92	**+2.10***
	STL-10 $(10, 128)$	**1.0047**/0.8910	**+0.1137***	**47.33**/14.81	**+32.52***
	EuroSAT RGB $(10, 128)$	**1.4171**/1.3882	**+0.0289***	**90.54**/80.52	**+10.01***
Normalized δ	CIFAR-100 $(100, 512)$	**0.3256**/0.2740	**+0.0516***	**51.40**/19.60	**+31.79***
	PCam $(2, 128)$	**0.5511**/0.4770	**+0.0741***	**69.48**/65.28	**+4.21***
	UCI HAR $(6, 128)$	**0.4251**/0.4104	**+0.0148***	**69.30**/61.26	**+8.05***
	Banking77 $(77, 128)$	**0.2091**/0.1781	**+0.0311***	**0.10**/0.00	**+0.10***

data before held-out test evaluation. Paired training seeds are the inferential units; examples and survivor sets are not independent replicates. Grid AUC is the trapezoidal area under the all-survivor certified-accuracy curve, and continuous-radius AUC is the corresponding continuous integral. Range-normalized grid AUC divides grid AUC by the applicable grid length: 1.48 for $\epsilon \in [0.02, 1.5]$ and 0.8 for $\delta \in [0, 0.8]$. The erasure-severity analysis instead uses common-δ AUC, computed for every dataset on the shared $\delta \in [0, 0.8]$ grid and divided by 0.8. We report certified accuracy at $\epsilon = 1.5$ or $\delta = 0.42$ and 95% confidence intervals for paired seed-level differences computed using Student's t distribution. We apply two-sided paired t-tests to five seed-level differences per dataset ($\mathrm{df} = 4$). Holm correction is applied across the seven dataset-level continuous-radius comparisons. We benchmark against six baselines: a Frame-Codes-inspired block unit-norm tight frame (BUTF), analytic task allocation, label-free max–min frame, block dropout, fixed-tight frame, and the matched fixed-factor control. Architectures, splits, grids, hyperparameters, baseline definitions, continuous-radius AUC, and computational costs are specified in Sec. C–Sec. E.

7 Results

7.1 TACBF improves all-survivor certification across seven datasets

TACBF has the highest mean grid AUC among seven methods on every dataset (Tab. 2). The paired fixed-factor control provides the closest comparison because it differs from TACBF only in holding B fixed; width, trace, decoder, allocation rule, objective, and schedule are matched. Learning B improves grid AUC and certified accuracy at $\epsilon = 1.5$ or $\delta = 0.42$ in all seven datasets (Tab. 1). Continuous-radius AUC also improves in all seven; every gain remains significant after Holm correction (Tab. 11).

Analytic task allocation outperforms the fixed-tight frame on all seven datasets, and learning B further improves on the matched fixed-factor control on all seven.

At $\epsilon = 1.5$, Fig. 2 shows gains of 2.10, 32.52, and 10.01 percentage points; normalized-radius curves are in Fig. 5. TACBF has higher mean grid AUC than BUTF on all seven datasets; the difference remains significant after Holm correction on five of seven datasets (not CIFAR-10 or UCI HAR). Dataset-level mean clean accuracy differs by at most 0.25 percentage points (Tab. 8); median training-time and maximum memory overheads are 5.9% and 2.4%, and inference is unchanged.

7.2 Closed-form allocation and learned geometry contribute complementary gains

Before factor optimization, D^* increases the grid AUC of the fixed-tight frame by factors of 4.13–4.71 on the three absolute-radius datasets.

The learned factor further improves grid AUC on CIFAR-10, STL-10, and EuroSAT RGB under both mean-codeword-energy matching ($+0.0059, +0.1066, +0.0255$, respectively) and a shared D^*

Table 2: Seven-method grid AUC on held-out test sets at $q = 2$: the unnormalized grid integral over $\epsilon \in [0.02, 1.5]$ or $\delta \in [0, 0.8]$. All methods use five seeds except analytic task allocation on PCam (four seeds) and Banking77 (three seeds), marked by †. The analytic allocation is feasible for four PCam seeds and three Banking77 seeds; the remaining seeds do not satisfy the required feasibility condition $\beta > \mu d/P$. Bold and underline mark the two highest means per dataset.

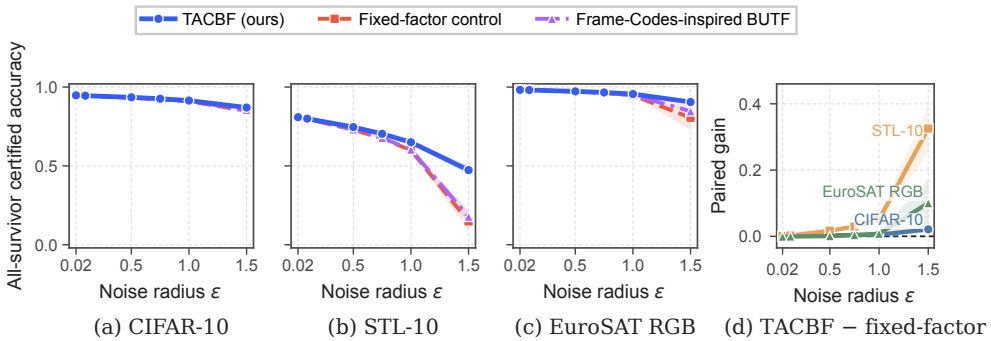
Absolute radius ϵ			
Method	CIFAR-10	STL-10	EuroSAT RGB
TACBF	**1.3606**	**1.0047**	**1.4171**
Fixed-factor control	1.3516	0.8910	1.3882
Frame-Codes-inspired BUTF	<u>1.3550</u>	0.9020	<u>1.3991</u>
Analytic task allocation	1.2422	<u>0.9038</u>	1.2866
Label-free max–min frame	0.5038	0.2811	0.5311
Block dropout	0.8613	0.2185	0.5460
Fixed-tight frame	0.2636	0.2187	0.2739

Control-energy-normalized radius δ				
Method	CIFAR-100	PCam	UCI HAR	Banking77
TACBF	**0.3256**	**0.5511**	**0.4251**	**0.2091**
Fixed-factor control	<u>0.2740</u>	0.4770	0.4104	0.1781
Frame-Codes-inspired BUTF	0.2715	<u>0.4850</u>	<u>0.4196</u>	<u>0.1813</u>
Analytic task allocation	0.2325	0.4381†	0.4109	0.1520†
Label-free max–min frame	0.1594	0.0560	0.1257	0.1668
Block dropout	0.2447	0.4163	0.0710	0.1441
Fixed-tight frame	0.1101	0.0533	0.0709	0.1441

(+0.0729, +0.1385, +0.0762, respectively); all six paired 95% CIs exclude zero (Sec. E.2). The shared-D^* comparison changes only B, holding model outputs and margins fixed. The learned factor lowers mean pairwise uncertainty by 9.0–11.4%; the certified set under the learned factor strictly contains that under the fixed factor in 85 of 90 seed–radius comparisons and is identical in the remaining five.

In a fixed-model 2×2 decomposition, both allocation and factor main effects improve grid AUC for all 35 dataset–seed pairs; nonzero, dataset-dependent interactions show that the effects are not strictly additive (Tab. 14).

For noise uniformly distributed on the unit sphere, the expected squared LS reconstruction error is $\mathrm{tr}(G_C^{-1})/((m-q)r)$. Across datasets, TACBF increases mean expected squared LS reconstruction



(a) CIFAR-10 (b) STL-10 (c) EuroSAT RGB (d) TACBF − fixed-factor

Figure 2: Five-seed absolute-radius certification. Panels (a–c) show all-survivor accuracy and (d) paired TACBF-minus-fixed-factor gains; bands are 95% CIs. BUTF uses the same downstream allocation and training protocol.

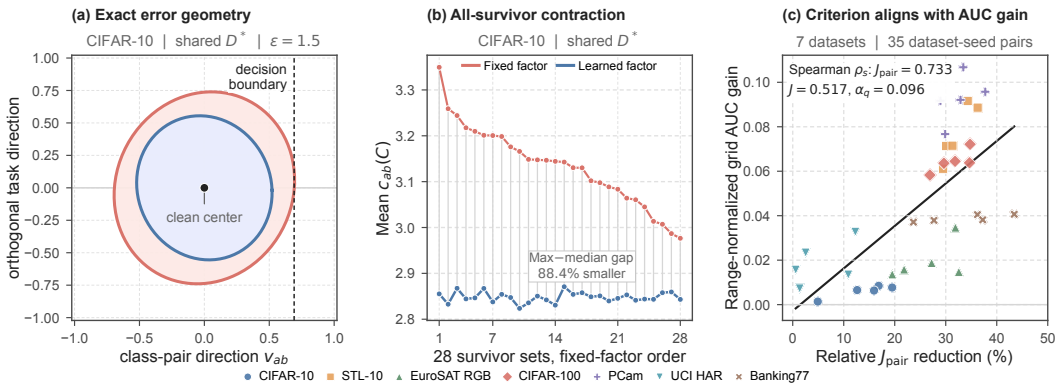


Figure 3: Certificate geometry and criterion alignment. (a) For an illustrative CIFAR-10 validation example at $\epsilon = 1.5$ and shared D^*, the fixed-factor uncertainty ellipse crosses the class-pair decision boundary, whereas the learned-factor ellipse remains inside it. (b) Mean $c_{ab}(C)$ over class pairs for all 28 survivor sets, ordered by the fixed-factor value. (c) Relative J_{pair} reduction versus held-out test-set range-normalized grid AUC gain across 35 dataset–seed pairs at $q = 2$.

error by 3.4%–65.2% while improving range-normalized grid AUC by 0.0061–0.0926 relative to the fixed-factor control (Fig. 6).

7.3 Certificate criteria align with certification gains

Learning B reduces J/J_{relax} on every dataset, from 1.93–9.30 for the control to 1.30–4.64 for TACBF. Every ratio remains strictly above one, consistent with the rank-constrained gap. In an illustrative CIFAR-10 validation example under shared D^*, Fig. 3 shows learned geometry keeps the uncertainty set inside the boundary, lowers uncertainty for all 28 survivor sets, and contracts their maximum-minus-median gap by 88.4%.

For each of the 35 dataset–seed pairs, we compute the relative reductions $1 - J_{\text{pair}}^{\text{TACBF}}/J_{\text{pair}}^{\text{fixed}}$, $1 - J^{\text{TACBF}}/J^{\text{fixed}}$, and $1 - \alpha_q^{\text{TACBF}}/\alpha_q^{\text{fixed}}$. The relative reduction in J_{pair} has a Spearman correlation of $\rho_s = 0.733$ with range-normalized grid AUC gain, exceeding the correlations for the J reduction (0.517) and survivor-floor reduction (0.096).

7.4 Certification gains extend across tasks and erasure levels

The evaluated tasks span $p = 1$–99, $d = 128$–512, and three modalities: images, inertial sensing, and frozen language-model representations (Sec. E.1). For a fixed encoder, head, and allocation, B cancels algebraically at $q = 0$ but enters after erasures through $G_C = D^*(I - \sum_{j \in E} B_j^\top B_j)D^*$, where $E = [m] \setminus C$. The nonzero $q = 0$ difference therefore reflects separately trained models rather than a direct geometric effect of B. Across these models, the mean common-δ AUC difference rises from $+0.0040$ at $q = 0$ to $+0.0723$ at $q = 3$ (Fig. 7).

8 Limitations

The model assumes known erasure locations, bounded total ℓ_2 noise, full-column-rank survivor frames with unregularized LS decoding, and a linear classifier. Unknown masks, alternative noise sets, rank-deficient or regularized decoding, and nonlinear heads require model-specific certificates. The closed-form D^* keeps training scalable. The class-pair SDP gives a tighter reference at greater computational cost. Branch-and-bound provides valid intervals at any computation budget and the exact optimum upon termination, but its worst-case leaf count remains $\binom{m}{q}$. The frozen-model $q = 3$ analysis measures transfer from frames trained with the $q = 2$ survivor-floor constraint.

9 Conclusion

At zero noise, all admissible frames recover the same latent representation; bounded noise makes their geometry govern decision sensitivity. D^* optimally allocates trace under the survivor-floor constraint, while B optimizes the remaining rank-constrained block geometry. Across seven datasets, TACBF improves certified robustness at fixed width and trace while preserving exact recovery and the LS decoder. Exact latent recovery thus leaves exploitable geometric freedom for task-aware robustness.

Reproducibility Statement

Complete proofs, algorithms, and experimental details are provided in the appendix; code and scripts for reproducing the reported results and figures will be released upon acceptance.

AI Use Statement

Large language models were used for code assistance and language polishing. The authors developed the research questions and mathematical derivations and designed and conducted the experiments. The authors reviewed all model-assisted code and text and take full responsibility for the manuscript.

References

Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge L. Reyes-Ortiz. A public domain dataset for human activity recognition using smartphones. In *Proceedings of the 21st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, pp. 437–442, 2013.

Bernhard G. Bodmann. Optimal linear transmission by loss-insensitive packet encoding. *Applied and Computational Harmonic Analysis*, 22(3):274–285, 2007. doi: 10.1016/j.acha.2006.07.003.

Bernhard G. Bodmann and Vern I. Paulsen. Frames, graphs and erasures. *Linear Algebra and its Applications*, 404:118–146, 2005. doi: 10.1016/j.laa.2005.02.016.

Chang Cai, Xiaojun Yuan, and Ying-Jun Angela Zhang. Multi-device task-oriented communication via maximal coding rate reduction. *IEEE Transactions on Wireless Communications*, 23(12): 18096–18110, 2024. doi: 10.1109/TWC.2024.3461336.

Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. Efficient intent detection with dual sentence encoders. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pp. 38–45. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.nlp4convai-1.5.

Peter G. Casazza and Jelena Kovačević. Equal-norm tight frames with erasures. *Advances in Computational Mathematics*, 18(2–4):387–430, 2003. doi: 10.1023/A:1021349819855.

Peter G. Casazza and Gitta Kutyniok. Robustness of fusion frames under erasures of subspaces and of local frame vectors. In *Radon Transforms, Geometry, and Wavelets*, volume 464 of *Contemporary Mathematics*, pp. 149–160. American Mathematical Society, 2008. doi: 10.1090/conm/464/09082.

Peter G. Casazza, Gitta Kutyniok, and Shidong Li. Fusion frames and distributed processing. *Applied and Computational Harmonic Analysis*, 25(1):114–132, 2008. doi: 10.1016/j.acha.2007.10.001.

Moustapha Cisse, Piotr Bojanowski, Edouard Grave, Yann Dauphin, and Nicolas Usunier. Parseval networks: Improving robustness to adversarial examples. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 854–863, 2017.

Adam Coates, Andrew Y. Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pp. 215–223, 2011.

Jeremy M. Cohen, Elan Rosenfeld, and J. Zico Kolter. Certified adversarial robustness via randomized smoothing. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1310–1320, 2019.

I. S. Dhillon, R. W. Heath, Thomas Strohmer, and Joel A. Tropp. Constructing packings in Grassmannian manifolds via alternating projection. *Experimental Mathematics*, 17(1):9–35, 2008. doi: 10.1080/10586458.2008.10129018.

Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936. doi: 10.1007/BF02288367.

Alan Edelman, Tomás A. Arias, and Steven T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM Journal on Matrix Analysis and Applications*, 20(2):303–353, 1998. doi: 10.1137/S0895479895290954.

G. Elfving. Optimum allocation in linear regression theory. *The Annals of Mathematical Statistics*, 23(2):255–262, 1952. doi: 10.1214/aoms/1177729442.

Homa Esfahanizadeh, Nargis Fayaz, Jinfeng Du, and Harish Viswanathan. Block erasure-aware semantic multimedia compression via JSCC autoencoder. *arXiv preprint arXiv:2601.20707*, 2026.

Nargis Fayaz, Homa Esfahanizadeh, Matin Mortaheb, Jinfeng Du, and Harish Viswanathan. Towards robust semantic video transmission over block erasure channels. *arXiv preprint arXiv:2607.07823*, 2026.

Vivek K. Goyal, Jelena Kovačević, and Jonathan A. Kelner. Quantized frame expansions with erasures. *Applied and Computational Harmonic Analysis*, 10(3):203–233, 2001. doi: 10.1006/acha.2000.0340.

Deniz Gündüz, Zhijin Qin, Iñaki Estella Aguerri, Harpreet S. Dhillon, Zhaohui Yang, Aylin Yener, Kai Kit Wong, and Chan-Byoung Chae. Beyond transmitting bits: Context, semantics, and task-oriented communications. *IEEE Journal on Selected Areas in Communications*, 41(1):5–41, 2023. doi: 10.1109/JSAC.2022.3223408.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016. doi: 10.1109/CVPR.2016.90.

Matthias Hein and Maksym Andriushchenko. Formal guarantees on the robustness of a classifier against adversarial manipulation. In *Advances in Neural Information Processing Systems 30*, pp. 2266–2276, 2017.

Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. EuroSAT: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019. doi: 10.1109/JSTARS.2019.2918242.

Nicholas J. Higham. Computing the polar decomposition—with applications. *SIAM Journal on Scientific and Statistical Computing*, 7(4):1160–1174, 1986. doi: 10.1137/0907079.

Roderick B. Holmes and Vern I. Paulsen. Optimal frames for erasures. *Linear Algebra and its Applications*, 377:31–51, 2004. doi: 10.1016/j.laa.2003.07.012.

Itamar Jacoby and Ram Zamir. Frame Codes for the block-erasure channel. In *2024 IEEE International Symposium on Information Theory (ISIT)*, pp. 1125–1130, 2024. doi: 10.1109/ISIT57864.2024.10619295.

Siddharth Joshi and Stephen Boyd. Sensor selection via convex optimization. *IEEE Transactions on Signal Processing*, 57(2):451–462, 2009. doi: 10.1109/TSP.2008.2007095.

J. Kiefer. General equivalence theory for optimum designs (approximate theory). *The Annals of Statistics*, 2(5):849–879, 1974. doi: 10.1214/aos/1176342810.

Lukas Köhldorfer, Peter Balazs, Pete Casazza, Sigrid Heineken, Clara Hollomey, Patricia Morillas, and Mitra Shamsabadi. A survey of fusion frames in Hilbert spaces. In *Sampling, Approximation, and Signal Analysis*, Applied and Numerical Harmonic Analysis, pp. 245–328. Springer International Publishing, 2023. doi: 10.1007/978-3-031-41130-4_11.

Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

Nikolina Kubiak and Simon Hadfield. TACTIC: Joint rate-distortion-accuracy optimisation for low bitrate compression. *arXiv preprint arXiv:2109.10658*, 2021.

Emrecan Kutay and Aylin Yener. Classification-oriented semantic wireless communications. In *2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 9096–9100, 2024. doi: 10.1109/ICASSP48485.2024.10447549.

Gitta Kutyniok, Ali Pezeshki, Robert Calderbank, and Taotao Liu. Robust dimension reduction, fusion frames, and Grassmannian packings. *Applied and Computational Harmonic Analysis*, 26 (1):64–76, 2009. doi: 10.1016/j.acha.2008.03.001.

David R. Larson and Sam L. Scholze. Signal reconstruction from frame and sampling erasures. *Journal of Fourier Analysis and Applications*, 21(5):1146–1167, 2015. doi: 10.1007/s00041-015-9404-3.

Yi Ma, Chunmei Xu, Zhenyu Liu, Siqi Zhang, and Rahim Tafazolli. Importance-aware source-channel coding for multi-modal task-oriented semantic communication. In *2025 IEEE International Conference on Machine Learning for Communication and Networking (ICMLCN)*, pp. 1–6, 2025. doi: 10.1109/ICMLCN64995.2025.11140558.

Pedro G. Massey. Optimal reconstruction systems for erasures and for the q-potential. *Linear Algebra and its Applications*, 431(8):1302–1316, 2009. doi: 10.1016/j.laa.2009.05.001.

Tom Needham and Clayton Shonkwiler. Fusion frame homotopy and tightening fusion frames by gradient descent. *Journal of Fourier Analysis and Applications*, 29(4):51, 2023. doi: 10.1007/s00041-023-10028-0.

Francesco Pezone, Sergio Barbarossa, and Paolo Di Lorenzo. Goal-oriented communication for edge learning based on the information bottleneck. In *2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8832–8836, 2022. doi: 10.1109/ICASSP43922.2022.9746468.

Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pp. 3982–3992, 2019. doi: 10.18653/v1/D19-1410.

Jiawei Shao, Yuyi Mao, and Jun Zhang. Learning task-oriented communication for edge inference: An information bottleneck approach. *IEEE Journal on Selected Areas in Communications*, 40(1): 197–211, 2022. doi: 10.1109/JSAC.2021.3126087.

Nir Shlezinger and Yonina C. Eldar. Task-based quantization with application to MIMO receivers. *Communications in Information and Systems*, 20(2):131–162, 2020. doi: 10.4310/CIS.2020.v20.n2.a3.

Moirangthem Tiken Singh and Adnan Arif. Reinforcement-learned unequal error protection for quantized semantic embeddings. *Array*, 29:100727, 2026. doi: 10.1016/j.array.2026.100727.

Thomas Strohmer and Robert W. Heath. Grassmannian frames with applications to coding and communication. *Applied and Computational Harmonic Analysis*, 14(3):257–275, 2003. doi: 10.1016/S1063-5203(03)00023-X.

Yusuke Tsuzuku, Issei Sato, and Masashi Sugiyama. Lipschitz-margin training: Scalable certification of perturbation invariance for deep neural networks. In *Advances in Neural Information Processing Systems 31*, pp. 6542–6551, 2018.

Bastiaan S. Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation equivariant CNNs for digital pathology. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, pp. 210–218, 2018. doi: 10.1007/978-3-030-00934-2_24.

Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. MiniLM: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. In *Advances in Neural Information Processing Systems 33*, pp. 5776–5788, 2020.

Task-Aware Certified Block Frames for Exact Latent Recovery

Supplementary Material

CONTENTS

A Notation **14**

B Proofs and Supporting Results **14**

B.1 Proof of the exact certificate and robust radius 14

B.2 Reusable pair coefficients and certified survivor search 15

B.3 Dimensional conditions and generic feasibility 19

B.4 Proofs of the class-pair coverage bound and relaxed optimum 19

B.5 Proof of the anisotropic lifting 21

B.6 Rank-constrained objective gap and Gram separation 24

C Experimental Details **24**

C.1 Architectures and optimization 24

C.2 Preprocessing and data splits 25

C.3 Certification metrics and statistical analysis 25

C.4 Factorized objective and training algorithm 26

C.5 Experimental controls and computational cost 27

D Exact Survivor Search and Scalability **27**

D.1 Search bounds and convergence 27

E Additional Results **28**

E.1 Results across datasets and baselines 29

E.2 Controlled comparisons of allocation and factor geometry 31

E.3 Factor geometry and certification gains 33

F Erasure-Severity Transfer and Relaxation Gaps **33**

F.1 Erasure-severity transfer 33

F.2 Relaxation gap and class-pair coverage 33

A NOTATION

Tab. 3 collects the symbols used repeatedly across the certificate, allocation, factorization, and rank-gap analyses. Dataset-specific quantities and auxiliary proof variables are defined locally where they first appear.

Table 3: Core notation used throughout the paper and supplementary material.

Symbol	Description
$[L]$, K, d	Index set $\{1, \dots, L\}$, number of classes, and latent dimension.
x, y, $s = f_\theta(x)$	Input, class label, and unit-normalized latent representation.
m, r, q, $n = m - q$	Number of blocks, rows per block, erased blocks, and surviving blocks.
A_j, A, A_C	Block analysis operator, stacked analysis operator, and survivor operator.
C, $E = [m] \setminus C$, $G_C = A_C^\top A_C$	Survivor index set, erased-block set, and survivor Gram.
P, μ, $\alpha_q(A)$	Trace budget, required survivor-floor threshold, and frame survivor floor.
η_C, ϵ, $\widehat{s}_C$, $\widehat{s}^0_{i,C}$	Observation noise, its total ℓ_2 bound, noisy LS estimate, and nominal LS solution.
W, w_k, $v_{ab} = w_a - w_b$	Linear head, class-k row, and class-pair decision direction.
M^0_{ik}, $M_{ik}(C)$	Clean pairwise margin and reconstructed pairwise margin.
$R_{ab}(C;\epsilon)$, $c_{ab}(C)$, $\overline{c}_{ab,q}$	Score uncertainty, per-survivor uncertainty coefficient, and worst-survivor uncertainty coefficient with ϵ factored out.
ϵ^*_i, $\mathrm{CertAcc}_{\mathrm{all}}$	Exact all-survivor radius and all-survivor certified-accuracy curve.
δ, $\varrho^{(\epsilon)}_i$, $\varrho^{(\delta)}_i$	Control-energy-normalized radius and absolute/normalized sample radii.
$\mathcal{U}$, p, U, $U_\perp$, $\Pi = UU^\top$	Task subspace, its rank, bases for the task subspace and its orthogonal complement, and its projector.
$\mathcal{D}$, $\widehat{v}_\ell$, z_ℓ, κ_W	Normalized class-pair directions, their task coordinates, and class-pair coverage.
$\underline{\kappa}_W$, $\underline{\kappa}^{\mathrm{opt}}_W$	Uniform lower bound and its exact optimized counterpart.
$J(A)$, $J_{\mathrm{pair}}(A)$	Spectral task envelope and class-pair envelope.
$J_{\mathrm{relax}} = p/\mathcal{B}$	Rank-relaxed lower bound for the spectral task envelope.
$\mathcal{R}_\beta(D)$, $\gamma^{\mathrm{rel}}_{\mathrm{pair}}$	Scalar-floor robust spectral envelope and class-pair relaxation optimum.
$H_j = A_j^\top A_j$, G_*, H_*	Block Gram, relaxed target survivor Gram, and homogeneous relaxed block Gram.
$\mathcal{B} = nP/m - (d-p)\mu$	Trace available to the task subspace in the rank-relaxed problem.
B_j, B_C, $F_C = B_C^\top B_C$, $\beta = \beta_q(B)$	Blocks of the Parseval factor, stacked surviving factor blocks and their Gram matrix, and the Parseval-factor survivor floor.
D, D^*, D_{iso}, ρ, τ	General SPD allocation, closed-form allocation, isotropic allocation, and the eigenvalues of $(D^*)^2$ on the task subspace and its orthogonal complement.
Γ_{alloc}	Allocation gain factor: squared-radius ratio and inverse squared-uncertainty ratio of D^* relative to isotropic scaling.
$\mathcal{F}_r$	Rank-constrained block-Gram family satisfying the trace and survivor-floor constraints.

B PROOFS AND SUPPORTING RESULTS

B.1 PROOF OF THE EXACT CERTIFICATE AND ROBUST RADIUS

Proof of Theorem 1. Write $G_C = A_C^\top A_C$. Then

$$\widehat{s}_C - s = G_C^{-1} A_C^\top \eta_C, \qquad v_{ab}^\top(\widehat{s}_C - s) = \left(A_C G_C^{-1} v_{ab}\right)^\top \eta_C. \tag{23}$$

The coefficient vector has squared norm

$$\left\| A_C G_C^{-1} v_{ab} \right\|_2^2 = v_{ab}^\top G_C^{-1} v_{ab}. \tag{24}$$

The Cauchy–Schwarz inequality bounds the absolute score change by $\epsilon \| A_C G_C^{-1} v_{ab} \|_2$. If $v_{ab} = 0$, both sides of Eq. (3) are zero. Otherwise, the perturbations $\eta_C = \pm \epsilon A_C G_C^{-1} v_{ab} / \| A_C G_C^{-1} v_{ab} \|_2$ attain this bound, proving Theorem 1. □

Corollary 6 (Exact all-survivor robust radius) *Fix nominal survivor observations $z^0_{i,C}$ and LS centers $\widehat{s}^0_{i,C} = G_C^{-1} A_C^\top z^0_{i,C}$. If all corresponding true-class margins are positive and $v_{y_i k} \neq 0$ for every $k \neq y_i$, the largest open-ball radius preserving label y_i for every survivor set is*

$$\epsilon_i^* = \min_{|C|=m-q} \min_{k \neq y_i} \frac{(w_{y_i} - w_k)^\top \widehat{s}^0_{i,C}}{c_{y_i k}(C)}. \tag{25}$$

At equality an explicit perturbation reaches a decision tie.

Proof. An additional perturbation ξ_C changes the y_i-versus-k nominal margin by $\left(A_C G_C^{-1} v_{y_i k}\right)^\top \xi_C$. The coefficient norm is $c_{y_i k}(C) > 0$. Therefore, over the closed perturbation ball of radius ϵ, the smallest possible pairwise margin is exactly

$$v_{y_i k}^\top \widehat{s}^0_{i,C} - \epsilon c_{y_i k}(C). \tag{26}$$

It remains strictly positive if and only if $\epsilon < (v_{y_i k}^\top \widehat{s}^0_{i,C})/c_{y_i k}(C)$. At equality, the perturbation

$$\xi_C = -\epsilon \frac{A_C G_C^{-1} v_{y_i k}}{c_{y_i k}(C)} \tag{27}$$

has norm ϵ and reaches the corresponding decision hyperplane. Minimizing over competitors and survivor sets proves both necessity and sufficiency. □

Support functions use closed ℓ_2 balls, while certification requires strict preservation of the true-class decision, so the attainable boundary tie is not counted as certified. For N evaluation examples, the all-survivor certified accuracy is

$$\mathrm{CertAcc}_{\mathrm{all}}(\epsilon) = \frac{1}{N} \sum_{i=1}^{N} \min_{|C|=m-q} \mathbf{1}\left\{ \min_{k \neq y_i} \left[v_{y_i k}^\top \widehat{s}^0_{i,C} - \epsilon c_{y_i k}(C) \right] > 0 \right\}. \tag{28}$$

Lemma 7 (Worst-case erasure level) *Suppose every survivor Gram with $m - q$ blocks is positive definite. For any $v \in \mathbb{R}^d$, the largest value of $v^\top G_C^{-1} v$ over all patterns with at most q erasures is attained by a pattern with exactly q erasures.*

Proof. Consider any survivor set C with $|C| \geq m - q$, and choose $C_0 \subseteq C$ with $|C_0| = m - q$. Then

$$A_C^\top A_C = A_{C_0}^\top A_{C_0} + \sum_{j \in C \setminus C_0} A_j^\top A_j \succeq A_{C_0}^\top A_{C_0}. \tag{29}$$

The assumed positive definiteness and Loewner-order reversal under inversion give $G_C^{-1} \preceq G_{C_0}^{-1}$. Hence $v^\top G_C^{-1} v \leq v^\top G_{C_0}^{-1} v$. Taking the maximum proves the claim. □

B.2 Reusable pair coefficients and certified survivor search

This subsection establishes coefficient reuse and erasure-level monotonicity before deriving certified bounds for exact worst-survivor search.

Proof of the reusable-pair reduction. At the nominal LS solution, $v_{y_i k}^\top \widehat{s}^0_{i,C} = M^0_{ik}$ is independent of C. Because every clean-correct pairwise margin is positive and the survivor family is finite,

$$\min_{|C|=m-q} \frac{M^0_{ik}}{c_{y_i k}(C)} = \frac{M^0_{ik}}{\max_{|C|=m-q} c_{y_i k}(C)}. \tag{30}$$

Taking the minimum over competitors k yields Eq. (5). □

For each erasure count q whose survivor Grams are all positive definite, let $\epsilon^*_{i,q}$ denote Eq. (5) with $\overline{c}_{ab,q}$, and define $\mathrm{CertAcc}^{(q)}_{\mathrm{all}}(\epsilon) := N^{-1} \sum_i \mathbf{1}\{\widehat{y}_i = y_i,\ \epsilon < \epsilon^*_{i,q}\}$. Let $q_{\max}$ be the largest such erasure count.

Corollary 8 (Monotonicity in erasure level) *If all relevant survivor Grams are positive definite, then for $0 \leq q < q_{\max}$,*

$$\overline{c}_{ab,q+1} \geq \overline{c}_{ab,q}, \qquad \epsilon^*_{i,q+1} \leq \epsilon^*_{i,q}, \qquad \mathrm{CertAcc}^{(q+1)}_{\mathrm{all}}(\epsilon) \leq \mathrm{CertAcc}^{(q)}_{\mathrm{all}}(\epsilon). \tag{31}$$

Proof. Fix an $(m-q)$-block survivor set C attaining $\overline{c}_{ab,q}$ and remove any $j \in C$ to obtain $C' = C \setminus \{j\}$. Then $G_{C'} = G_C - A_j^\top A_j \preceq G_C$. Positive definiteness and reversal of the Loewner order under inversion give

$$v_{ab}^\top G_{C'}^{-1} v_{ab} \geq v_{ab}^\top G_C^{-1} v_{ab}. \tag{32}$$

Maximizing over all $(m-q-1)$-block survivor sets proves $\overline{c}_{ab,q+1} \geq \overline{c}_{ab,q}$. At the nominal LS solution, the clean pairwise margins are independent of the survivor set and erasure level. Substitution into Eq. (5) proves the radius inequality, and thresholding these radii proves the certified-accuracy inequality in Eq. (31). □

Exact worst-survivor search with anytime bounds. The reusable-pair reduction requires one combinatorial search per unordered class pair. Searching survivor sets C is equivalent to searching erased sets $E = [m] \setminus C$. Assume $G_C \succeq \mu I$ for every admissible C, where $\mu > 0$ is the required survivor-floor threshold. For a partial erased set E, define $G(E) := \sum_{j \notin E} A_j^\top A_j$, which equals G_C at a complete erasure set, and let $h = q - |E|$. Among the remaining blocks, let $s_A(E)$ be the sum of the h largest values of $\|A_j\|_2^2$. Every completion $F \subseteq [m] \setminus E$ with $|F| = h$ satisfies

$$G(E \cup F) \succeq G(E) - s_A(E) I. \tag{33}$$

If the right-hand side is positive definite, its inverse gives a node-wise upper bound on $v^\top G(E \cup F)^{-1} v$; otherwise the frame floor gives the bound $\|v\|_2^2/\mu$. Completed patterns provide the incumbent lower bound, and the largest open-node bound provides the global upper bound. Thus, if $L_{ab} \leq \overline{c}_{ab,q}^2 \leq U_{ab}$, then

$$\min_{k \neq y_i} \frac{M_{ik}^0}{\sqrt{U_{y_i k}}} \leq \epsilon_i^* \leq \min_{k \neq y_i} \frac{M_{ik}^0}{\sqrt{L_{y_i k}}}. \tag{34}$$

The left endpoint is a certified lower bound on ϵ_i^* at every stopping time. At a target radius, evaluation can stop for a sample as soon as this lower endpoint exceeds the radius (certified), or as soon as the upper endpoint does not exceed it (not certified). Exhausting the queue recovers the global optimum. Greedy and beam completions can accelerate incumbent discovery, while the retained node bounds provide the global guarantee.

A complementary whitened bound uses the direction being evaluated. Write $H_j = A_j^\top A_j$ and $G = G(E)$. For each remaining block define

$$\theta_j = \lambda_{\max}(G^{-1/2} H_j G^{-1/2}), \qquad r_j(v) = v^\top G^{-1} H_j G^{-1} v = \|A_j G^{-1} v\|_2^2. \tag{35}$$

Let $\Theta_h(E)$ and $R_h(E, v)$ be the sums of the h largest remaining θ_j and $r_j(v)$ values, respectively. If $\Theta_h(E) < 1$, every completion F satisfies

$$v^\top G(E \cup F)^{-1} v \leq v^\top G^{-1} v + \frac{R_h(E, v)}{1 - \Theta_h(E)}. \tag{36}$$

Taking the minimum of this expression, the scalar spectral-load bound, and $\|v\|_2^2/\mu$ preserves certification while tightening any node for which the new bound is smaller.

Algorithm 1 Anytime branch-and-bound for one class-pair direction. A node (E, a) represents exactly the q-subsets whose increasing erased-index prefix is E and whose remaining indices are greater than a.

Input: block Grams $H_j = A_j^\top A_j$, direction v, number of erased blocks q, and required survivor-floor threshold μ.
Obtain a complete q-set by greedy/beam completion; set the incumbent L to its exact value $v^\top G_C^{-1} v$.
Create the root $(E, a) = (\emptyset, 0)$; compute its certified upper bound and insert it into a max-priority queue $\mathcal{Q}$.
while $\mathcal{Q} \neq \emptyset$ and $\max_{N \in \mathcal{Q}} \mathrm{UB}(N) > L$ **do**
  Pop the node (E, a) with largest upper bound and set $k \leftarrow |E|$.
  for $j = a + 1, \ldots, m - q + k + 1$ **do**
    Set $E' \leftarrow E \cup \{j\}$.
    if $|E'| = q$ **then** evaluate its exact objective and update L.
    else compute the minimum of the certified node bounds for completions using indices greater than j; insert (E', j) only if its bound exceeds L.
    end if
  end for
end while
Output at any stopping point: $[L, U]$, where $U = L$ if $\mathcal{Q} = \emptyset$ and otherwise $U = \max\{L, \max_{N \in \mathcal{Q}} \mathrm{UB}(N)\}$; at termination $L = U = \bar{c}^2_{ab,q}$.

Let $\mathcal{C}_{\mathrm{eval}}$ denote the cost of forming and solving one survivor-Gram quadratic for a fixed direction. Tab. 4 separates the combinatorial worst case from the reductions used in practice.

Table 4: Worst-case and practical complexity of survivor search. A node budget b counts evaluated or bounded search nodes for one class-pair direction.

Procedure	Work per direction	Guarantee
Exhaustive enumeration	$\binom{m}{q}\mathcal{C}_{\mathrm{eval}}$	Global optimum
Whitened greedy ordering	$O(m \log m)$ after block scores	$(1 - \Theta_q)$ incumbent guarantee when $\Theta_q < 1$
Budgeted beam search	$O(b\mathcal{C}_{\mathrm{eval}})$	Approximate unless completed
Branch-and-bound	$O(b\mathcal{C}_{\mathrm{eval}})$; at most $\binom{m}{q}$ leaves	Certified interval at every budget; exact at termination
Class-pair reuse	At most $K(K-1)/2$ searches instead of one search per sample–competitor pair	Exact coefficient sharing

Proposition 9 (Greedy incumbent guarantee) *Let $G_0 = \sum_j H_j$, $X_j = G_0^{-1/2} H_j G_0^{-1/2}$, and $u = G_0^{-1/2} v$. Define $r_j = u^\top X_j u$ and let S_{g} contain the q largest r_j. Let Θ_q be the sum of the q largest $\lambda_{\max}(X_j)$. For*

$$G(S) := G_0 - \sum_{j \in S} H_j, \qquad v^\top G(S)^{-1} v = u^\top u + \Delta(S), \tag{37}$$

where

$$\Delta(S) := u^\top \big[(I - \textstyle\sum_{j \in S} X_j)^{-1} - I \big] u, \tag{38}$$

if $\Theta_q < 1$, then

$$\Delta(S_{\mathrm{g}}) \geq (1 - \Theta_q) \max_{|S| = q} \Delta(S). \tag{39}$$

Thus maximizing $\Delta(S)$ is equivalent to maximizing the original erased-set objective $v^\top G(S)^{-1} v$.

Proof. For every q-set S, write $X_S = \sum_{j \in S} X_j$. Since $0 \preceq X_S \preceq \Theta_q I$, functional calculus gives

$$u^\top X_S u \leq \Delta(S) \leq \frac{u^\top X_S u}{1 - \Theta_q}. \tag{40}$$

The definition of S_{g} gives $u^\top X_{S_{\mathrm{g}}} u \geq u^\top X_S u$ for every q-set S. Combining the lower bound for S_{g} with the upper bound for an optimal set proves Eq. (39). □

Proof of the anytime search bounds. **Step 1: scalar node bound.** For a partial erased set E and any completion F of size $h = q - |E|$,

$$\sum_{j\in F} A_j^\top A_j \preceq \sum_{j\in F} \|A_j\|_2^2 I \preceq s_A(E) I. \tag{41}$$

Subtracting from $G(E)$ proves Eq. (33). If its right-hand side is positive definite, Loewner-order reversal gives

$$G(E\cup F)^{-1} \preceq \left(G(E) - s_A(E) I\right)^{-1}. \tag{42}$$

If it is not positive definite, the assumed frame floor gives $G(E\cup F)^{-1} \preceq \mu^{-1} I$. These two cases upper-bound every completion represented by the node. Moreover, any partial deletion set satisfies $G(E) \succeq G(E\cup F) \succ 0$.

Step 2: direction-aware node bound. Set $X = G^{-1/2}(\sum_{j\in F} H_j)G^{-1/2}$. By construction, $0 \preceq X \preceq \Theta_h(E) I$. When $\Theta_h(E) < 1$, scalar functional calculus gives

$$(I - X)^{-1} \preceq I + \frac{X}{1 - \Theta_h(E)}. \tag{43}$$

Because

$$G(E\cup F) = G^{1/2}(I - X)G^{1/2}, \tag{44}$$

Sandwiching by $G^{-1/2}v$ and using $\sum_{j\in F} r_j(v) \le R_h(E, v)$ proves Eq. (36).

Step 3: updates and global certification. The frame floor makes every complete survivor Gram positive definite, and each partial Gram dominates its completions; hence the following downdate is well defined. The matrix inversion lemma applied to $G(E\cup\{j\}) = G(E) - A_j^\top A_j$ gives

$$G(E\cup\{j\})^{-1} = G(E)^{-1} + G(E)^{-1}A_j^\top \left(I - A_j G(E)^{-1} A_j^\top\right)^{-1} A_j G(E)^{-1}. \tag{45}$$

Thus each child requires an $r \times r$ inverse rather than a fresh $d \times d$ inverse.

Every completed erasure set provides a feasible lower bound. The largest bound among the open nodes upper-bounds all unevaluated completions. Because branching partitions the finite search space, exhausting the queue closes the gap and returns the global maximum; pruning a node whose upper bound does not exceed the incumbent cannot remove an improving completion. Every q-set has a unique increasing index sequence and therefore a unique root-to-leaf path in Alg. 1. Finally, $x \mapsto M/\sqrt{x}$ is decreasing for $M > 0$. Applying it to $L_{ab} \le \bar{c}_{ab,q}^2 \le U_{ab}$ and then minimizing over competitors proves Eq. (34). □

Parseval floor identities and a computable bound. The Parseval condition yields exact floor identities, a computable lower bound, and a conservative lifting based on that bound. For $E = [m]\setminus C$, the Parseval condition gives the exact factor-floor identity

$$\beta_q(B) = 1 - \max_{|E|=q} \lambda_{\max}\left(\sum_{j\in E} B_j^\top B_j\right). \tag{46}$$

Consequently, $\beta_1 = 1 - \max_j \|B_j\|_2^2$, $\beta_{q+1} \le \beta_q$, and $\beta_q \le (m-q)/m$, strictly when every block has rank at most $r < d$ and $1 \le q \le m-1$.

Proof of the Parseval floor identities and bound. For the Parseval factor,

$$B_C^\top B_C = I - \sum_{j\in E} B_j^\top B_j. \tag{47}$$

Taking the smallest eigenvalue and maximizing over q-element erased sets proves Eq. (46); the $q = 1$ identity and monotonicity follow immediately. Uniformly averaging the erased-block sum over all q-subsets gives $(q/m)I$, and hence

$$\max_{|E|=q} \lambda_{\max}\left(\sum_{j\in E} B_j^\top B_j\right) \ge q/m. \tag{48}$$

This proves $\beta_q \le (m-q)/m$. Equality would force every q-subset sum to equal $(q/m)I$: under equality, $(q/m)I - \sum_{j\in E} B_j^\top B_j \succeq 0$ for every E, and these positive-semidefinite differences average to zero, so every difference vanishes. When $1 \le q \le m-1$, compare two q-subsets $E = S \cup \{i\}$ and $E' = S \cup \{j\}$ that differ in exactly one block. Equality of their sums forces $B_i^\top B_i = B_j^\top B_j$, and hence $B_j^\top B_j = I/m$ for every j, which is impossible when $\operatorname{rank}(B_j) \le r < d$. Subadditivity of $\lambda_{\max}$ for positive-semidefinite matrices gives

$$\lambda_{\max}\left(\sum_{j\in E} B_j^\top B_j\right) \le \sum_{j\in E} \|B_j\|_2^2 \le \sum_{\ell=1}^{q} h_{(\ell)}, \tag{49}$$

where $h_{(1)} \ge \cdots \ge h_{(m)}$ sorts $h_j = \|B_j\|_2^2$. Therefore

$$\beta_q(B) \ge \underline{\beta}_q := \left[1 - \sum_{\ell=1}^{q} h_{(\ell)}\right]_+ . \tag{50}$$

If $\underline{\beta}_q > \mu d/P$, substituting $\underline{\beta}_q$ for β in both τ and ρ yields a conservative feasible lifting, because the true floor satisfies $\beta_q \tau \ge \underline{\beta}_q \tau = \mu$. □

B.3 Dimensional conditions and generic feasibility

Lemma 10 (Dimension and generic feasibility) *Exact recovery after every q-block erasure requires $(m-q)r \ge d$. Conversely, if this inequality holds and the stacked operator has a jointly absolutely continuous distribution, then every admissible survivor operator has full column rank almost surely. Right whitening produces a Parseval operator without changing any survivor rank.*

Proof. Necessity follows from $\operatorname{rank}(A_C) \le (m-q)r$. For sufficiency, fix an admissible C. Because $(m-q)r \ge d$, at least one $d \times d$ minor of A_C is a nonzero polynomial in its entries: assigning an identity matrix to that minor shows that the polynomial is not identically zero. Its zero set has Lebesgue measure zero, so A_C has full column rank almost surely. The survivor family is finite; the union of its measure-zero exceptional sets still has measure zero.

The stacked A is then full column rank. Let $\widetilde{A} = A(A^\top A)^{-1/2}$. We have $\widetilde{A}^\top \widetilde{A} = I$, and each $\widetilde{A}_C = A_C(A^\top A)^{-1/2}$ has the same rank as A_C because the right factor is invertible. □

B.4 Proofs of the class-pair coverage bound and relaxed optimum

Proof of Proposition 2. Fix C and write $M_C = U^\top G_C^{-1} U \succeq 0$. Each $\widehat{v}_\ell$ lies in $\mathcal{U}$, so $z_\ell = U^\top \widehat{v}_\ell$ is a unit vector. Hence $z_\ell^\top M_C z_\ell \le \lambda_{\max}(M_C)$, which gives the upper bound in Eq. (7). If u_C is a unit top eigenvector of M_C, then

$$\max_{\ell} (u_C^\top z_\ell)^2 \ge \kappa_W. \tag{51}$$

Choose ℓ_C to maximize $|u_C^\top z_\ell|$. The spectral decomposition of M_C gives $z_{\ell_C}^\top M_C z_{\ell_C} \ge \lambda_{\max}(M_C)(u_C^\top z_{\ell_C})^2$. Maximizing over C gives the lower bound in Eq. (7). If A_J globally minimizes J over a feasible family and A_{pair} minimizes J_{pair}, then $J_{\mathrm{pair}}(A_J) \le J(A_J) \le J(A_{\mathrm{pair}}) \le J_{\mathrm{pair}}(A_{\mathrm{pair}})/\kappa_W$. □

Corollary 11 (Computable class-pair coverage bounds) *The class-pair coverage coefficient satisfies*

$$\kappa_W \ge \underline{\kappa}_W^{\mathrm{opt}} := \max_{\pi \in \Delta_{|\mathcal{D}|}} \lambda_{\min}\left(\sum_{\ell} \pi_\ell z_\ell z_\ell^\top\right) \ge \underline{\kappa}_W > 0,$$
$$\underline{\kappa}_W := \lambda_{\min}\left(\frac{1}{|\mathcal{D}|}\sum_{\ell} z_\ell z_\ell^\top\right). \tag{52}$$

The optimized lower bound is computed by an SDP and dominates uniform weighting.

Proof. For any $\pi \in \Delta_{|\mathcal{D}|}$ and unit u, the maximum over ℓ dominates the π-weighted average. For fixed π, minimizing $u^\top(\sum_\ell \pi_\ell z_\ell z_\ell^\top)u$ over unit u gives the smallest eigenvalue. Maximizing this value over π yields $\underline{\kappa}_W^{\mathrm{opt}}$. Uniform weighting is feasible and yields $\underline{\kappa}_W$, proving the full chain in Eq. (52). □

Proof of Theorem 3. **Lower bound.** Complete U to an orthogonal basis $[U, U_\perp]$ and let $T_C = U^\top G_C U$. Uniform averaging over all n-subsets gives

$$\mathbb{E}_C G_C = \frac{n}{m}\sum_j H_j, \qquad \mathbb{E}_C \operatorname{tr}(G_C) = \frac{nP}{m}. \tag{53}$$

Because $G_C \succeq \mu I$, $\operatorname{tr}(U_\perp^\top G_C U_\perp) \geq (d-p)\mu$. Hence $\mathbb{E}_C \operatorname{tr}(T_C) \leq \mathcal{B}$, and some C_* satisfies $\operatorname{tr}(T_{C_*}) \leq \mathcal{B}$. In the basis $[U, U_\perp]$, the upper-left block of $G_{C_*}^{-1}$ is the inverse of a Schur complement no larger than T_{C_*}. Thus the compression-of-inverse inequality yields

$$U^\top G_{C_*}^{-1} U \succeq T_{C_*}^{-1}, \qquad \lambda_{\max}(T_{C_*}^{-1}) \geq p/\operatorname{tr}(T_{C_*}) \geq p/\mathcal{B}. \tag{54}$$

Achievability. Let $G_* = (\mathcal{B}/p)UU^\top + \mu(I - UU^\top)$ and set $H_j = G_*/n$. Every n-block survivor Gram then equals G_*, while $\operatorname{tr}(G_*) = \mathcal{B} + (d-p)\mu = nP/m$, so $\sum_j \operatorname{tr}(H_j) = (m/n)\operatorname{tr}(G_*) = P$. This construction satisfies the floor and attains $p/\mathcal{B}$ on every survivor set.

Equality and uniqueness. Suppose a feasible design attains $L = p/\mathcal{B}$. For every C,

$$L \geq \lambda_{\max}(U^\top G_C^{-1} U) \geq \lambda_{\max}(T_C^{-1}) \geq p/\operatorname{tr}(T_C). \tag{55}$$

Thus $\operatorname{tr}(T_C) \geq \mathcal{B}$. Its survivor-set average is at most $\mathcal{B}$, so equality holds for every C. Equality in the last spectral inequality forces $T_C = (\mathcal{B}/p)I_p$. Let $S_C = U_\perp^\top G_C U_\perp$. The floor gives $S_C \succeq \mu I$, while its survivor-set average satisfies

$$\mathbb{E}_C \operatorname{tr}(S_C) = \mathbb{E}_C \operatorname{tr}(G_C) - \mathbb{E}_C \operatorname{tr}(T_C) = \frac{nP}{m} - \mathcal{B} = (d-p)\mu. \tag{56}$$

Hence $S_C = \mu I$ for every C. A positive-semidefinite block matrix with a zero diagonal block must have a zero corresponding cross block. Thus, in the basis $[U, U_\perp]$, positivity of $G_C - \mu I$ forces the cross block to vanish, and $G_C = G_*$ for every survivor set.

If $1 \leq q \leq m-1$, then for any $i \neq j$ we can choose an $(n-1)$-subset excluding both indices and add either i or j. Equality of the resulting survivor Grams gives $H_i = H_j$, so $H_j = G_*/n$ for every j. □

Proposition 12 (Exact class-pair relaxation) *Under the rank-relaxed feasible family of Theorem 3,*

$$\begin{aligned}\gamma_{\mathrm{pair}}^{\mathrm{rel}} &= \min_{\substack{T \succeq \mu I_p \\ \operatorname{tr}(T)=\mathcal{B}}} \max_\ell z_\ell^\top T^{-1} z_\ell \\ &= \min_{T,\gamma} \gamma \quad \textit{subject to} \quad T \succeq \mu I_p,\ \operatorname{tr}(T) = \mathcal{B}, \quad \begin{bmatrix} T & z_\ell \\ z_\ell^\top & \gamma \end{bmatrix} \succeq 0 \quad \forall \ell. \end{aligned} \tag{57}$$

Moreover, $\gamma_{\mathrm{pair}}^{\mathrm{rel}} = p/\mathcal{B}$ if and only if

$$\exists \pi \in \Delta_{|\mathcal{D}|} : \qquad \sum_\ell \pi_\ell z_\ell z_\ell^\top = \frac{1}{p} I_p. \tag{58}$$

Otherwise the exact classifier-specific optimum is strictly smaller than $p/\mathcal{B}$.

Proof of Proposition 12. Let $\phi(G) = \max_\ell \widehat{v}_\ell^\top G^{-1} \widehat{v}_\ell$, which is convex on the positive-definite cone. For any feasible rank-relaxed block Grams,

$$\overline{G} := \mathbb{E}_C G_C = \frac{n}{m}\sum_j H_j, \qquad \operatorname{tr}(\overline{G}) = \frac{nP}{m}. \tag{59}$$

Jensen's inequality gives $\max_C \phi(G_C) \geq \mathbb{E}_C \phi(G_C) \geq \phi(\overline{G})$. Conversely, any $G \succeq \mu I$ with $\operatorname{tr}(G) = nP/m$ is attained homogeneously by $H_j = G/n$. Hence the rank-relaxed class-pair problem equals

$$\min_{G \succeq \mu I,\ \operatorname{tr}(G) = nP/m} \phi(G). \tag{60}$$

For a feasible G, let $T_0 = U^\top G U$. The floor gives $T_0 \succeq \mu I_p$ and $\mathrm{tr}(T_0) \le \mathcal{B}$. The compression-of-inverse inequality yields $U^\top G^{-1} U \succeq T_0^{-1}$. Increasing T_0 isotropically until its trace is $\mathcal{B}$ can only decrease every inverse quadratic. Conversely, every feasible T is realized by $G = UTU^\top + \mu(I - UU^\top)$. This proves the first line of Eq. (57); the second follows from the Schur complement.

For the equality criterion, define $Q_\pi = \sum_\ell \pi_\ell z_\ell z_\ell^\top$. Convex minimax duality gives

$$\gamma_{\mathrm{pair}}^{\mathrm{rel}} = \max_{\pi \in \Delta_{|\mathcal{D}|}} \min_{\substack{T \succeq \mu I_p \\ \mathrm{tr}(T) = \mathcal{B}}} \mathrm{tr}(Q_\pi T^{-1}). \tag{61}$$

Specifically, Sion's theorem applies because the probability simplex and the feasible T-set are compact and convex, the objective is continuous, affine in π, and convex in T on $T \succeq \mu I_p$. The isotropic matrix $T_0 = (\mathcal{B}/p) I_p$ is strictly above the floor and gives $p/\mathcal{B}$ for every Q_π, because $\mathrm{tr}(Q_\pi) = 1$. If $Q_\pi = I_p/p$, the arithmetic–harmonic mean inequality shows that no feasible T can do better, proving sufficiency. Conversely, equality in the primal and strong duality produce a π_* for which T_0 minimizes the inner problem. The interior first-order condition under the trace constraint requires $T_0^{-1} Q_{\pi_*} T_0^{-1}$ to be a scalar multiple of the identity, and the unit trace of Q_{π_*} then forces $Q_{\pi_*} = I_p/p$. This proves necessity. □

Water-filling solution. If $\lambda_i(Q_\pi)$ are the eigenvalues, the inner minimizer in Eq. (61) has the water-filling form $t_i = \max\{\mu, c\sqrt{\lambda_i(Q_\pi)}\}$, where c is chosen so that $\sum_i t_i = \mathcal{B}$.

Proposition 13 (Exact task–reconstruction tradeoff) *Let $\bar{t} = nP/m$. Every $G \succ 0$ with $\mathrm{tr}(G) = \bar{t}$ satisfies*

$$\mathrm{tr}(G^{-1}) \ge d^2/\bar{t}, \tag{62}$$

with equality only at $(\bar{t}/d) I$. For G_*, define $\Delta_{\mathrm{task}} = d/\bar{t} - p/\mathcal{B}$ and $\Delta_{\mathrm{rec}} = \mathrm{tr}(G_*^{-1}) - d^2/\bar{t}$. These quantities satisfy

$$\Delta_{\mathrm{rec}} = \frac{\bar{t} - d\mu}{\mu} \Delta_{\mathrm{task}} > 0.$$

Thus the relaxed optimum's reduction in task-direction uncertainty incurs the reconstruction increase $\Delta_{\mathrm{rec}} = ((\bar{t} - d\mu)/\mu)\Delta_{\mathrm{task}}$.

Proof of Proposition 13. Let $\lambda_1, \dots, \lambda_d > 0$ be the eigenvalues of G. The arithmetic–harmonic mean inequality gives

$$\left(\sum_{i=1}^d \lambda_i\right)\left(\sum_{i=1}^d \lambda_i^{-1}\right) \ge d^2, \tag{63}$$

with equality if and only if all eigenvalues equal $\bar{t}/d$. For G_*, the task eigenvalue is $a = (\bar{t} - (d-p)\mu)/p$. Since $\mu < \bar{t}/d$, we have $a > \bar{t}/d$ and hence $1/a < d/\bar{t}$. The two eigenvalues a and μ are unequal, so the inverse-trace inequality is strict for G_*, yielding $p/a + (d-p)/\mu > d^2/\bar{t}$.

Exact exchange. Substituting $\mathcal{B} = \bar{t} - (d-p)\mu$ and collecting terms gives

$$\frac{d}{\bar{t}} - \frac{p}{\mathcal{B}} = \frac{(d-p)(\bar{t} - d\mu)}{\bar{t}\mathcal{B}}, \tag{64}$$

$$\frac{p^2}{\mathcal{B}} + \frac{d-p}{\mu} - \frac{d^2}{\bar{t}} = \frac{(d-p)(\bar{t} - d\mu)^2}{\mu \bar{t} \mathcal{B}}. \tag{65}$$

The ratio of the second expression to the first is $(\bar{t} - d\mu)/\mu$, which proves Eq. (10). □

B.5 Proof of the anisotropic lifting

Proposition 14 (Exact class-pair allocation under a scalar factor floor) *Under the assumptions of Theorem 4, let $\tau = \mu/\beta$, $P_{\mathcal{U}} = P - (d-p)\tau$, and define the robust class-pair envelope*

$$\mathcal{R}_{\beta,\mathrm{pair}}(D) := \sup_{M \succeq \beta I} \max_\ell \widehat{v}_\ell^\top D^{-1} M^{-1} D^{-1} \widehat{v}_\ell. \tag{66}$$

Then the minimum over all SPD matrices D satisfying $\mathrm{tr}(D^2) = P$ and $\beta D^2 \succeq \mu I$ is

$$\min_D \mathcal{R}_{\beta,\mathrm{pair}}(D) = \frac{1}{\beta} \min_{\substack{T \succeq \tau I_p \\ \mathrm{tr}(T) = P_{\mathcal{U}}}} \max_\ell z_\ell^\top T^{-1} z_\ell. \tag{67}$$

If T_{pair} solves the right-hand problem, an optimal allocation is

$$D_{\mathrm{pair}}^2 = UT_{\mathrm{pair}}U^\top + \tau(I - UU^\top). \tag{68}$$

Moreover, D^* is optimal for Eq. (67) if and only if the isotropic-weighting condition Eq. (58) holds. Otherwise, D_{pair} strictly improves the scalar-floor robust class-pair envelope relative to D^*.

Proof. For $M \succeq \beta I$, inversion gives $M^{-1} \preceq \beta^{-1} I$, with equality at $M = \beta I$. Thus, writing $S = D^2$,

$$\mathcal{R}_{\beta,\mathrm{pair}}(D) = \frac{1}{\beta} \max_{\ell} \widehat{v}_\ell^\top S^{-1} \widehat{v}_\ell. \tag{69}$$

Complete U to an orthogonal basis and write

$$[U, U_\perp]^\top S [U, U_\perp] = \begin{bmatrix} S_{11} & S_{12} \\ S_{12}^\top & S_{22} \end{bmatrix}, \qquad K = S_{11} - S_{12} S_{22}^{-1} S_{12}^\top. \tag{70}$$

The task block of S^{-1} is K^{-1}. Since $S \succeq \tau I$, we have $S^{-1} \preceq \tau^{-1} I$ and hence $K \succeq \tau I_p$. Also $K \preceq S_{11}, S_{22} \succeq \tau I$, and therefore

$$\mathrm{tr}(K) \le \mathrm{tr}(S_{11}) = P - \mathrm{tr}(S_{22}) \le P - (d-p)\tau = P_{\mathcal{U}}. \tag{71}$$

Thus every feasible S induces $T = K + [P_{\mathcal{U}} - \mathrm{tr}(K)] I_p / p$, which satisfies $T \succeq \tau I_p$ and $\mathrm{tr}(T) = P_{\mathcal{U}}$. Because $T \succeq K$, every inverse quadratic at T is no larger than the corresponding task-block quadratic at S. This maps each feasible S to a feasible point of the reduced problem with an objective no larger. Conversely, every feasible T in the reduced problem induces the block-diagonal feasible matrix $S = UTU^\top + \tau(I - UU^\top)$ with the same objective, proving equality in Eq. (67).

The spectral allocation corresponds to $T = \rho I_p$, where $\rho = P_{\mathcal{U}}/p > \tau$. Applying the equality criterion in Proposition 12 after replacing $(\mu, \mathcal{B})$ by $(\tau, P_{\mathcal{U}})$ proves that this isotropic task allocation is optimal if and only if Eq. (58) holds; without such a weighting, the optimum is strictly below $1/(\beta\rho)$. □

Lemma 15 (Parseval–SPD factorization) *Every full-column-rank stacked frame A admits the unique factorization in Eq. (12) among factorizations with D symmetric positive definite and B Parseval. Its block partition satisfies $A_j = B_j D$, $G_C = D(B_C^\top B_C)D$, and $\mathrm{tr}(A^\top A) = \mathrm{tr}(D^2)$.*

Proof of Lemma 15. Full column rank makes $A^\top A$ positive definite, so its unique positive-definite square root $D = (A^\top A)^{1/2}$ is invertible. Setting $B = AD^{-1}$ gives

$$B^\top B = D^{-1} A^\top A D^{-1} = I. \tag{72}$$

The inherited row partition gives $A_j = B_j D$ and hence $G_C = D(B_C^\top B_C)D$; the trace identity follows from $D^2 = A^\top A$. For uniqueness, if $A = \widetilde{B}\widetilde{D}$ with $\widetilde{B}^\top \widetilde{B} = I$ and $\widetilde{D} \succ 0$, then $A^\top A = \widetilde{D}^2$. Uniqueness of the positive-definite square root gives $\widetilde{D} = D$ and then $\widetilde{B} = B$. □

Proof of Theorem 4. **Trace and frame survivor floor.** Because B is Parseval, $\sum_j \|B_j D^*\|_F^2 = \mathrm{tr}((D^*)^2) = p\rho + (d-p)\tau = P$. For every survivor set, $B_C^\top B_C \succeq \beta I$ and hence

$$G_C = D^* B_C^\top B_C D^* \succeq \beta (D^*)^2 \succeq \beta\tau I = \mu I. \tag{73}$$

Task-direction identity. $G_C^{-1} = (D^*)^{-1}(B_C^\top B_C)^{-1}(D^*)^{-1}$. For $v \in \mathcal{U}$, $(D^*)^{-1} v = v/\sqrt{\rho}$, which yields the stated inverse-Gram identity. Finally, $\mu < \beta P/d$ implies $\tau < P/d$. Hence $\rho > P/d$ when $p < d$, while $\rho = P/d$ when $p = d$.

Robust envelope. For every $M \succeq \beta I$, inversion gives $M^{-1} \preceq \beta^{-1} I$ and therefore

$$U^\top D^{-1} M^{-1} D^{-1} U \preceq \beta^{-1} U^\top D^{-2} U. \tag{74}$$

Equality is attained at $M = \beta I$, proving Eq. (15).

Global optimality over SPD allocations. Write $S = D^2$, so the certified allocation constraints are $S \succeq \tau I$ and $\mathrm{tr}(S) = P$. First suppose $p < d$, complete U to an orthogonal basis $[U, U_\perp]$, and write

$$[U, U_\perp]^\top S [U, U_\perp] = \begin{bmatrix} S_{11} & S_{12} \\ S_{12}^\top & S_{22} \end{bmatrix}, \qquad K := S_{11} - S_{12} S_{22}^{-1} S_{12}^\top. \tag{75}$$

The task block of S^{-1} is K^{-1}. Because $S \succeq \tau I$, its principal block satisfies $S_{22} \succeq \tau I$. Also $K \preceq S_{11}$, and hence

$$\lambda_{\min}(K) \le \lambda_{\min}(S_{11}) \le \frac{\operatorname{tr}(S_{11})}{p} = \frac{P - \operatorname{tr}(S_{22})}{p} \le \frac{P - (d-p)\tau}{p} = \rho. \tag{76}$$

Thus every feasible SPD allocation has $\mathcal{R}_\beta(D) \ge 1/(\beta\rho)$. Equality is attained by $S^* = \rho\Pi + \tau(I - \Pi) = (D^*)^2$. Equality throughout Eq. (76) first forces $S_{22} = \tau I$ and $S_{11} = \rho I$. Then $K = \rho I - \tau^{-1} S_{12} S_{12}^\top$ can have smallest eigenvalue ρ only if $S_{12} = 0$. Thus $S = S^*$, and the positive-definite square root gives $D = D^*$. If $p = d$, then $\lambda_{\min}(S) \le \operatorname{tr}(S)/d = P/d = \rho$, with equality only for $S = (P/d)I = (D^*)^2$. This proves optimality and uniqueness in all cases. □

Corollary 16 (Exact fixed-factor scaling and reconstruction decomposition)

(i) *Under the fixed-factor comparison in Eq. (16), squared uncertainty and uncertainty coefficients scale by* $\Gamma_{\text{alloc}}^{-1}$ *and* $\Gamma_{\text{alloc}}^{-1/2}$*, respectively, and Eq. (17) holds. For any absolute-radius interval* $[t_0, t_{\max}]$*,*

$$\int_{t_0}^{t_{\max}} \operatorname{CertAcc}_{\text{all}}^{D^*}(\epsilon)\, d\epsilon = \sqrt{\Gamma_{\text{alloc}}} \int_{t_0/\sqrt{\Gamma_{\text{alloc}}}}^{t_{\max}/\sqrt{\Gamma_{\text{alloc}}}} \operatorname{CertAcc}_{\text{all}}^{D_{\text{iso}}}(u)\, du. \tag{77}$$

(ii) *For* $F_C = B_C^\top B_C$ *and an orthonormal basis* $U_\perp$ *of* $\mathcal{U}^\perp$*,*

$$\begin{aligned} \operatorname{tr}(U^\top G_{C,D^*}^{-1} U) &= \rho^{-1} \operatorname{tr}(U^\top F_C^{-1} U), \\ \operatorname{tr}(U_\perp^\top G_{C,D^*}^{-1} U_\perp) &= \tau^{-1} \operatorname{tr}(U_\perp^\top F_C^{-1} U_\perp), \end{aligned} \tag{78}$$

When $p < d$*, their ratios to isotropic scaling are*

$$\frac{P}{d\rho} = \Gamma_{\text{alloc}}^{-1} < 1, \qquad \frac{P}{d\tau} = \frac{\beta P}{d\mu} > 1. \tag{79}$$

When $p = d$*,* $U_\perp$ *is empty and the task ratio is* $P/(d\rho) = \Gamma_{\text{alloc}}^{-1} = 1$*.*

(iii) *Moreover,*

$$1 < \Gamma_{\text{alloc}} < d/p \quad (p < d), \qquad \Gamma_{\text{alloc}} = 1 \quad (p = d), \tag{80}$$

and, holding all other scalar parameters fixed, it decreases with p *and* μ *and increases with* β*.*

Proof of Corollary 16. **Radius, curve, and continuous-radius AUC scaling.** For $v \in \mathcal{U}$, isotropic scaling gives

$$v^\top G_{C,\text{iso}}^{-1} v = \frac{d}{P} v^\top (B_C^\top B_C)^{-1} v = \Gamma_{\text{alloc}}\, v^\top G_{C,D^*}^{-1} v. \tag{81}$$

Thus every uncertainty coefficient scales by $\Gamma_{\text{alloc}}^{-1/2}$ under D^*. The nominal LS solution and all pairwise margins are unchanged, so minimizing the margin-to-uncertainty ratios proves Eq. (17). The change of variables $u = \epsilon/\sqrt{\Gamma_{\text{alloc}}}$ proves Eq. (18).

Reconstruction decomposition. For $F_C = B_C^\top B_C$, $G_{C,D^*}^{-1} = (D^*)^{-1} F_C^{-1} (D^*)^{-1}$. Because $(D^*)^{-1} U = U/\sqrt{\rho}$ and $(D^*)^{-1} U_\perp = U_\perp/\sqrt{\tau}$, taking traces proves Eq. (19). Isotropic scaling multiplies both components by d/P, yielding the ratios in Eq. (79).

Range and monotonicity. Let $x := \mu d/(\beta P) \in (0, 1)$. Then

$$\Gamma_{\text{alloc}} = 1 + \frac{d-p}{p}(1 - x), \qquad \frac{d}{p} - \Gamma_{\text{alloc}} = \frac{d-p}{p} x. \tag{82}$$

These identities prove Eq. (80). For $p < d$, the first expression decreases with p and μ, while it increases with β; for $p = d$ it equals one. □

B.6 Rank-constrained objective gap and Gram separation

Proposition 17 (Quantitative survivor-Gram separation) *Under the assumptions of Theorem 5, let $H_* = G_*/n$ and write its eigenvalues as $\lambda_1(H_*) \geq \cdots \geq \lambda_d(H_*)$. Every feasible design satisfies*

$$\left(\frac{1}{\binom{m}{n}} \sum_{|C|=n} \|G_C - G_*\|_F^2 \right)^{1/2} \geq \sqrt{\frac{nq}{m-1}} \left(\sum_{\ell=r+1}^{d} \lambda_\ell(H_*)^2 \right)^{1/2}. \tag{83}$$

Proof of Theorem 5 and Proposition 17. **Step 1: equality contradicts rank-constrained blocks.** Let $L = p/\mathcal{B}$ and suppose that a feasible rank-constrained design attains L. The equality characterization in the proof of Theorem 3 forces every survivor Gram to equal G_* and, because $q \geq 1$, every block Gram to equal G_*/n. This matrix has rank d, contradicting $r < d$.

Step 2: unattainability makes the gap strict. By assumption, $\mathcal{F}_r$ is nonempty. It is also compact: the positive-semidefinite, rank, trace, and floor constraints define a closed set; the fixed total trace bounds every H_j. Because every $G_C \succeq \mu I$, the objective is continuous and therefore attains its minimum on $\mathcal{F}_r$. Since L cannot be attained, this minimum strictly exceeds L.

Step 3: quantify survivor-Gram separation. Let $\overline{H} = m^{-1} \sum_j H_j$, $H_* = G_*/n$, and $X_j = H_j - \overline{H}$. For a uniformly sampled n-subset, $G_C - n\overline{H} = \sum_{j \in C} X_j$. Using $\sum_j X_j = 0$ and hence $\sum_{i \neq j} \langle X_i, X_j \rangle_F = -\sum_j \|X_j\|_F^2$ gives

$$\begin{aligned} \mathbb{E}_C \|G_C - n\overline{H}\|_F^2 &= \frac{n}{m} \sum_j \|X_j\|_F^2 + \frac{n(n-1)}{m(m-1)} \sum_{i \neq j} \langle X_i, X_j \rangle_F \\ &= \frac{nq}{m-1} \frac{1}{m} \sum_{j=1}^{m} \|H_j - \overline{H}\|_F^2. \end{aligned} \tag{84}$$

Using $m^{-1} \sum_j \|H_j - H_*\|_F^2 = m^{-1} \sum_j \|H_j - \overline{H}\|_F^2 + \|\overline{H} - H_*\|_F^2$ and $n^2 \geq nq/(m-1)$ yields

$$\mathbb{E}_C \|G_C - G_*\|_F^2 = \frac{nq}{m-1} \frac{1}{m} \sum_j \|H_j - \overline{H}\|_F^2 + n^2 \|\overline{H} - H_*\|_F^2 \tag{85}$$

$$\geq \frac{nq}{m-1} \frac{1}{m} \sum_j \|H_j - H_*\|_F^2. \tag{86}$$

By the Eckart–Young theorem (Eckart & Young, 1936), every rank-at-most-r matrix H_j has Frobenius distance at least $\left(\sum_{\ell=r+1}^{d} \lambda_\ell(H_*)^2\right)^{1/2}$ from H_*. Averaging over j and taking square roots yields Eq. (83). Since $\mathcal{B}/p > \mu$, H_* has p leading eigenvalues $\mathcal{B}/(pn)$ and $d - p$ eigenvalues μ/n. □

The spectral-tail term in Eq. (83) has two explicit forms. When $r \geq p$, the lower bound becomes

$$\mu \sqrt{\frac{q(d-r)}{n(m-1)}}. \tag{87}$$

When $r < p$, its squared spectral-tail factor is

$$(p-r) \left(\frac{\mathcal{B}}{pn}\right)^2 + (d-p) \left(\frac{\mu}{n}\right)^2. \tag{88}$$

C Experimental Details

C.1 Architectures and optimization

For each of the five image datasets, we train ResNet-18 (He et al., 2016) from scratch with batch size 128 and Nesterov stochastic gradient descent (SGD). The initial learning rate is 0.1, momentum is 0.9, and weight decay is 5×10^{-4}; schedules use cosine annealing and fixed dataset-specific augmentation policies. Evaluation transforms are deterministic. UCI HAR uses a one-dimensional

convolutional neural network (1D CNN) with three stages (channels $64, 128, 128$; kernels $7, 5, 3$), batch normalization, ReLU, two max-pooling stages, adaptive average pooling, and a linear 128-dimensional latent projection. For Banking77, the frozen `all-MiniLM-L6-v2` encoder (Wang et al., 2020; Reimers & Gurevych, 2019) produces unit-normalized 384-dimensional embeddings by attention-mask mean pooling. A bias-free learned projection maps them to $d = 128$. Both non-image tasks use the same batch size and optimizer family with initial learning rate 0.05. For every architecture and dataset, the final projected latent $s = f_\theta(x)$ is ℓ_2-normalized before it enters the frame and classifier.

For factorized methods, reduced QR factorization maps the raw block parameters to a Parseval factor. At the start of each epoch, a singular value decomposition (SVD) of the centered classifier weights defines p by retaining singular values $\sigma_\ell > 10^{-6}\sigma_1$. The retained rank is $p = K - 1$ (Tab. 5), so the numerical basis spans all class-pair directions. The current head and Parseval factor determine the anisotropic scaling. D^* is updated whenever $\beta > \mu d/P$ and remains fixed within the epoch. In the $q = 2$ runs, survivor-trained methods sample two of the 28 subsets per minibatch. Evaluation uses unregularized least squares throughout. Learned and fixed Parseval factors use the same reduced-QR map and seed-matched Gaussian initialization. Schedules are fixed at 200 epochs, except PCam at 35 and Banking77 at 100; results use the model from the final epoch.

C.2 Preprocessing and data splits

CIFAR-10/100 use random 32×32 crops with four-pixel padding and horizontal flips. STL-10 uses random 96×96 crops with 12-pixel padding and horizontal flips. EuroSAT uses random resized 32×32 crops with scale $[0.7, 1]$ and both horizontal and vertical flips; PCam uses random 96×96 crops with 12-pixel padding and the same two flips. UCI HAR uses raw inertial windows without augmentation; Banking77 trains on frozen embeddings without text augmentation.

CIFAR-10 and CIFAR-100 divide their official training sets into 45,000 training and 5,000 development examples. Their official 10,000-example test sets provide the final evaluation. STL-10 uses the same fixed split rule for a 4,000/1,000 division of its official training set and evaluates on the official 8,000-example test set. EuroSAT uses 17,000 training, 5,000 development, and 5,000 disjoint test examples; PCam uses the official 262,144/32,768/32,768 train/development/test partition. UCI HAR uses raw 9×128 inertial windows and the official subject-disjoint test partition, with subjects $6, 15, 23, 29$ held out from the official training subjects for development and the remaining subjects used for training. Banking77 uses the official 10,003/3,080 train/test split and a deterministic stratified development split of ten training examples per class, yielding 9,233 training and 770 development examples.

C.3 Certification metrics and statistical analysis

Certification is evaluated using the final model's certified-accuracy curve. Objective and configuration choices use development data; official test sets provide the final evaluation.

The absolute-radius grid, used for CIFAR-10, STL-10, and EuroSAT, is

$$\epsilon \in \{0.02, 0.10, 0.50, 0.75, 1.00, 1.50\}.$$

CIFAR-100, PCam, UCI HAR, and Banking77 use control-energy normalization to compare different task ranks, latent dimensions, and modalities. This choice and both grids were chosen using training and development data before held-out test evaluation. Let $\overline{E}_{\mathrm{control,train}}$ denote the fixed-factor control's mean training-set codeword energy. CIFAR-100, PCam, UCI HAR, and Banking77 use

$$\delta \in \{0, 0.02, 0.05, 0.10, 0.20, 0.30, 0.42, 0.60, 0.80\},$$

where $\delta = \epsilon/\sqrt{\overline{E}_{\mathrm{control,train}}}$. For every example, define the nonnegative absolute and control-energy-normalized radii

$$\varrho_i^{(\epsilon)} := \max\left\{0, \min_{\substack{|C|=m-q \\ k \neq y_i}} \frac{(w_{y_i} - w_k)^\top \widehat{s}^0_{i,C}}{c_{y_i k}(C)}\right\}, \qquad \varrho_i^{(\delta)} := \frac{\varrho_i^{(\epsilon)}}{\sqrt{\overline{E}_{\mathrm{control,train}}}}. \tag{89}$$

On the absolute grid, $t = \epsilon$; on the normalized grid, $\epsilon = t\sqrt{\overline{E}_{\text{control,train}}}$. For a radius grid $t_0 < \cdots < t_L$, the grid tables and figures report grid AUC as the unnormalized trapezoidal integral

$$\text{AUC} := \sum_{\ell=1}^{L} \frac{\text{CertAcc}_{\text{all}}(t_{\ell-1}) + \text{CertAcc}_{\text{all}}(t_\ell)}{2}(t_\ell - t_{\ell-1}). \tag{90}$$

This metric has a direct radius interpretation. Write $\widehat{y}_i$ for the clean prediction and let $\varrho_i = \varrho_i^{(\epsilon)}$ on an absolute grid and $\varrho_i = \varrho_i^{(\delta)}$ on a normalized grid. For the underlying continuous certification curve and $t_{\max} = t_L$, the tail-integral identity gives

$$\int_{t_0}^{t_{\max}} \text{CertAcc}_{\text{all}}(t)\, dt = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\{\widehat{y}_i = y_i\}\,[\min(\varrho_i, t_{\max}) - t_0]_+. \tag{91}$$

Continuous-radius AUC is therefore the dataset-wide sample mean of $[\min(\varrho_i, t_{\max}) - t_0]_+$ for clean-correct examples, with clean-incorrect examples contributing zero; Eq. (90) is its grid approximation. For cross-dataset comparisons, the range-normalized grid AUC is the grid AUC divided by the applicable grid range: 1.48 on the absolute-ϵ grid and 0.8 on the normalized-δ grid. The erasure-severity analysis separately evaluates all datasets on the latter shared δ grid; we call its range-normalized quantity common-δ AUC. In finite-precision evaluation, define $\widehat{y}_{i,C} := \arg\max_k w_k^\top \widehat{s}^0_{i,C}$. We use $\mathbf{1}\{\widehat{y}_{i,C} = y_i \text{ for every admissible } C\}$ as the clean-correct indicator and certify only strict inequalities. Under exact LS invariance, this is identical to the indicator in Eq. (91). The independent unit is the training seed; samples and survivor sets are evaluation conditions.

C.4 Factorized objective and training algorithm

The term $\text{CE}_{i,C}$ is the cross-entropy of logits computed from $\widehat{s}^0_{i,C}$, $[t]_+ = \max(t, 0)$, and γ_{tr} is the training margin buffer. The score uncertainty is evaluated at the fixed radius ϵ_{train}. The robust-loss configuration uses $\epsilon_{\text{train}} = 0.02$, $\gamma_{\text{tr}} = 0.05$, $\mu = 0.01$, $\lambda_{\text{cert}} = 1$, and $\lambda_{\text{floor}} = 10$ wherever the corresponding terms apply across methods and datasets. For factorized methods, choose an exact worst-survivor set $C_\beta \in \arg\min_{|C|=m-q} \lambda_{\min}(B_C^\top B_C)$ for the current Parseval factor, and let $\mathcal{S}$ be the training survivor family and $\xi_{\text{floor}} = 10^{-4}$. In the $q = 2$ setting, $\mathcal{S}$ contains two sampled sets. The floor term in Eq. (22) is

$$\mathcal{L}_{\text{floor}} = \left[\mu - \min_{C \in \mathcal{S}\cup\{C_\beta\}} \lambda_{\min}(A_C^\top A_C)\right]_+ + \left[\mu d/P + \xi_{\text{floor}} - \lambda_{\min}(B_{C_\beta}^\top B_{C_\beta})\right]_+. \tag{92}$$

The first term penalizes required-floor violations on the sampled sets and the most recently recomputed C_β; the second penalizes violations of the factor-floor condition required by the analytic allocation. The exact C_β selected at each allocation update is held fixed within the epoch. Both terms are evaluated on the current QR-projected factor during optimization. After training, we exhaustively verify the frame survivor-floor constraint over all admissible sets.

Algorithm 2 Training and deployment of TACBF. Here $\text{qf}(\widetilde{B})$ denotes the thin-QR Q factor reshaped into m blocks.

Input: training data, encoder f_θ, head W, unconstrained factor parameters $\widetilde{B}$, and (m, q, r, P, μ).
Initialize: $B \leftarrow \text{qf}(\widetilde{B})$ and $D^* \leftarrow \sqrt{P/d}\, I$.
for epoch $t = 1, \ldots, T_{\text{ep}}$ **do**
  Form the centered-head basis U and compute $\beta = \min_{|C|=m-q} \lambda_{\min}(B_C^\top B_C)$.
  If $\beta \le \mu d/P$, retain the current D^* (isotropic at initialization); otherwise set $\tau = \mu/\beta$, $\rho = [P - (d-p)\tau]/p$, and $D^* = \sqrt{\rho} UU^\top + \sqrt{\tau}(I - UU^\top)$.
  for each minibatch **do**
    Set $B \leftarrow \text{qf}(\widetilde{B})$ and $A_j \leftarrow B_j D^*$; select the training survivor family $\mathcal{S}$.
    Recover each $\widehat{s}^0_{i,C}$ by noiseless exact least squares from $A_C s_i$ and evaluate the objective in Eq. (22).
    Update $(\theta, W, \widetilde{B})$ by stochastic gradient descent (SGD) through the QR map.
  end for
end for
Output: Recompute B, U, and β from the final parameters; require $\beta > \mu d/P$, recompute D^*, and construct $A_j = B_j D^*$ for certification and deployment.

C.5 Experimental controls and computational cost

All trained methods use the same dataset-specific encoder, preprocessing, augmentation, optimization schedule, and LS decoder. Unless stated otherwise, experiments use $m = 8$, $q = 2$, $r = d/4$, and $P = d$, and evaluation enumerates all $\binom{8}{2} = 28$ survivor sets. Hyperparameters are selected using training and development data; held-out test sets are used only for final reporting. Each TACBF/fixed-factor pair shares the training seed and initial Parseval factor. The mean-codeword-energy comparison separately allows the trace to vary.

Table 5: Dataset configurations and relative computational cost. All listed configurations use $m = 8$, $q = 2$, $mr/d = 2$, and $P = d$. Training-time and peak-memory ratios compare TACBF with the paired fixed-factor control.

Dataset	K	p	d	p/d	Radius grid	Relative train time	Relative peak memory
CIFAR-10	10	9	128	0.070	Absolute ϵ	1.201	1.0005
STL-10	10	9	128	0.070	Absolute ϵ	1.015	1.0001
EuroSAT RGB	10	9	128	0.070	Absolute ϵ	1.110	1.0005
CIFAR-100	100	99	512	0.193	Normalized δ	1.033	1.0183
PCam	2	1	128	0.008	Normalized δ	1.017	1.0001
UCI HAR	6	5	128	0.039	Normalized δ	1.115	1.0142
Banking77	77	76	128	0.594	Normalized δ	1.059	1.0239

Table 6: Realized frame constraints across five seeds per dataset. The first two columns take minima over seeds and survivor sets. Because $G_C \preceq \rho I$, the third column reports the conservative condition-number bound $\kappa_{\mathrm{ub}} := \rho/\lambda_{\min}(G_C) \geq \kappa(G_C)$, maximized over seeds and survivor sets. Trace error is $|\operatorname{tr}(A^\top A) - P|/P$.

Dataset	$\min \alpha_q(A)$	$\min \beta_q(B)$	Upper bound on worst-survivor condition number	Max. relative trace error
CIFAR-10	0.010039	0.081594	1276.5	5.6×10^{-16}
STL-10	0.010041	0.066238	1236.9	1.6×10^{-15}
EuroSAT RGB	0.010035	0.077570	1255.9	4.4×10^{-16}
CIFAR-100	0.010594	0.054709	418.7	2.2×10^{-16}
PCam	0.010002	0.095483	11648.4	4.4×10^{-16}
UCI HAR	0.010138	0.051625	2125.6	6.7×10^{-16}
Banking77	0.013967	0.106431	116.5	4.4×10^{-16}

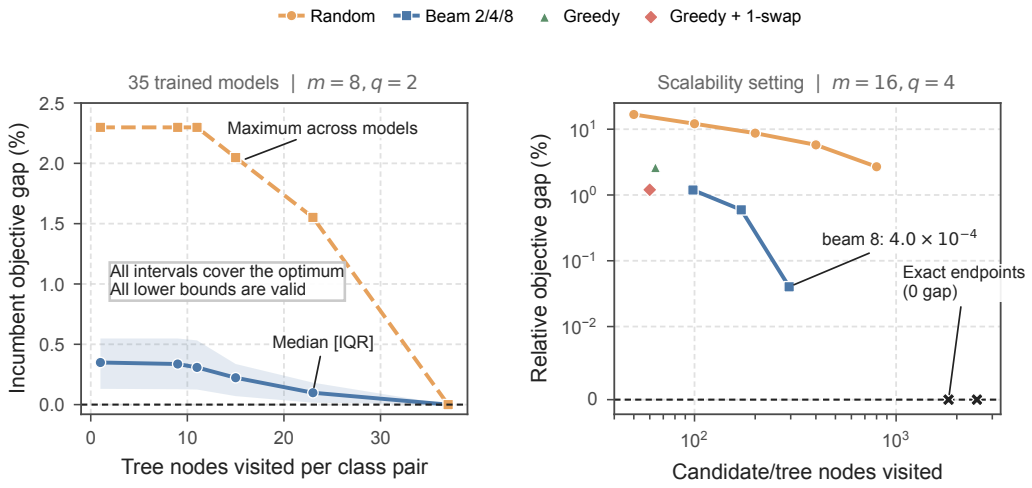
Variation in task rank and latent dimension. CIFAR-100 uses $(d, r, P) = (512, 128, 512)$, whereas PCam, UCI HAR, and Banking77 use $(128, 32, 128)$. All preserve $mr/d = 2$, the $q/m = 1/4$ erased fraction, and trace per latent coordinate; CIFAR-100 also satisfies $(m - q)r = 768 \geq 512$. The 512-dimensional setting uses ResNet-18's native final feature. Every baseline receives the same dataset-specific dimensions, schedule, and energy-normalized δ grid. Banking77 lies in the $r < p$ regime of the rank-gap bound; the other six $q = 2$ configurations satisfy $r \geq p$.

D Exact Survivor Search and Scalability

D.1 Search bounds and convergence

On the trained TACBF models, completed branch-and-bound matches exhaustive enumeration at $q = 2$ and $q = 3$. At each $q = 2$ search budget, the intervals contain the optimum and the radius lower bounds yield valid certifications. The $q = 3$ intervals also contain every enumerated optimum; this analysis measures transfer from frames trained with the $q = 2$ survivor-floor constraint.

In the $m = 16$ scalability experiment, full enumeration takes 10.4 seconds, whereas beam width eight visits 295 tree nodes per class pair on average and reaches a mean relative objective gap of 4.0×10^{-4}. Completed branch-and-bound closes the gap to zero.



(a) Anytime incumbent convergence

(b) Search quality versus computation

Figure 4: Exact-search convergence analysis. (a) Across the trained $q = 2$ models, the median incumbent objective gap and its interquartile range contract as branch-and-bound visits more tree nodes; the dashed curve is the largest model-averaged gap. All objective intervals contain the exhaustive optimum and every radius lower bound yields a valid certification. (b) In the separate $(m, q) = (16, 4)$ scalability experiment, larger beam widths reduce objective error with substantially fewer visited nodes than exhaustive enumeration; exhaustive search and completed branch-and-bound attain the zero-gap endpoints.

E Additional Results

Table 7: Seven-method comparison setup. Certificate and survivor-floor penalties refer to the regularizers in Eq. (22).

Method	Frame and allocation	Training or derivation
TACBF	Learned Parseval factor B obtained by reduced QR; D^* updated each epoch when the current head and survivor floor β satisfy the feasibility condition	Maximum survivor-set cross-entropy plus certificate and survivor-floor penalties
Fixed-factor control	Seed-paired Parseval factor obtained by reduced QR and held fixed; the same conditional D^* update	Same objective as TACBF; the encoder and head are optimized
Frame-Codes-inspired BUTF	Hadamard-derived Parseval factor selected by simulated annealing of the survivor-capacity objective; the same conditional D^* update	Same matched training objective as the fixed-factor control
Analytic task allocation	Closed-form D^* applied to the trained model of the matched fixed-tight frame	Closed-form evaluation from the trained model
Label-free max–min frame	Label-free optimization of a max–min Parseval frame, then held fixed during classifier training	Full-block clean cross-entropy with the optimized frame fixed
Block dropout	Learned trace-normalized frame A	Maximum survivor-set cross-entropy under sampled block erasures
Fixed-tight frame	Fixed seed-matched random tight frame	Full-block clean cross-entropy

Frame-Codes-inspired BUTF. Let Φ_C contain the surviving columns of the row-selected and block-permuted Hadamard frame. The search maximizes the exact mean survivor capacity $\binom{8}{2}^{-1} \sum_C \log_2 \det(I_d + 100\Phi_C\Phi_C^\top)$ at 20 dB. Simulated annealing proposes a row swap with probability 0.5, otherwise swapping two columns from different blocks. Proposal scores use four

sampled survivor sets. The temperature decreases geometrically from 0.05 to 10^{-4} over 2,000 steps for $d = 128$ and 1,000 for $d = 512$. Every 50 steps and after the final update, we evaluate the objective over all 28 sets. Among candidates satisfying $\min_C \lambda_{\min}(\Phi_C \Phi_C^\top) \geq \mu + \xi_{\text{floor}} = 0.0101$, we select the one with the largest exact mean capacity. The search is initialized from a feasible randomized Hadamard construction.

Label-free max–min frame. This baseline solves $\max_B \min_{|C|=6} \lambda_{\min}(B_C^\top B_C)$ before classifier training. Optimization uses the differentiable surrogate $-T \log \sum_C \exp[-\lambda_{\min}(B_C^\top B_C)/T]$ with $T = 0.002$, all 28 survivor sets, and 500 Adam steps at learning rate 0.01. From a Gaussian initialization, each step applies a reduced-QR Parseval projection and clips the gradient norm at 10. The iterate with the largest exact survivor floor is then held fixed during classifier training.

E.1 Results across datasets and baselines

TACBF's positive grid-AUC gains over the fixed-factor control extend from $p = 1$ to $p = 99$ and from $d = 128$ to $d = 512$ across image, inertial, and text representations.

Table 8: Paired five-seed results for the comparison in Tab. 1. Cells report TACBF/fixed-factor and their difference with a 95% CI. Grid AUC is the unnormalized grid integral; certified accuracy uses $\epsilon = 1.5$ or $\delta = 0.42$ and is reported in percent (differences and CIs in pp).

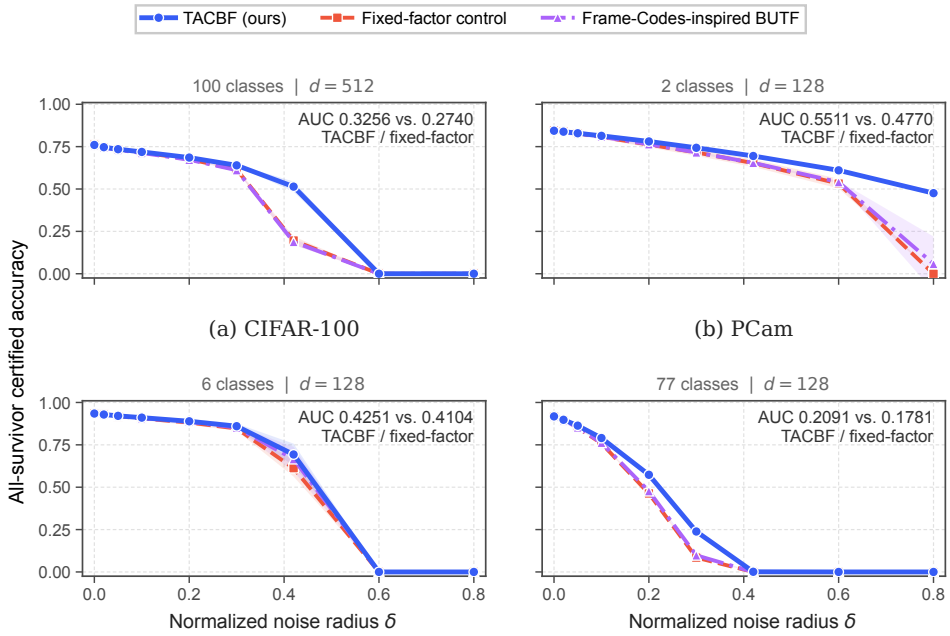
Grid AUC ↑				
Setting	Dataset (K, d)	TACBF	Fixed-factor control	Δ [95% CI]
Absolute ϵ	CIFAR-10 $(10, 128)$	$\mathbf{1.3606}$	1.3516	$+\mathbf{0.0090}$ $[0.0039, 0.0141]$
	STL-10 $(10, 128)$	$\mathbf{1.0047}$	0.8910	$+\mathbf{0.1137}$ $[0.0900, 0.1374]$
	EuroSAT RGB $(10, 128)$	$\mathbf{1.4171}$	1.3882	$+\mathbf{0.0289}$ $[0.0130, 0.0449]$
Normalized δ	CIFAR-100 $(100, 512)$	$\mathbf{0.3256}$	0.2740	$+\mathbf{0.0516}$ $[0.0466, 0.0565]$
	PCam $(2, 128)$	$\mathbf{0.5511}$	0.4770	$+\mathbf{0.0741}$ $[0.0634, 0.0848]$
	UCI HAR $(6, 128)$	$\mathbf{0.4251}$	0.4104	$+\mathbf{0.0148}$ $[0.0050, 0.0245]$
	Banking77 $(77, 128)$	$\mathbf{0.2091}$	0.1781	$+\mathbf{0.0311}$ $[0.0295, 0.0327]$
Certified accuracy at $\epsilon = 1.5$ or $\delta = 0.42$ ↑ (%)				
Setting	Dataset (K, d)	TACBF	Fixed-factor control	Δ [95% CI] (pp)
Absolute ϵ	CIFAR-10 $(10, 128)$	$\mathbf{87.02}$	84.92	$+\mathbf{2.10}$ $[1.03, 3.17]$
	STL-10 $(10, 128)$	$\mathbf{47.33}$	14.81	$+\mathbf{32.52}$ $[25.41, 39.63]$
	EuroSAT RGB $(10, 128)$	$\mathbf{90.54}$	80.52	$+\mathbf{10.01}$ $[2.78, 17.24]$
Normalized δ	CIFAR-100 $(100, 512)$	$\mathbf{51.40}$	19.60	$+\mathbf{31.79}$ $[28.90, 34.69]$
	PCam $(2, 128)$	$\mathbf{69.48}$	65.28	$+\mathbf{4.21}$ $[2.58, 5.83]$
	UCI HAR $(6, 128)$	$\mathbf{69.30}$	61.26	$+\mathbf{8.05}$ $[1.86, 14.24]$
	Banking77 $(77, 128)$	$\mathbf{0.10}$	0.00	$+\mathbf{0.10}$ $[0.05, 0.15]$

Table 9: Clean accuracy and reconstruction error at $q = 2$ (five-seed means). Clean accuracy is reported in percent. Expected squared reconstruction error is $\mathbb{E}_C \operatorname{tr}(G_C^{-1})/((m-q)r)$ for unit-sphere observation noise. Entries are reported as TACBF/fixed-factor control.

Dataset	Clean accuracy (%)	Expected squared reconstruction error
CIFAR-10	$94.800/94.790$	$10.2574/7.2368$
STL-10	$81.178/81.150$	$8.8557/7.1912$
EuroSAT RGB	$98.260/98.368$	$9.7799/7.2006$
CIFAR-100	$75.990/76.136$	$7.3056/7.0627$
PCam	$84.380/84.625$	$12.5943/7.6342$
UCI HAR	$93.465/93.458$	$7.6547/7.3957$
Banking77	$91.779/91.812$	$5.9747/3.6544$

The largest absolute difference between dataset-level mean clean accuracies is 0.2441 percentage points. For the three absolute-radius datasets, we additionally match TACBF to the fixed-factor control's mean codeword energy; the corresponding certification comparisons are reported in Tab. 13.

Results across task ranks and modalities. Under mean-codeword-energy matching, the paired five-seed grid-AUC gains are $+0.0657$ $[0.0527, 0.0788]$ on CIFAR-100 and $+0.0412$ $[0.0301, 0.0524]$ on PCam. The UCI HAR experiment changes both the input domain and encoder family while preserving the latent block setting; the Banking77 experiment evaluates frozen language-model representations and 77-way intent classification. TACBF attains the highest grid AUC on both. Mean-codeword-energy matching on UCI HAR preserves a paired gain of $+0.0175$ $[0.0055, 0.0296]$; on Banking77, it yields $+0.0347$ $[0.0325, 0.0368]$.



(a) CIFAR-100 (b) PCam

(c) UCI HAR (d) Banking77

Figure 5: Five-seed certification on the normalized δ grid. Lines are seed means and bands are 95% CIs. Every panel compares TACBF with the fixed-factor control and Frame-Codes-inspired BUTF designs under matched width and trace.

Table 10: Paired comparison with Frame-Codes-inspired BUTF. Grid AUC is the unnormalized grid integral; intervals use the same five training/construction seeds. Bold means mark the higher grid AUC in each row; a bold difference has a CI excluding zero.

Dataset	**TACBF**	BUTF	Δ [95% CI]
CIFAR-10	**1.3606**	1.3550	$+0.0056$ $[-0.0022, 0.0133]$
STL-10	**1.0047**	0.9020	$+\mathbf{0.1027}$ $[0.0771, 0.1283]$
EuroSAT RGB	**1.4171**	1.3991	$+\mathbf{0.0180}$ $[0.0129, 0.0231]$
CIFAR-100	**0.3256**	0.2715	$+\mathbf{0.0541}$ $[0.0475, 0.0606]$
PCam	**0.5511**	0.4850	$+\mathbf{0.0661}$ $[0.0478, 0.0844]$
UCI HAR	**0.4251**	0.4196	$+0.0055$ $[-0.0095, 0.0206]$
Banking77	**0.2091**	0.1813	$+\mathbf{0.0278}$ $[0.0268, 0.0289]$

Multiplicity-adjusted inference. For grid AUC, Holm correction is applied across the 14 dataset–baseline comparisons, comprising seven datasets for each of the fixed-factor and BUTF contrasts.

All fixed-factor differences and five BUTF differences remain significant; the BUTF contrasts on CIFAR-10 and UCI HAR have adjusted $p = 0.2322$ and 0.3645.

Table 11: Continuous-radius AUC from Eq. (91) for the paired five-seed comparison. Values integrate over the same stated radius ranges as the grid metric and use the same all-survivor clean-correct convention. Intervals are paired 95% CIs; two-sided paired t-tests use the five seed-level differences ($\mathrm{df} = 4$), and p_{Holm} adjusts over the seven dataset comparisons.

Dataset	**TACBF**	Fixed-factor control	Δ [95% CI]	p_{Holm}
CIFAR-10	**1.3629**	1.3550	**+0.0078** [0.0030, 0.0126]	0.0212
STL-10	**1.0113**	0.9316	**+0.0797** [0.0623, 0.0972]	0.0011
EuroSAT RGB	**1.4200**	1.4048	**+0.0153** [0.0127, 0.0178]	0.0004
CIFAR-100	**0.3058**	0.2701	**+0.0357** [0.0274, 0.0440]	0.0011
PCam	**0.5528**	0.4966	**+0.0562** [0.0400, 0.0724]	0.0019
UCI HAR	**0.3960**	0.3770	**+0.0190** [0.0074, 0.0307]	0.0212
Banking77	**0.2061**	0.1753	**+0.0308** [0.0293, 0.0324]	$< 10^{-5}$

Table 12: Paired seed-level continuous-radius AUC differences (TACBF minus fixed-factor control). Columns correspond to the five paired training seeds; all 35 differences are positive.

Dataset	Seed 1	Seed 2	Seed 3	Seed 4	Seed 5
CIFAR-10	0.008350	0.010231	0.011221	0.001312	0.007938
STL-10	0.090574	0.070631	0.098358	0.065281	0.073775
EuroSAT RGB	0.016105	0.013813	0.018480	0.013750	0.014206
CIFAR-100	0.039754	0.034213	0.024600	0.038928	0.040890
PCam	0.051666	0.073004	0.062366	0.056228	0.037810
UCI HAR	0.030310	0.008201	0.011304	0.026059	0.019344
Banking77	0.029971	0.032005	0.030034	0.029721	0.032398

E.2 Controlled comparisons of allocation and factor geometry

The shared-D^* comparison holds the trained model and D^* fixed and differs only in B.

Table 13: Learned-factor certification gains at $\epsilon = 1.5$ under two five-seed controlled comparisons. Entries are accuracy fractions with paired 95% intervals; bold differences have intervals excluding zero.

Dataset	Mean-codeword-energy matched ↑	Fixed model, shared D^* ↑
CIFAR-10	**+0.0126** [0.0007, 0.0245]	**+0.2598** [0.0551, 0.4645]
STL-10	**+0.3090** [0.2408, 0.3772]	**+0.4014** [0.3557, 0.4471]
EuroSAT RGB	**+0.0898** [0.0203, 0.1592]	**+0.2736** [0.1251, 0.4222]

Grid AUC complements the certified-accuracy results at $\epsilon = 1.5$ in Tab. 13. After mean-codeword-energy matching, the gains are $+0.0059$ $[0.0008, 0.0110]$, $+0.1066$ $[0.0838, 0.1294]$, and $+0.0255$ $[0.0107, 0.0404]$ on CIFAR-10, STL-10, and EuroSAT RGB. Under a shared D^*, the corresponding learned-minus-fixed-factor differences are $+0.0729$ $[0.0213, 0.1244]$, $+0.1385$ $[0.1224, 0.1545]$, and $+0.0762$ $[0.0392, 0.1131]$. Both controlled comparisons produce positive grid-AUC gains on all three datasets.

These gains persist after matching mean codeword energy on all three absolute-radius datasets; dataset-level mean clean accuracy differs by at most 0.25 percentage points.

Illustrative geometry. Panel (a) of Fig. 3 visualizes a validation example chosen to illustrate the geometric contrast between the two methods. The aggregate comparisons use all examples and

survivor sets. The displayed plane spans the normalized class-pair direction and the highest-variance orthogonal task direction under the fixed survivor inverse Gram.

Separating allocation and factor effects. For each trained TACBF model, we fix its encoder, classifier, latent samples, predictions, and margins. We evaluate all four combinations of $B \in \{B_{\text{fixed}}, B_{\text{learned}}\}$ and $D \in \{D_{\text{iso}}, D^*_{\text{common}}\}$. Specifically, let $\beta_{\text{common}} := \min\{\beta(B_{\text{fixed}}), \beta(B_{\text{learned}})\}$ and define D^*_{common} by the closed-form allocation in Theorem 4, using β_{common} and that model's classifier subspace. This allocation is identical and feasible for both factors. Tab. 14 reports the allocation main effect, the factor main effect, and their difference-in-differences interaction. Both main effects improve grid AUC for all 35 dataset–seed pairs. Nonzero, dataset-dependent interactions show that the effects are not strictly additive.

Table 14: Fixed-model 2×2 decomposition across seven datasets. Main effects average the two simple effects over the other factor; interaction is the difference between the allocation effects under the learned and fixed factors, i.e., a difference in differences. Entries are five-seed means with paired 95% CIs in unnormalized grid-AUC units.

Dataset	Allocation main effect	Factor main effect	Interaction
CIFAR-10	$+0.9177$ $[0.8607, 0.9746]$	$+0.1772$ $[0.1155, 0.2388]$	-0.1117 $[-0.1877, -0.0357]$
STL-10	$+0.6758$ $[0.6622, 0.6894]$	$+0.1082$ $[0.0806, 0.1358]$	$+0.0744$ $[0.0372, 0.1115]$
EuroSAT RGB	$+0.9561$ $[0.9277, 0.9845]$	$+0.1845$ $[0.1583, 0.2106]$	-0.1304 $[-0.1861, -0.0747]$
CIFAR-100	$+0.1537$ $[0.1527, 0.1546]$	$+0.0673$ $[0.0625, 0.0721]$	$+0.0314$ $[0.0300, 0.0329]$
PCam	$+0.4459$ $[0.4370, 0.4549]$	$+0.0477$ $[0.0403, 0.0550]$	$+0.0862$ $[0.0737, 0.0987]$
UCI HAR	$+0.3337$ $[0.3295, 0.3379]$	$+0.0200$ $[0.0096, 0.0305]$	-0.0023 $[-0.0102, 0.0057]$
Banking77	$+0.0387$ $[0.0381, 0.0392]$	$+0.0247$ $[0.0231, 0.0264]$	$+0.0070$ $[0.0058, 0.0082]$

Task–reconstruction tradeoff. For noise uniformly distributed on the unit sphere in $\mathbb{R}^{(m-q)r}$, the expected squared LS reconstruction error is $\mathbb{E}\|\widehat{s}_C - s\|_2^2 = \operatorname{tr}(G_C^{-1})/((m-q)r)$; for an orthonormal basis $U_\perp$ of $\mathcal{U}^\perp$, the corresponding component on $\mathcal{U}^\perp$ is $\operatorname{tr}(U_\perp^\top G_C^{-1} U_\perp)/((m-q)r)$.

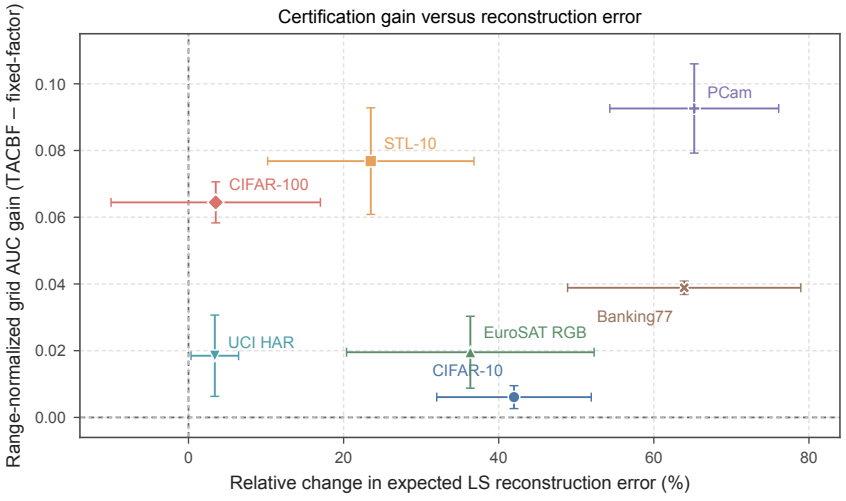


Figure 6: Certification–reconstruction tradeoff at $q = 2$ across seven datasets. Each point compares TACBF with its paired fixed-factor control: the horizontal axis shows the relative change in mean expected squared LS reconstruction error, and the vertical axis shows the range-normalized grid AUC gain. Error bars are paired 95% CIs over five seeds.

Across the seven dataset-level five-seed means, the cross-factor comparison at the same trace gives 3.4%–65.2% higher mean expected squared LS reconstruction error and $+0.0061$–$+0.0926$ higher

range-normalized grid AUC for TACBF. The matched STL-10 $q = 3$ setting reproduces the same qualitative tradeoff.

E.3 Factor geometry and certification gains

Panels (a) and (b) of Fig. 3 quantify how learning the Parseval factor changes exact survivor uncertainty under shared anisotropic scaling. Replacing the fixed factor with the learned factor in the same CIFAR-10 model reduces decision-direction uncertainty while leaving the survivor spectral floor nearly unchanged. Averaging $c_{ab}(C)$ over off-diagonal class pairs for each survivor set, the learned factor lowers the maximum by 14.3% and contracts the maximum-minus-median gap by 88.4%. Panel (c) of Fig. 3 relates the paired reduction in J_{pair} to certification gains across all seven datasets.

F Erasure-Severity Transfer and Relaxation Gaps

F.1 Erasure-severity transfer

The frozen-model evaluation covers all seven datasets at $q = 0, 1, 2, 3$. Fig. 7 reports paired common-δ AUC gains on each held-out split. The mean gain increases from $+0.0040$ to $+0.0723$ between the endpoints. TACBF's common-δ AUC advantage is larger at $q = 3$ than at $q = 0$ for all 35 dataset–seed pairs, although intermediate trends vary by dataset. At $q = 3$, the dataset-level gains range from $+0.0144$ to $+0.1265$.

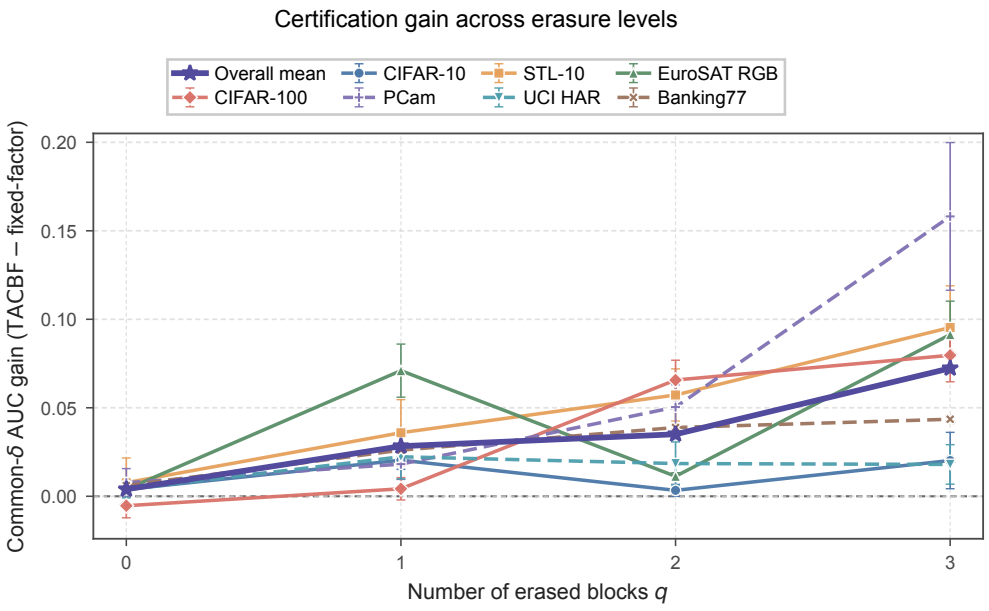


Figure 7: Erasure-severity transfer of models trained at $q = 2$. Curves show dataset-level mean paired common-δ AUC gains when evaluated at $q = 0, 1, 2, 3$; bars are paired 95% CIs over five seeds. The thick curve averages the 35 dataset–seed pairs.

Direct training at $q = 3$. To distinguish direct optimization from frozen-model transfer, we additionally train three matched STL-10 pairs at $q = 3$ using all 56 survivor sets. TACBF achieves mean grid AUC 0.3728, compared with 0.1781 for the fixed-factor control, and outperforms the control in all three paired runs.

F.2 Relaxation gap and class-pair coverage

Here J is the measured spectral task envelope at $q = 2$, and J_{relax} is the lower bound in Theorem 3. Tab. 15 combines the spectral and class-pair relaxation measurements at the final models.

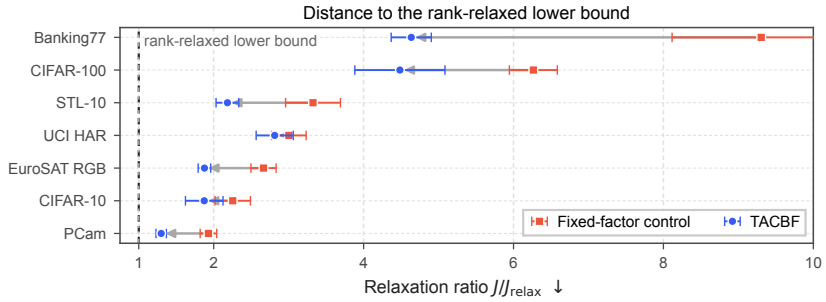


Figure 8: Relaxation ratios at $q = 2$ across seven datasets. Points and arrows show five-seed means with 95% CIs for the fixed-factor control and TACBF. Learning the factor moves every dataset toward the rank-relaxed lower bound $J/J_{\text{relax}} = 1$, while all ratios remain above it.

Table 15: Spectral and class-pair relaxation measurements at $q = 2$ (five-seed means). The first column lists J/J_{relax} for TACBF and the fixed-factor control, respectively; the remaining columns describe TACBF. † marks the midpoint of a certified relative bracket of width at most 0.5%.

Dataset	J/J_{relax} (TACBF/fixed-factor)	J_{pair}/J	$\gamma^{\text{rel}}_{\text{pair}}/J_{\text{relax}}$
CIFAR-10	1.873/2.254	0.910	1.0000
STL-10	2.183/3.325	0.688	0.8302
EuroSAT RGB	1.876/2.665	0.914	1.0000
CIFAR-100	4.484/6.265	0.322	$0.4304^{\dagger}$
PCam	1.299/1.930	1.000	1.0000
UCI HAR	2.814/3.003	0.911	1.0000
Banking77	4.636/9.305	0.419	0.6604

We solve Eq. (57) independently for each final classifier. Learning the block geometry decreases J/J_{relax} on every dataset, while all ratios remain above the rank-relaxed lower bound. The class-pair and spectral relaxations coincide numerically on CIFAR-10, EuroSAT RGB, PCam, and UCI HAR; the class-pair reference is tighter on STL-10, CIFAR-100, and Banking77. The class-pair coverage bounds remain analytically useful but are loose on several higher-dimensional tasks.

CIFAR-100 has 4,950 class-pair constraints in a $p = 99$ task subspace, making the full positive-semidefinite constraint system substantially more expensive than for the other heads. We use constraint generation with full-pair separation and report a certified 0.5% bracket rather than a solver-tolerance point estimate.

Alignment with certification gains. Across the 35 dataset–seed pairs at $q = 2$, the relative reduction $1 - J^{\text{TACBF}}_{\text{pair}}/J^{\text{fixed}}_{\text{pair}}$ is positive in every pair and has Spearman correlation $\rho_s = 0.733$ with the range-normalized grid AUC gain on held-out test sets; leave-one-dataset-out correlations remain 0.640–0.788. The correlations are 0.517 for $1 - J^{\text{TACBF}}/J^{\text{fixed}}$ and 0.096 for $1 - \alpha_q^{\text{TACBF}}/\alpha_q^{\text{fixed}}$, making J_{pair} the criterion most strongly associated with certification gains.