WHY AND WHEN SCORE-FUNCTION LEARNING FAILS UNDER DISCRETE DOWNSTREAM EVALUATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Score-function learning estimates gradients from scalar feedback and thereby enables optimization through nondifferentiable and even black-box evaluators. Although its estimator is unbiased for a smoothed objective, this guarantee does not determine whether the current gradient can be recovered under a finite evaluation budget. We identify two failure mechanisms under discrete downstream evaluation. *Boundary starvation* occurs when sampled candidates rarely cross loss-changing decision interfaces, causing centered group updates to vanish. *Tangential noise* arises because perturbation components tangent to these interfaces contribute variance without changing the mean gradient. Based on this characterization, we introduce two budget-dependent quantities. *Activation evidence* measures whether the available queries are likely to reveal an informative loss comparison, while *gradient resolvability* compares the squared norm of the smoothed gradient with the estimation variance remaining after budgeted averaging. Both quantities can be estimated from pilot groups through a confidence-aware diagnosis. Numerical experiments verify the existence of both failure mechanisms and the effectiveness of the proposed quantities. Code is avaliable at `https://anonymous.4open.science/r/SFBO-70C9/`.

1 INTRODUCTION

Many learning and optimization methods follow the same iterative pattern: generate candidate solutions, evaluate their downstream performance, and use the resulting feedback to improve future candidates. When the downstream evaluator is differentiable, this update can be obtained by ordinary backpropagation. The situation is more difficult when evaluation involves a discrete decision, a combinatorial optimizer, or an external black box, because its output may provide no usable derivative with respect to the candidate (Berthet et al., 2020; Mohamed et al., 2020).

Score-function learning (SFL) offers a general way to optimize such pipelines. It samples candidates from a parameterized distribution, evaluates each candidate using the downstream system, and updates the distribution according to the association between the sampled candidates and their scalar feedback. More precisely, the update combines the observed feedback with the sensitivity of each sample's log probability to the distribution parameters. It therefore does not differentiate through the evaluator itself (Williams, 1992; Mohamed et al., 2020). This makes SFL applicable to nondifferentiable and black-box pipelines, including black-box optimization (BBO) (Wierstra et al., 2014) and decision-focused learning (DFL) (Silvestri et al., 2026).

Under standard conditions, the resulting estimator is unbiased for the gradient of an expected objective (Mohamed et al., 2020). For Gaussian candidate perturbations, this objective is a smoothed version of the downstream loss (Nesterov & Spokoiny, 2017). Population unbiasedness, however, does not imply finite-budget reliability. In practice, downstream evaluations are often costly, so each update must be estimated from a limited number of samples. The estimated updating direction may repeatedly vanish or vary so strongly across sample groups that learning has no stable direction. We study a central question: *why do these failures occur under discrete evaluations, and when can the current smoothed gradient be recovered under a specified query budget?*

The key observation is that a discrete evaluator maps many nearby candidates to the same decision and hence often to the same loss. The continuous candidate space is partitioned into decision regions, and the loss changes only when a candidate crosses an interface between regions with different

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.

feedback values (Berthet et al., 2020; Silvestri et al., 2026). We show that the Gaussian-smoothed gradient is generated by these loss-changing interfaces. This characterization reveals two distinct failure mechanisms.

The first mechanism concerns whether the sampled group contains any useful comparison. The grouped SFL estimator we study centers the observed feedback, so a group in which every candidate receives the same loss produces an exactly zero update. When loss-changing interfaces are rarely crossed, most groups are therefore uninformative. We call this phenomenon *boundary starvation*. The second mechanism concerns the quality of the update after a loss change is observed. Crossing an interface may depend on movement in only a small number of directions, but the score-function update contains the entire perturbation vector. Perturbation components tangent to the interface add estimator variance without changing the mean gradient. We call this effect *tangential noise*. Thus, boundary starvation limits the availability of comparative feedback, while tangential noise reduces the precision of the resulting gradient estimate, as shown in Figure 1.

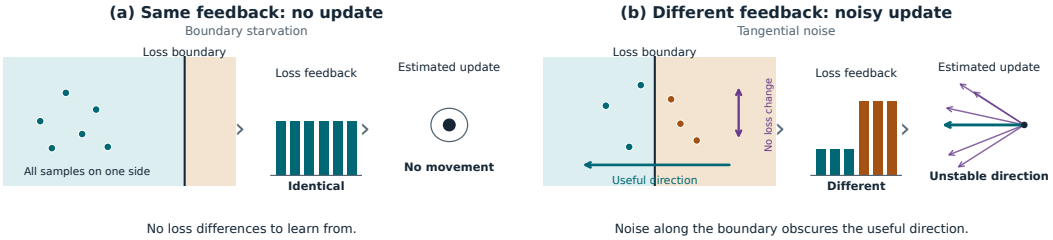


Figure 1: **Two mechanisms that hinder learning from discrete feedback.** (a) Identical losses yield no centered update. (b) Crossing the boundary reveals loss differences, but tangential noise obscures the descent direction (teal), producing variable updates (purple).

These mechanisms lead to two budget-level questions. First, is the available future budget likely to contain any group with distinguished feedback values? Second, if such groups occur, can averaging all estimates available under the budget recover the smoothed gradient with sufficient accuracy? We quantify the first requirement through *activation evidence* and the second through *gradient resolvability*. Activation evidence determines the probability that the budget reveals at least one informative comparison. Gradient resolvability compares the squared norm of the smoothed gradient with the estimation variance remaining after the available groups are averaged. Its reciprocal is exactly the relative mean squared error of the budget-averaged gradient.

Both quantities can be estimated from independent pilot groups using only sampled candidates, scalar feedback, and their score-function updates. We construct confidence bounds that diagnose boundary starvation, insufficient gradient resolution, or gradient recoverability. When the pilot does not provide enough evidence to support any of these conclusions, the diagnostic explicitly abstains. Numerical experiments verify the two failure mechanisms and show that the proposed quantities predict the reliability of independently sampled future gradients in DFL and BBO.

Our contributions are twofold:

1. We explain why grouped SFL becomes unreliable under discrete downstream evaluation. We show that the Gaussian-smoothed gradient is generated by loss-changing decision interfaces and use this representation to characterize boundary starvation and tangential noise. The resulting analysis quantifies both the scarcity of informative groups and the variance accumulated in directions that do not change the mean gradient.

2. We characterize when the current smoothed gradient is recoverable under a specified future query budget. We introduce activation evidence and gradient resolvability, establish the exact connection between resolvability and relative gradient MSE, and estimate both quantities from grouped pilot evaluations using a confidence-aware rule with explicit abstention.

2 PRELIMINARIES AND PROBLEM STATEMENT

The grouped SFL pipeline. Let p_η be a candidate distribution over $\mathbb{R}^D$, parameterized by $\eta \in \mathbb{R}^P$. SFL samples candidates from p_η, evaluates their downstream losses, and uses the resulting feedback to update the distribution parameters. The grouped estimator studied here draws $K \geq 2$ candidates $U_1, \ldots, U_K \overset{\text{i.i.d.}}{\sim} p_\eta$ and obtains scalar feedback $Y_i = L(U_i)$ from an evaluator $L : \mathbb{R}^D \to \mathbb{R}$. These observations are converted into a gradient estimate G_K, giving the pipeline

$$\eta \xrightarrow{p_\eta} \{U_i\}_{i=1}^K \xrightarrow{L} \{Y_i\}_{i=1}^K \xrightarrow{\text{SFL}} G_K \xrightarrow{\text{update}} \eta^+.$$

The expected objective is $J(\eta) = \mathbb{E}_{U\sim p_\eta}[L(U)]$. Define the score $s_\eta(U) = \nabla_\eta \log p_\eta(U)$, which describes how the log density of a sampled candidate changes with the distribution parameters. Under standard regularity conditions, $\nabla_\eta J(\eta) = \mathbb{E}_{U\sim p_\eta}[L(U)s_\eta(U)]$. This identity combines scalar feedback with the score to estimate the objective gradient without differentiating through L (Williams, 1992; Mohamed et al., 2020).

For grouped feedback, define the leave-one-out (LOO) baseline (Kool et al., 2019) $\bar{Y}_{-i} = (K-1)^{-1}\sum_{j\neq i} Y_j$. The corresponding estimator is

$$G_K = \frac{1}{K}\sum_{i=1}^{K}(Y_i - \bar{Y}_{-i})s_\eta(U_i). \tag{1}$$

Because $\bar{Y}_{-i}$ is independent of U_i and $\mathbb{E}[s_\eta(U_i)] = 0$, the baseline preserves unbiasedness: $\mathbb{E}[G_K] = \nabla_\eta J(\eta)$. For loss minimization, the parameters can then be updated as $\eta^+ = \eta - \alpha G_K$, where $\alpha > 0$ is the learning rate.

Discrete downstream evaluation. We consider an evaluator $L(u) = \ell(\mathcal{D}(u))$, where $\mathcal{D} : \mathbb{R}^D \to \mathcal{Z}$ is a discrete decision map and $\ell : \mathcal{Z} \to \mathbb{R}$ is a scalar decision loss. We assume deterministic tie-breaking and write the attainable decision set as $\mathcal{Z} = \{z_1, \ldots, z_J\}$. Each decision z_j defines a region $\mathcal{C}_j = \{u \in \mathbb{R}^D : \mathcal{D}(u) = z_j\}$, within which the feedback is constant: $L(u) = \ell(z_j)$. The evaluator is therefore piecewise constant over the candidate space, so ordinary differentiation inside a decision region provides no useful gradient (Berthet et al., 2020; Silvestri et al., 2026). Different decision regions may share the same loss; our analysis therefore focuses on interfaces separating regions with different feedback values.

Gaussian smoothing. For this piecewise-constant evaluator, Gaussian sampling defines a differentiable expected objective (Nesterov & Spokoiny, 2017; Silvestri et al., 2026). We specialize p_η to an isotropic Gaussian, optimizing its mean $\mu \in \mathbb{R}^D$ while holding the scale $\sigma > 0$ fixed: $U_i = \mu + \sigma\xi_i,\ \xi_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_D)$. The smoothed objective is $F_\sigma(\mu) = \mathbb{E}_\xi[L(\mu + \sigma\xi)]$. Since the Gaussian score is $\nabla_\mu \log p_{\mu,\sigma}(U) = \xi/\sigma$, the score-function identity gives its population gradient $m \coloneqq \nabla F_\sigma(\mu) = \sigma^{-1}\mathbb{E}_\xi[L(\mu + \sigma\xi)\xi]$.

Writing $Y_i = L(\mu + \sigma\xi_i)$ and $\bar{Y} = K^{-1}\sum_{i=1}^K Y_i$, substitution into equation 1 gives

$$G_K = \frac{1}{K\sigma}\sum_{i=1}^{K}(Y_i - \bar{Y}_{-i})\xi_i = \frac{1}{(K-1)\sigma}\sum_{i=1}^{K}(Y_i - \bar{Y})\xi_i,\ \ \mathbb{E}[G_K] = m. \tag{2}$$

Thus, the grouped update combines each candidate's relative feedback with its perturbation direction. If all Y_i are equal, every centered feedback term vanishes and $G_K = 0$.

Problem statement. We study finite-budget recovery of the smoothed gradient at a fixed sampling configuration (μ, σ, K). Let $B \geq K$ denote the number of downstream evaluations allocated to a future gradient update. The budget provides $M_B = \lfloor B/K \rfloor$ independent group estimates, whose average is $\bar{G}_B = M_B^{-1}\sum_{b=1}^{M_B} G_K^{(b)}$. Throughout the paper, B refers to this future update budget; the independent pilot evaluations introduced in Section 4 are accounted for separately. Our goal is to determine whether $\bar{G}_B$ recovers the population gradient m with sufficient accuracy under the specified budget, and whether this can be diagnosed from independent pilot groups at the same sampling configuration before the future budget is used.

3 WHY SFL FAILS UNDER DISCRETE EVALUATION

equation 2 exposes two requirements for finite-budget recovery of m. The losses within a group must first contain an informative comparison, and the associated perturbations must then translate that comparison into a sufficiently stable gradient estimate. Discrete evaluation can violate these requirements through *boundary starvation* and *tangential noise*, respectively. Proof details and technical conditions are deferred to Appendix B.

3.1 WHERE THE SMOOTHED GRADIENT COMES FROM

The evaluator L is constant inside each decision region $\mathcal{C}_j$, yet its Gaussian smoothing can have a nonzero gradient. To identify where this gradient comes from, let $\ell_j = \ell(z_j)$ and denote by $\mathcal{F}_{jk} = \partial\mathcal{C}_j \cap \partial\mathcal{C}_k$ the interface between regions j and k. Let $n_{jk}(u)$ be the unit normal on this interface, oriented from $\mathcal{C}_j$ to $\mathcal{C}_k$, and let $\varphi_{\sigma,D}$ denote the density of $\mathcal{N}(0, \sigma^2 I_D)$. Under the regularity conditions in Appendix B.2, the smoothed gradient admits the following boundary representation.

Proposition 3.1 (Boundary representation)**.** *The gradient of the smoothed objective satisfies*

$$m = \sum_{j<k} (\ell_k - \ell_j) \int_{\mathcal{F}_{jk}} \varphi_{\sigma,D}(u - \mu) n_{jk}(u) \, \mathrm{d}\mathcal{H}^{D-1}(u), \tag{3}$$

where $\mathcal{H}^{D-1}$ is surface measure on each interface.

Proposition 3.1 shows that each interface contributes through three factors: the feedback difference across the interface, its Gaussian weight around the current mean, and its normal direction. An interface contributes nothing when its adjacent regions have the same feedback, even if they correspond to different decisions. Interfaces far from the current perturbation distribution receive little weight, while contributions from different interfaces may also cancel. This representation provides the geometric basis for the two failure modes below: finite samples may rarely encounter relevant feedback changes, and the interface normals determine the directions in which population signal can arise.

3.2 FAILURE MODE I: BOUNDARY STARVATION

We first ask whether one sampled group contains any comparative feedback. Let $Y = L(\mu + \sigma\xi)$ take distinct values $v_1, \ldots, v_R$ with probabilities $\rho_r = \Pr(Y = v_r)$. A group is *active* when its K feedback values are not all equal, and we denote its activation probability by q_K.

Proposition 3.2 (Group activation)**.** *The activation probability is $q_K = 1 - \sum_{r=1}^{R} \rho_r^K$. If $q_K^{\mathcal{D}}$ is the probability that the group contains at least two distinct decisions, then $\Pr(G_K \neq 0) \leq q_K \leq q_K^{\mathcal{D}}$.*

An inactive group necessarily gives $G_K = 0$, whereas an active group may still cancel vectorially. Moreover, distinct decisions need not produce distinct feedback, so decision diversity can overstate the comparison signal available to SFL.

To connect activation with the local geometry of a discrete evaluator and isolate the role of a loss-changing interface, consider the most common binary switching model

$$L(u) = L_0 + \Delta \, \mathbf{1}\{n^\top u < \beta\}, \;\; \|n\| = 1, \; \Delta \neq 0. \tag{4}$$

Here L_0 is the baseline loss, Δ is the loss jump across the interface, n is its unit normal, and β determines the interface location. Suppose $d = n^\top \mu - \beta > 0$. The standardized distance from the Gaussian mean to the feedback change is $a = d/\sigma$, and one perturbation reaches the other side with probability $p = \Phi(-a)$. Here Φ is the standard normal distribution function. A group is active unless all samples fall on the same side, giving $q_K(p) = 1-(1-p)^K - p^K$. If independent groups are sampled until the first active group, the expected number of evaluator calls is $\mathcal{Q}_K(p) = K/q_K(p)$.

Theorem 3.3 (Rare-boundary query complexity)**.** *For every $K \geq 2$ and $p \in (0, \frac{1}{2}]$, $\mathcal{Q}_K(p) \geq \frac{1}{p}$, $\inf \mathcal{Q}_{K_{>2}}(p) = \Theta\left(\frac{1}{p}\right)$ $(p \downarrow 0)$. For Gaussian perturbations, $\frac{1}{p} = \frac{1}{\Phi(-a)} \sim \sqrt{2\pi}\, a e^{a^2/2}$ $(a \to \infty)$.*

Theorem 3.3 shows that observing even one group with comparative feedback requires at least $\Theta(1/p)$ queries in the rare-transition regime, regardless of how the group size is chosen. For Gaussian perturbations, this cost grows exponentially in the squared standardized distance a^2.

Increasing K does not remove the bottleneck under a fixed total budget. When $Kp \to 0$, $q_K(p) = Kp[1 + o(1)]$, while a budget B contains B/K groups. The expected number of active groups is therefore $Bp[1+o(1)]$: the larger activation probability of each group is offset by the smaller number of groups. We call this finite-budget shortage of comparative feedback *boundary starvation*.

3.3 Failure Mode II: Tangential Noise

We next ask whether an active group provides a reliable gradient direction. In the binary model, feedback depends only on the normal coordinate $n^\top \xi$, whereas the score estimator uses the full perturbation vector. Under isotropic Gaussian perturbations, the tangential components carry no population signal but still contribute sampling variance.

Let ϕ denote the standard normal density, and let I_D denote the $D \times D$ identity matrix. The following result makes this normal–tangential decomposition explicit.

Theorem 3.4 (Normal–tangential variance decomposition)**.** *For the binary evaluator equation 4,* $m = -\frac{\Delta}{\sigma}\phi(a)n,\ \mathrm{Cov}(G_K) = v_\parallel nn^\top + v_\perp(I_D - nn^\top)$, *where* $v_\perp = \frac{\Delta^2}{\sigma^2(K-1)}p(1-p), v_\parallel = \frac{\Delta^2}{\sigma^2 K}\left[a\phi(a)(1-2p) + \frac{K}{K-1}p(1-p) - \frac{K-2}{K-1}\phi(a)^2\right]$. *For fixed K, as $a \to \infty$,* $v_\parallel = \frac{\Delta^2}{\sigma^2 K}a^2 p[1+o(1)]$.

Here $v_\parallel$ is the variance along the interface normal, while $v_\perp$ is the variance along each tangential direction. Every tangential direction therefore carries zero mean signal but positive variance $v_\perp$. Although this variance is incurred separately in each direction, it accumulates across dimensions. To isolate this effect, let P_r, $1 \le r \le D$, be a rank-r orthogonal projector satisfying $P_r n = n$. It retains the normal direction, which contains the entire mean gradient, together with $r-1$ tangential directions; $r = D$ recovers the full estimator. The projector is used only for analysis.

Corollary 3.5 (Effective dimension)**.** *The signal-to-noise ratio (SNR) in the retained subspace is*

$$\mathrm{SNR}_r := \frac{\|m\|^2}{\mathbb{E}\|P_r(G_K - m)\|^2} = \frac{\Delta^2\phi(a)^2/\sigma^2}{v_\parallel + (r-1)v_\perp}. \tag{5}$$

For fixed K and r, the rare-boundary limit is $\mathrm{SNR}_r = \frac{K(K-1)a^2 p}{(K-1)a^2 + K(r-1)}[1 + o(1)]$.

The key consequence is that adding tangential directions increases estimation noise without increasing population signal. When the tangential contribution dominates the denominator of equation 5, the SNR decreases approximately as $1/r$. Boundary starvation and tangential noise therefore limit different stages of gradient recovery: the former determines whether comparative feedback is observed, the latter determines whether the resulting estimate is accurate enough to recover the population gradient. We next turn these two requirements into observable budget criteria.

4 When Is the Smoothed Gradient Recoverable?

The geometry of loss-changing interfaces is typically unavailable in realistic evaluators. We therefore characterize finite-budget recovery through group activation and gradient moments, which can be estimated from pilot queries without reconstructing the interfaces.

4.1 Activation Evidence and Gradient Resolvability

Recall that q_K is the probability that one group is active and a budget B contains M_B independent groups. The probability of observing at least one active group is $1-(1-q_K)^{M_B}$. We define the corresponding *activation evidence* as $A_B := -M_B \log(1-q_K)$, so that $\Pr(\text{at least one active group}) = 1-e^{-A_B}$. The transformation preserves the activation probability while providing a convenient budget scale. In particular, when $q_K \ll 1$, $A_B \approx M_B q_K$, the expected number of active groups. Thus A_B measures whether the future budget is likely to reveal any comparative feedback.

Activation alone does not determine gradient accuracy. Let $V_K = \mathbb{E}\|G_K - m\|^2 = \mathrm{tr}(\mathrm{Cov}(G_K))$ denote the total variance of one group estimate. Averaging M_B independent groups reduces this variance to V_K/M_B. For $0 < V_K < \infty$, we therefore define the *gradient resolvability* $\Lambda_B = M_B\|m\|^2/V_K$. It compares the squared norm of the population gradient with the variance remaining after the future groups are averaged, with $\Lambda_B = 0$ when $m = 0$. Equivalently, Λ_B is M_B times the one-group SNR. In the binary model of Section 3.3, $\Lambda_B = M_B\,\mathrm{SNR}_D$.

Theorem 4.1 (Budgeted gradient resolvability). *If $m \neq 0$, then for $M_B \geq 1$ independent groups, $\mathbb{E}\|\bar{G}_B - m\|^2/\|m\|^2 = 1/\Lambda_B$. Consequently, for any tolerated relative root mean squared error $\varepsilon > 0$, $\mathbb{E}\|\bar{G}_B - m\|^2 \leq \varepsilon^2\|m\|^2$ if and only if $\Lambda_B \geq \varepsilon^{-2}$.*

When $m \neq 0$, $1/\Lambda_B$ is exactly the relative mean squared error of the budget-averaged gradient (Bollapragada et al., 2018). Although Λ_B directly characterizes recovery accuracy for a nonzero gradient, it does not replace activation evidence. The latter identifies the distinct event that a finite budget may contain no comparative feedback at all and remains meaningful when the population gradient is zero.

The two quantities are nevertheless linked. When active groups are rare, the population mean must be carried entirely by the estimates from these active groups, which necessarily increases variance.

Proposition 4.2 (Activation limits resolvability). *If $m \neq 0$, G_K is square integrable, and $q_K < 1$, then $\Lambda_B \leq M_B q_K/(1 - q_K)$.*

For $q_K \ll 1$, this gives $\Lambda_B \lesssim M_B q_K$: severe starvation directly prevents accurate gradient recovery. The converse does not hold. For example, in the binary model, increasing the number of tangential directions leaves q_K unchanged while increasing estimation variance. Activation evidence therefore asks whether comparative feedback is observed, whereas gradient resolvability asks whether the resulting vector estimate is accurate enough to recover m.

4.2 Estimating Both Quantities from a Pilot

The population quantities q_K, $\|m\|^2$, and V_K are unknown. We estimate them from $M_0 \geq 2$ independent pilot groups drawn at the same fixed configuration (μ, σ, K) as the future groups. Let I_b indicate whether pilot group b is active. Then $\widehat{q}_K = M_0^{-1} \sum_{b=1}^{M_0} I_b$.

Let $G^{(1)}, \ldots, G^{(M_0)}$ denote the pilot group estimates, and define their first and second empirical moments by $\bar{G}_0 = M_0^{-1} \sum_{b=1}^{M_0} G^{(b)}$ and $\widehat{T} = M_0^{-1} \sum_{b=1}^{M_0} \|G^{(b)}\|^2$. For two independent groups, $\mathbb{E}\langle G^{(b)}, G^{(c)}\rangle = \|m\|^2$ whenever $b \neq c$.

This allows the squared gradient norm $S \coloneqq \|m\|^2$ to be estimated without a reference gradient by $\widehat{S} = \{M_0(M_0-1)\}^{-1} \sum_{b \neq c} \langle G^{(b)}, G^{(c)}\rangle$, equivalently $\widehat{S} = (M_0\|\bar{G}_0\|^2 - \widehat{T})/(M_0 - 1)$ (Hoeffding, 1948). The corresponding variance estimator is $\widehat{V} = (M_0 - 1)^{-1} \sum_{b=1}^{M_0} \|G^{(b)} - \bar{G}_0\|^2 = \widehat{T} - \widehat{S}$.

For a future budget B, substituting the pilot estimates into the population definitions gives $\widehat{A}_B = -M_B \log(1 - \widehat{q}_K)$ when $\widehat{q}_K < 1$, and $\widehat{\Lambda}_B = M_B \widehat{S}/\widehat{V}$ when $\widehat{S} > 0$ and $\widehat{V} > 0$. We set $\widehat{A}_B = +\infty$ when $\widehat{q}_K = 1$. When $\widehat{\Lambda}_B$ is undefined, no artificial value is assigned; the categorical diagnosis instead uses the margin $\widehat{D}_B$ and its confidence bounds introduced in Section 4.3.

Proposition 4.3 (Pilot estimation). *If the pilot groups are independent and $\mathbb{E}\|G_K\|^2 < \infty$, then $\widehat{q}_K$, $\widehat{S}$, and $\widehat{V}$ are unbiased for q_K, $\|m\|^2$, and V_K, respectively, and converge almost surely to these quantities as $M_0 \to \infty$. If additionally $q_K < 1$, $m \neq 0$, and $0 < V_K < \infty$, then $\widehat{A}_B \to A_B$ and $\widehat{\Lambda}_B \to \Lambda_B$ almost surely for every fixed B.*

Because $\widehat{S}$ is unconstrained, it may be negative for a finite pilot. Confidence bounds use this unmodified estimate; a negative point estimate does not by itself provide evidence that $\|m\|^2$ is positive.

4.3 Confidence-Aware Diagnosis

The required activation and recovery levels depend on how the future gradient will be used. Let $1 - \delta_A$ be the required probability of observing at least one active group, and let ε be the tolerated relative root mean squared error. The activation and recovery conditions above give the corresponding thresholds $A_{\mathrm{req}} = -\log \delta_A$ and $\Lambda_{\mathrm{req}} = \varepsilon^{-2}$.

Let $\mathrm{LCB}(\cdot)$ and $\mathrm{UCB}(\cdot)$ denote lower and upper confidence bounds obtained from the pilot. Their construction is given in Appendix C. Activation bounds use exact binomial intervals (Clopper & Pearson, 1934), while moment bounds use a pointwise asymptotic multiplier approximation. For inference, define the threshold-equivalent margin $D_B = M_B S - \Lambda_{\mathrm{req}} V_K$ and $\widehat{D}_B = M_B \widehat{S} - \Lambda_{\mathrm{req}} \widehat{V}$.

When $V_K > 0$, $D_B \geq 0$ if and only if $\Lambda_B \geq \Lambda_{\text{req}}$. The margin avoids forming a ratio near $V_K = 0$ and admits confidence bounds directly from $\widehat{S}$ and $\widehat{V}$. Because relative gradient recovery concerns a nonzero gradient, the gradient-resolvable diagnosis below additionally requires evidence that $S = \|m\|^2$ is positive. This yields four diagnostic outcomes for a future budget B:

1. STARVED if $\mathrm{UCB}(A_B) \leq A_{\text{req}}$. The budget misses the prescribed activation requirement.
2. LOW-RESOLUTION if $\mathrm{LCB}(A_B) > A_{\text{req}}$ and $\mathrm{UCB}(D_B) < 0$. Comparative feedback is likely to occur, but the budget does not support recovery at the required accuracy.
3. GRADIENT-RESOLVABLE if $\mathrm{LCB}(A_B) > A_{\text{req}}$, $\mathrm{LCB}(S) > 0$, and $\mathrm{LCB}(D_B) \geq 0$. The budget supports recovery of a nonzero smoothed gradient at the required accuracy.
4. UNRESOLVED otherwise. The pilot does not support a definite conclusion.

A low-resolution case with $\mathrm{LCB}(S) > 0$ receives the more specific label NOISE-DOMINATED. When it is defined, $\widehat{\Lambda}_B$ provides a continuous estimate of gradient resolvability, whereas categorical diagnosis follows the confidence-bound rules above. All diagnoses are local to the current (μ, σ, K) and the specified future budget B.

5 EXPERIMENTS

We evaluate whether the two failure mechanisms occur in discrete evaluators and whether pilot statistics predict future gradient quality. Controlled shortest-path experiments isolate boundary starvation and tangential noise; SEMO Energy-Knapsack and UCI HAR feature selection test prospective diagnosis across sampling states and query budgets. Complete configurations and additional results are reported in Appendix D.

Evaluation protocol. Each configuration uses independent pilot, reference, and future samples. The pilot estimates two quantities: *Activation evidence* and *Gradient resolvability*. The reference sample provides a high accuracy population-gradient target, and the future sample is divided into non-overlapping budget-sized windows. Continuous calibration compares the pilot prediction with the relative MSE observed across these future windows.

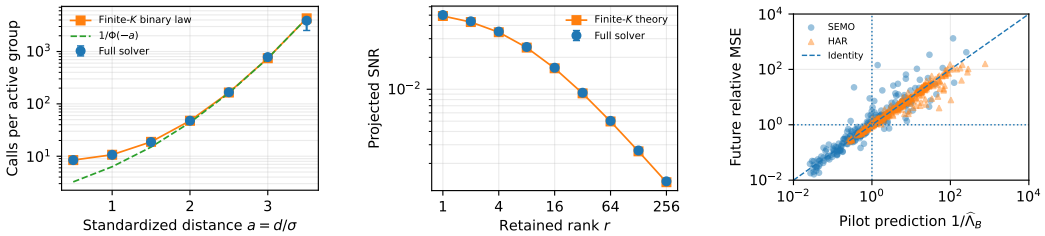
The same quantities define three operational states: STARVED, LOW-RESOLUTION and GRADIENT-RESOLVABLE. Future states are determined entirely from the independent reference and future samples. The pilot may instead abstain when its confidence bounds do not establish a state. We report coverage, selective accuracy among issued diagnoses, and class-wise recall. The primary experiments use a 95% activation requirement, unit relative-MSE threshold, and 90% pointwise confidence level. Appendix D.1 gives the estimators, future labels, and confidence construction.

5.1 VALIDATING THE TWO FAILURE MECHANISMS

We first construct isolated switching interfaces between optimal paths on a shortest-path problem with 16×16 grid ($D = 256$); every perturbation is evaluated by the full dynamic programming solver. To test boundary starvation, we vary the standardized boundary distance $a = d/\sigma \in \{0.5, 1, \ldots, 3.5\}$ and the group size $K \in \{2, 4, 8, 16, 32, 64\}$. Empirical activation probabilities match the finite-K prediction with mean absolute error 2.83×10^{-3}. At $a = 2.5$ and $K = 8$, only 4.58%–5.08% of groups are active, and the mean cost of obtaining one active group is 157.5–174.8 solver calls, close to the predicted 164.6. At $a = 3.5$, this cost rises to approximately 3.88×10^3 calls (Figure 2(a)). Thus, increasing the distance from a loss-changing boundary directly removes comparative feedback from a finite query budget.

To isolate tangential noise, we retain the same boundary normal and increase the number of sampled directions from $r = 1$ to 256 by adding coordinates tangent to the interface. These coordinates leave both activation and the population gradient unchanged. Nevertheless, at $a = 2.5$ and $K = 8$, the empirical SNR falls from 0.05014 to 0.001355, closely matching the predicted change from 0.04909 to 0.001330 (Figure 2(b)). The normal mean, normal variance, and tangential variance agree with theory within 3.65%, 4.02%, and 3.66% mean absolute relative error, respectively. Hence, tangential directions accumulate variance without changing the mean gradient, even when informative comparisons are present.

Why use observable group statistics? Two additional controls examine whether the relevant geometry can instead be recovered from a small pilot. In the first, a rare boundary with a large loss jump is absent from 17 of 32 pilots, close to its predicted miss probability of 0.599. Projecting onto pilot-observed directions reduces short-budget MSE from 9.043 to 1.332 at $K = 4$, but creates persistent bias: at $K = 4096$, the full estimator reaches 0.00781 while the projected estimator remains at 0.0353. In the second, a curved boundary and its nearest tangent plane have the same standardized distance, yet at high curvature the plane predicts activation 0.4249 while the exact and empirical probabilities are 0.03624 and 0.03606. These tests show that a finite pilot can miss consequential interfaces and that distance alone does not encode curvature at the smoothing scale. This motivates estimating activation and gradient moments directly; complete stress-test curves appear in Appendix D.



(a) Boundary starvation **(b)** Tangential noise **(c)** Future-gradient prediction

Figure 2: **Mechanism validation and prospective prediction. (a)** The calls required to obtain an active group grow rapidly with standardized boundary distance. **(b)** Adding tangent directions leaves the signal unchanged but reduces gradient SNR. **(c)** Pilot-predicted relative MSE tracks the error of an independent future gradient on SEMO and HAR; the diagonal denotes exact calibration.

5.2 Predicting Future Gradient Quality

SEMO (Ifrim et al., 2012) evaluates a learned cost predictor through an exact Energy-Knapsack solver (Mandi et al., 2020) with 48 items and 12 fixed validation contexts. We evaluate five checkpoints from each of three training trajectories at seven smoothing scales and future budgets $B \in \{128, 512, 2048\}$, with $K = 8$. This produces 315 configurations spanning different sampling scales, optimization states, and budgets. HAR (Anguita et al., 2013) provides a structurally different test: Gaussian perturbations are decoded into subsets of 30 among 561 features and evaluated by a fixed probabilistic 7-nearest-neighbor validation objective. Four smoothing scales, $K \in \{8, 16\}$, four future budgets, five checkpoints, and three trajectories give 480 configurations. Each cell uses 256 pilot groups, 1024 reference groups, and 4096 future groups; pilot statistics are computed before examining the future outcomes and are reused when evaluating candidate budgets.

Activation evidence accurately predicts whether future groups contain loss variation. Pilot and future activation probabilities have MAE 0.0118 on SEMO and 0.00449 on HAR; their log activation evidences have Pearson/Spearman correlations $0.994/0.994$ and $0.991/0.994$, respectively. Gradient resolvability also predicts the magnitude of future error. As shown in Figure 2(c) and Table 1, the log correlations between predicted and observed relative MSE are $0.923/0.945$ on SEMO and $0.976/0.981$ on HAR. Moreover, 79.9% and 94.9% of future errors fall within a factor of two of the pilot prediction.

The confidence-aware rule converts these continuous predictions into gradient states. Of the 315 SEMO configurations, 270 have a defined future state: 34 are starved, 118 low-resolution, and 118 gradient-resolvable. All 480 HAR configurations have defined states, distributed as 68, 321, and 91. The pilot issues diagnoses for 65.9% of resolved SEMO configurations and 88.5% of HAR configurations, with all issued diagnoses agreeing with the independent future state. Recall remains class-dependent, especially for the low-resolution SEMO regime; the remaining cases receive the UNRESOLVED outcome (Table 1). The reference and future labels themselves agree on 98.5% of resolved SEMO configurations and 99.2% of HAR configurations.

Finally, the observed regimes change systematically with available information. Increasing the future budget moves SEMO from $21/56/19$ starved/low-resolution/resolvable configurations at $B = 128$ to $5/22/59$ at $B = 2048$. On HAR, no configuration is resolvable at $B \leq 512$, whereas

Table 1: **Pilot prediction of independent future outcomes.** Correlations are Pearson/Spearman on the log scale. Recall is ordered as STARVED/LOW-RESOLUTION/GRADIENT-RESOLVABLE.

		Continuous prediction		State diagnosis			
Data	Configs.	RelMSE corr.	Within 2×	Coverage	Sel. acc.	Recall S/L/G	Target agree
SEMO	315	0.923/0.945	79.9%	65.9%	100.0%	82.4/56.8/70.3	98.5%
HAR	480	0.976/0.981	94.9%	88.5%	100.0%	86.8/93.8/71.4	99.2%

68 are resolvable at $B = 4096$. Increasing the pilot from 32 to 1024 groups raises diagnostic coverage from 17.4% to 84.8% on SEMO and from 50.8% to 96.7% on HAR, while selective accuracy remains at least 99.4% and 99.6%, respectively. These results connect the two mechanisms to the operational question: the budget must first expose loss differences and then average enough active but noisy estimates to resolve the smoothed gradient.

6 CONCLUSION AND LIMITATIONS

This work explains why grouped SFL fails under finite query budgets through boundary starvation and tangential noise, and characterizes smoothed-gradient recoverability using pilot-estimable activation evidence and gradient resolvability with confidence-aware abstention.

Our analysis is local to a fixed sampling configuration (μ, σ, K) and concerns recovery of the smoothed gradient; it does not guarantee global convergence, improvement in the original unsmoothed objective, or task-level optimality. The diagnosis requires independent pilot groups from the same sampling state and may abstain frequently when $\|m\|^2$ is small or the population quantities lie near the diagnostic thresholds; reusing a pilot after the sampling distribution is updated is not covered by the present theory. Finally, the analysis assumes Gaussian perturbations and deterministic piecewise constant feedback, while other advanced methods require separate treatment.

AI USE STATEMENT

Claude Opus4.7 is used to assist with the experiments coding. The authors take responsibility for the final manuscript.

REFERENCES

Philipp Andelfinger and Wentong Cai. Dimensional peeking for low-variance gradients in zeroth-order discrete optimization via simulation. In *Proceedings of the 40th ACM SIGSIM International Conference on Principles of Advanced Discrete Simulation*, pp. 78–87. Association for Computing Machinery, 2026. doi: 10.1145/3806789.3810252. URL https://doi.org/10.1145/3806789.3810252.

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 European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, pp. 437–442, 2013. URL https://www.esann.org/sites/default/files/proceedings/legacy/es2013-84.pdf.

Wyame Benslimane, Tinghan Ye, Pascal Van Hentenryck, and Paul Grigas. Decision-focused on-policy learning for contextual linear optimization with partial feedback, 2026. URL https://arxiv.org/abs/2606.01081.

Quentin Berthet, Mathieu Blondel, Olivier Teboul, Marco Cuturi, Jean-Philippe Vert, and Francis Bach. Learning with differentiable perturbed optimizers. In *Advances in Neural Information Processing Systems*, volume 33, pp. 9508–9519. Curran Associates, Inc., 2020. URL https://papers.nips.cc/paper/2020/hash/6bb56208f672af0dd65451f869fedfd9-Abstract.html.

Raghu Bollapragada, Richard Byrd, and Jorge Nocedal. Adaptive sampling strategies for stochastic optimization. *SIAM Journal on Optimization*, 28(4):3312–3343, 2018. doi: 10.1137/17M1154679. URL https://doi.org/10.1137/17M1154679.

Aymeric Capitaine, Maxime Haddouche, Eric Moulines, Michael Jordan, Etienne Boursier, and Alain Oliviero Durmus. Online decision-focused learning. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/4e14a2a225d5d186b7eb1ecc7605bc2d-Abstract-Conference.html.

C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934. doi: 10.1093/biomet/26.4.404. URL https://doi.org/10.1093/biomet/26.4.404.

Adam N. Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 68 (1):9–26, 2022. doi: 10.1287/mnsc.2020.3922. URL https://doi.org/10.1287/mnsc.2020.3922.

Haoyu Geng, Hang Ruan, Runzhong Wang, Yang Li, Yang Wang, Lei Chen, and Junchi Yan. Benchmarking PtO and PnO methods in the predictive combinatorial optimization regime. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/796076672b00f54fb01d05a2e5fde363-Abstract-Datasets_and_Benchmarks_Track.html. Datasets and Benchmarks Track.

Vishal Gupta and Michael Huang. Decision-focused learning with directional gradients. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/907a9fb75a408f6c3a2ae1bf84c39e44-Abstract-Conference.html.

Wassily Hoeffding. A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, 19(3):293–325, 1948. doi: 10.1214/aoms/1177730196. URL https://doi.org/10.1214/aoms/1177730196.

Georgiana Ifrim, Barry O'Sullivan, and Helmut Simonis. Properties of energy-price forecasts for scheduling. In Michela Milano (ed.), *Principles and Practice of Constraint Programming*, volume 7514 of *Lecture Notes in Computer Science*, pp. 957–972. Springer, 2012. doi: 10.1007/978-3-642-33558-7_68. URL https://doi.org/10.1007/978-3-642-33558-7_68.

Lingkai Kong, Wenhao Mu, Jiaming Cui, Yuchen Zhuang, B. Aditya Prakash, Bo Dai, and Chao Zhang. DF^2: Distribution-free decision-focused learning. In *Proceedings of the Forty-first Conference on Uncertainty in Artificial Intelligence*, volume 286 of *Proceedings of Machine Learning Research*, pp. 2269–2290. PMLR, 2025. URL https://proceedings.mlr.press/v286/kong25a.html.

Wouter Kool, Herke van Hoof, and Max Welling. Buy 4 REINFORCE samples, get a baseline for free! In *Deep Reinforcement Learning Meets Structured Prediction Workshop at ICLR*, 2019. URL https://openreview.net/forum?id=r1lgTGL5DE.

Jayanta Mandi, Emir Demirović, Peter J. Stuckey, and Tias Guns. Smart predict-and-optimize for hard combinatorial optimization problems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 1603–1610, 2020. doi: 10.1609/aaai.v34i02.5521. URL https://ojs.aaai.org/index.php/AAAI/article/view/5521.

Jayanta Mandi, Víctor Bucarey, Maxime Mulamba Ke Tchomba, and Tias Guns. Decision-focused learning: Through the lens of learning to rank. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 14935–14947. PMLR, 2022. URL https://proceedings.mlr.press/v162/mandi22a.html.

Jayanta Mandi, James Kotary, Senne Berden, Maxime Mulamba, Víctor Bucarey, Tias Guns, and Ferdinando Fioretto. Decision-focused learning: Foundations, state of the art, benchmark and future opportunities. *Journal of Artificial Intelligence Research*, 80:1623–1701, 2024. doi: 10.1613/jair.1.15320. URL https://www.jair.org/index.php/jair/article/view/15320.

Jayanta Mandi, Marianne Defresne, Senne Berden, and Tias Guns. Feasibility-aware decision-focused learning for predicting parameters in the constraints. In *Advances in Neural Information Processing Systems*, volume 38, 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/ef158a0d4f5364741a571b8d1d44fb1b-Abstract-Conference.html.

Shakir Mohamed, Mihaela Rosca, Michael Figurnov, and Andriy Mnih. Monte Carlo gradient estimation in machine learning. *Journal of Machine Learning Research*, 21(132):1–62, 2020. URL https://jmlr.org/papers/v21/19-346.html.

Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2):527–566, 2017. doi: 10.1007/s10208-015-9296-2. URL https://doi.org/10.1007/s10208-015-9296-2.

Felix Petersen, Christian Borgelt, Aashwin Ananda Mishra, and Stefano Ermon. Generalizing stochastic smoothing for differentiation and gradient estimation. In *International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=okzQ1x71pS.

Utsav Sadana, Abhilash Chenreddy, Erick Delage, Alexandre Forel, Emma Frejinger, and Thibaut Vidal. A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*, 320(2):271–289, 2025. doi: 10.1016/j.ejor.2024.03.020. URL https://doi.org/10.1016/j.ejor.2024.03.020.

Mattia Silvestri, Senne Berden, Gaetano Signorelli, Ali İrfan Mahmutoğulları, Jayanta Mandi, Brandon Amos, Tias Guns, and Michele Lombardi. Score function gradient estimation to widen the applicability of decision-focused learning. *Journal of Artificial Intelligence Research*, 85:10:1–10:34, 2026. doi: 10.1613/jair.1.19498. URL https://www.jair.org/index.php/jair/article/view/19498.

George Tucker, Andriy Mnih, Chris J. Maddison, John Lawson, and Jascha Sohl-Dickstein. REBAR: Low-variance, unbiased gradient estimates for discrete latent variable models. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/ebd6d2f5d60ff9afaeda1a81fc53e2d0-Abstract.html.

Marin Vlastelica, Anselm Paulus, Vít Musil, Georg Martius, and Michal Rolínek. Differentiation of blackbox combinatorial solvers. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=BkevoJSYPB.

Mingrui Wang, Prakash Chakraborty, and Uday V. Shanbhag. Complexity guarantees for zeroth-order methods via exponentially-shifted Gaussian smoothing: Mitigating dimension-dependence and incorporating decision-dependence, 2026. URL https://arxiv.org/abs/2604.15525.

Daan Wierstra, Tom Schaul, Tobias Glasmachers, Yi Sun, Jan Peters, and Jürgen Schmidhuber. Natural evolution strategies. *Journal of Machine Learning Research*, 15(27):949–980, 2014. URL https://jmlr.org/papers/v15/wierstra14a.html.

Klas Wijk, Ricardo Vinuesa, and Hossein Azizpour. Revisiting score function estimators for k-subset sampling. In *ICML 2024 Workshop on Differentiable Almost Everything: Differentiable Relaxations, Algorithms, Operators, and Simulators*, 2024. URL https://openreview.net/forum?id=opr1ijTcj7.

Bryan Wilder, Bistra Dilkina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1658–1665, 2019. doi: 10.1609/aaai.v33i01.33011658. URL https://ojs.aaai.org/index.php/AAAI/article/view/3982.

Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8:229–256, 1992. doi: 10.1007/BF00992696. URL https://doi.org/10.1007/BF00992696.

A Related Work

Score-function estimation and stochastic smoothing. Score-function (likelihood-ratio) estimators provide gradients of expectations without differentiating through the sampled computation, with REINFORCE as a canonical instance (Williams, 1992; Mohamed et al., 2020). When only function values are available, randomized zeroth-order methods instead differentiate a smoothed objective by perturbing its input, with Gaussian smoothing providing a standard construction (Nesterov & Spokoiny, 2017). A central difficulty in both families is variance: classical work develops control variates and related constructions that preserve unbiasedness while reducing Monte Carlo noise (Mohamed et al., 2020; Tucker et al., 2017). Recent work has revisited these ideas for increasingly structured and nonsmooth settings. For discrete k-subset sampling, Wijk et al. (2024) derive an exact score function and combine it with control variates; Petersen et al. (2026) develop a more general stochastic-smoothing framework for nondifferentiable black-box operators together with several variance-reduction strategies. In simulation optimization, Andelfinger & Cai (2026) exploit shared control-flow structure to reduce the variance of zeroth-order gradients, while Wang et al. (2026) study a modified Gaussian-smoothing estimator with improved dimension dependence. Most directly related to our setting, Silvestri et al. (2026) combine stochastic smoothing with score-function gradient estimation to make decision-focused learning applicable to a broad class of optimization problems. These works primarily address how to construct a usable gradient estimator, reduce its variance, or broaden the class of problems to which smoothing can be applied. Our question is complementary: *given* a score-function estimator and a finite solver-query budget, when is its population direction statistically recoverable at the current iterate? For losses induced by a discrete optimizer, the answer depends on whether queries encounter loss-changing decision interfaces and on how the estimator's variance is distributed relative to the signal-bearing directions. We therefore study finite-budget recoverability and prospective reliability certification, rather than proposing another score-function or smoothing estimator.

Learning through discrete optimization. Decision-focused learning (DFL), or end-to-end predict-then-optimize learning, trains predictors according to the quality of the downstream decisions rather than predictive error alone (Wilder et al., 2019; Elmachtoub & Grigas, 2022). Its central technical obstacle is that discrete optimization maps are typically piecewise constant or otherwise unsuitable for ordinary backpropagation. Existing approaches address this obstacle through continuous relaxations, black-box surrogate differentiation, random perturbations of the optimizer, or surrogate ranking losses (Wilder et al., 2019; Vlastelica et al., 2020; Berthet et al., 2020; Mandi et al., 2022). The resulting literature is now broad (Sadana et al., 2025); recent surveys and benchmarks organize both gradient-based and gradient-free DFL methods and document their behavior across predictive combinatorial optimization tasks (Mandi et al., 2024; Geng et al., 2024). Work from 2024–2026 has also expanded the scope of DFL in several orthogonal directions: directional-gradient surrogates (Gupta & Huang, 2024), distribution-free learning (Kong et al., 2025), prediction of uncertain constraint parameters with explicit feasibility control (Mandi et al., 2025), online and nonstationary decision-focused learning (Capitaine et al., 2026), and on-policy learning under partial feedback, where a score-function term is combined with a decision-focused plug-in gradient (Benslimane et al., 2026). Our work does not introduce a new differentiable optimizer, decision-focused loss, or training objective. Instead, it studies a statistical question that arises after such stochastic differentiation mechanisms have already been chosen: whether the current smoothed gradient can be trusted under a specified number of black-box solver calls. This leads to separate notions of activation evidence and directional resolvability, and to an explicit unresolved state when the available queries cannot certify either conclusion. In this sense, prior DFL work largely asks *how to obtain or optimize with a gradient through a discrete decision layer*; we ask *when a finite-budget score-function gradient through that layer is actually informative enough to use*.

B Proofs and Technical Details

This appendix restates the definitions and results needed by the proofs. It can therefore be read independently of the main text. Throughout, $\phi(t) = (2\pi)^{-1/2}e^{-t^2/2}$ and $\Phi(t) = \int_{-\infty}^{t}\phi(s)\,\mathrm{d}s$ denote the standard normal density and distribution function. We write

$$\varphi_{\sigma,D}(v) = (2\pi\sigma^2)^{-D/2}\exp\left(-\frac{\|v\|^2}{2\sigma^2}\right)$$

for the density of $\mathcal{N}(0,\sigma^2 I_D)$.

B.1 Setup and Estimator Identities

Let p_η be a distribution over candidates $U \in \mathbb{R}^D$, and let $L(U)$ be the scalar feedback returned by the evaluator. Define

$$J(\eta) = \mathbb{E}_{U\sim p_\eta}[L(U)], \;\; s_\eta(U) = \nabla_\eta \log p_\eta(U).$$

Whenever differentiation and integration can be interchanged,

$$\nabla_\eta J(\eta) = \mathbb{E}[L(U)s_\eta(U)]. \tag{6}$$

For $K \geq 2$, draw $U_1,\ldots,U_K \overset{\text{i.i.d.}}{\sim} p_\eta$, set $Y_i = L(U_i)$, and define

$$\bar{Y}_{-i} = \frac{1}{K-1}\sum_{j\neq i} Y_j.$$

The grouped score-function estimator is

$$G_K = \frac{1}{K}\sum_{i=1}^{K}(Y_i - \bar{Y}_{-i})s_\eta(U_i). \tag{7}$$

Unbiasedness. Assume that the expectations in equation 6–equation 7 are finite. Then

$$\mathbb{E}[G_K] = \nabla_\eta J(\eta). \tag{8}$$

Proof. The score has mean zero:

$$\mathbb{E}[s_\eta(U)] = \int \nabla_\eta p_\eta(u)\,\mathrm{d}u = \nabla_\eta \int p_\eta(u)\,\mathrm{d}u = 0.$$

For each i, $\bar{Y}_{-i}$ depends only on the other candidates and is independent of U_i. Hence

$$\mathbb{E}[\bar{Y}_{-i}s_\eta(U_i)] = \mathbb{E}[\bar{Y}_{-i}]\mathbb{E}[s_\eta(U_i)] = 0.$$

Using identical distributions across i and equation 6,

$$\mathbb{E}[G_K] = \frac{1}{K}\sum_{i=1}^{K}\mathbb{E}[Y_i s_\eta(U_i)] = \mathbb{E}[L(U)s_\eta(U)] = \nabla_\eta J(\eta). \tag{9}$$

□

Gaussian specialization. We use

$$U_i = \mu + \sigma\xi_i, \;\; \xi_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_D), \;\; Y_i = L(\mu + \sigma\xi_i). \tag{10}$$

The smoothed objective and its population gradient are

$$F_\sigma(\mu) = \mathbb{E}[L(\mu+\sigma\xi)], \;\; m = \nabla F_\sigma(\mu) = \frac{1}{\sigma}\mathbb{E}[L(\mu+\sigma\xi)\xi].$$

Let $\bar{Y} = K^{-1}\sum_i Y_i$. Since $Y_i - \bar{Y}_{-i} = K(Y_i - \bar{Y})/(K-1)$, the grouped estimator becomes

$$G_K = \frac{1}{K\sigma}\sum_{i=1}^{K}(Y_i - \bar{Y}_{-i})\xi_i = \frac{1}{(K-1)\sigma}\sum_{i=1}^{K}(Y_i - \bar{Y})\xi_i, \;\; \mathbb{E}[G_K] = m. \tag{11}$$

A future evaluation budget B contains $M_B = \lfloor B/K \rfloor$ complete groups. Their average is

$$\bar{G}_B = \frac{1}{M_B}\sum_{b=1}^{M_B} G_K^{(b)}, \tag{12}$$

where the group estimates are independent copies of G_K. The pilot data used for diagnosis are independent of these future groups.

B.2 Boundary Representation

Suppose a discrete decision map partitions $\mathbb{R}^D$ into regions $\{\mathcal{C}_j\}$. Region $\mathcal{C}_j$ returns the constant loss ℓ_j. For two adjacent regions, write

$$\mathcal{F}_{jk} = \partial\mathcal{C}_j \cap \partial\mathcal{C}_k$$

and orient the unit normal $n_{jk}(u)$ from $\mathcal{C}_j$ toward $\mathcal{C}_k$.

Assumption B.1 (Regular decision partition)**.** The regions cover $\mathbb{R}^D$ up to a Lebesgue-null set and form a locally finite partition. Pairwise interfaces are piecewise C^1, intersections of three or more regions have zero $(D-1)$-dimensional Hausdorff measure, and the Gaussian-weighted volume and surface integrals below are absolutely summable.

Boundary representation. Under Assumption B.1,

$$m = \sum_{j<k}(\ell_k - \ell_j)\int_{\mathcal{F}_{jk}} \varphi_{\sigma,D}(u-\mu)n_{jk}(u)\,\mathrm{d}\mathcal{H}^{D-1}(u). \tag{13}$$

Proof. We first consider a finite partition. Since $L(u) = \ell_j$ on $\mathcal{C}_j$,

$$F_\sigma(\mu) = \sum_{j=1}^{J}\ell_j\int_{\mathcal{C}_j}\varphi_{\sigma,D}(u-\mu)\,\mathrm{d}u.$$

The Gaussian density and its first derivatives are integrable, so

$$\nabla F_\sigma(\mu) = \sum_{j=1}^{J}\ell_j\int_{\mathcal{C}_j}\nabla_\mu\varphi_{\sigma,D}(u-\mu)\,\mathrm{d}u.$$

Using $\nabla_\mu\varphi_{\sigma,D}(u-\mu) = -\nabla_u\varphi_{\sigma,D}(u-\mu)$,

$$\nabla F_\sigma(\mu) = -\sum_{j=1}^{J}\ell_j\int_{\mathcal{C}_j}\nabla_u\varphi_{\sigma,D}(u-\mu)\,\mathrm{d}u.$$

Truncate each region by the ball $B_R = \{u : \|u\| \le R\}$ and apply the divergence theorem componentwise. The flux through the truncating sphere vanishes as $R \to \infty$ because of the Gaussian tail. If $\nu_j(u)$ is the outward unit normal of $\mathcal{C}_j$, then

$$\nabla F_\sigma(\mu) = -\sum_{j=1}^{J}\ell_j\int_{\partial\mathcal{C}_j}\varphi_{\sigma,D}(u-\mu)\nu_j(u)\,\mathrm{d}\mathcal{H}^{D-1}(u).$$

On $\mathcal{F}_{jk}$, the two outward normals satisfy $\nu_j = n_{jk}$ and $\nu_k = -n_{jk}$. Their contributions combine as

$$-\ell_j\varphi_{\sigma,D}n_{jk} - \ell_k\varphi_{\sigma,D}(-n_{jk}) = (\ell_k - \ell_j)\varphi_{\sigma,D}n_{jk}.$$

Higher-order intersections have zero surface measure. Summing over all interfaces proves equation 13. Assumption B.1 extends the argument to locally finite partitions by truncation and dominated convergence. □

Direct differentiation also gives

$$\nabla_\mu \varphi_{\sigma,D}(u-\mu) = \frac{u-\mu}{\sigma^2}\varphi_{\sigma,D}(u-\mu),$$

and therefore

$$\nabla F_\sigma(\mu) = \int L(u)\frac{u-\mu}{\sigma^2}\varphi_{\sigma,D}(u-\mu)\,\mathrm{d}u = \frac{1}{\sigma}\mathbb{E}[L(\mu+\sigma\xi)\xi],$$

which agrees with equation 11.

Flat interfaces. If an interface lies in a complete hyperplane H at Euclidean distance δ from μ, rotate coordinates so that the first axis is normal to H. Integrating the remaining coordinates gives

$$\int_H \varphi_{\sigma,D}(u-\mu)\,\mathrm{d}\mathcal{H}^{D-1}(u) = \frac{1}{\sigma}\phi\left(\frac{\delta}{\sigma}\right). \tag{14}$$

An interface contained in H has no larger Gaussian surface mass.

B.3 Activation and Rare-Boundary Complexity

Let $Y = L(\mu+\sigma\xi)$ take distinct values $v_1,\dots,v_R$ with probabilities $\rho_r = \Pr(Y = v_r)$. A group is *active* when its K feedback values are not all equal. Denote its activation probability by q_K.

Group activation. The activation probability and its relation to the gradient estimator are

$$q_K = 1 - \sum_{r=1}^{R}\rho_r^K, \tag{15}$$

$$\Pr(G_K \neq 0) \le q_K \le q_K^{\mathcal{D}}, \tag{16}$$

where $q_K^{\mathcal{D}}$ is the probability that the group contains at least two distinct downstream decisions.

Proof. A group is inactive exactly when all K losses equal the same value v_r. These events are disjoint and have probabilities ρ_r^K, which proves equation 15. In an inactive group, $Y_i = \bar{Y}_{-i}$ for every i, so $G_K = 0$ and $\Pr(G_K \neq 0) \le q_K$.

Let π_j be the probability of decision z_j, and let $\mathcal{I}_r$ contain the decisions with loss v_r. Then $\rho_r = \sum_{j\in\mathcal{I}_r}\pi_j$, so

$$\rho_r^K = \left(\sum_{j\in\mathcal{I}_r}\pi_j\right)^K \ge \sum_{j\in\mathcal{I}_r}\pi_j^K.$$

Summing over r gives $\sum_r \rho_r^K \ge \sum_j \pi_j^K$, and hence $q_K \le 1 - \sum_j \pi_j^K = q_K^{\mathcal{D}}$. □

Activation under a future budget. The number N_B of active groups among M_B independent future groups is binomial with parameters M_B and q_K. Define

$$A_B = -M_B\log(1-q_K),\ \ \Pr(N_B \ge 1) = 1 - e^{-A_B}, \tag{17}$$

with $A_B = +\infty$ when $q_K = 1$. For $q_K \ll 1$, $A_B = M_B q_K[1+o(1)]$.

Binary switching model. Consider

$$L(u) = L_0 + \Delta\,\mathbf{1}\{n^\top u < \beta\},\ \ \|n\| = 1,\ \ \Delta \neq 0. \tag{18}$$

Suppose $d = n^\top\mu - \beta > 0$, and define

$$a = \frac{d}{\sigma},\ \ p = \Phi(-a).$$

One perturbation reaches the rare side with probability p. A group is active unless all samples fall on one side, so

$$q_K(p) = 1 - (1-p)^K - p^K. \tag{19}$$

Rare-boundary query complexity. Let $\mathcal{Q}_K(p)$ be the expected number of evaluator calls until the first active group. For every $K \geq 2$ and $p \in (0, 1/2]$,

$$\mathcal{Q}_K(p) = \frac{K}{q_K(p)} \geq \frac{1}{p},$$
$$\inf_{K \geq 2} \mathcal{Q}_K(p) = \Theta\left(\frac{1}{p}\right) \quad (p \downarrow 0). \tag{20}$$

For Gaussian perturbations,

$$\frac{1}{p} = \frac{1}{\Phi(-a)} \sim \sqrt{2\pi}\, a e^{a^2/2} \ \ (a \to \infty). \tag{21}$$

Proof. Independent groups produce a geometric waiting time with success probability $q_K(p)$, and each group costs K calls. Every active group contains at least one rare-side sample. Therefore

$$q_K(p) \leq 1 - (1-p)^K \leq Kp,$$

which gives $\mathcal{Q}_K(p) \geq 1/p$.

For the matching order, set $K_p = \lceil 1/p \rceil$. Then $K_p p \to 1$ and

$$K_p \log(1-p) = -K_p p[1 + o(1)] \to -1.$$

Thus $(1-p)^{K_p} \to e^{-1}$, while $p^{K_p} \to 0$. Hence $q_{K_p}(p) \to 1 - e^{-1} > 0$ and $\mathcal{Q}_{K_p}(p) = O(1/p)$. This proves the second part of equation 20.

Finally, Mills' ratio gives

$$\Phi(-a) = \frac{\phi(a)}{a}[1 + o(1)],$$

which yields equation 21. □

Increasing K does not remove the rare-transition bottleneck under a fixed total budget. If $Kp \to 0$, then

$$q_K(p) = Kp[1 + o(1)].$$

Writing $B = M_B K + r_B$ with $0 \leq r_B < K$,

$$M_B q_K(p) = (B - r_B)p[1 + o(1)].$$

Apart from the incomplete final group, the expected number of active groups is therefore Bp to first order.

B.4 Tangential Noise

We continue with the binary model in equation 18. Write

$$Z_i = n^\top \xi_i, \ \ B_i = \mathbf{1}\{Z_i < -a\}, \ \ \bar{B} = \frac{1}{K}\sum_{i=1}^{K} B_i.$$

The constant L_0 disappears after centering, and $B_i - \bar{B}_{-i} = K(B_i - \bar{B})/(K-1)$. Thus

$$G_K = \frac{\Delta}{\sigma(K-1)} \sum_{i=1}^{K} (B_i - \bar{B})\xi_i. \tag{22}$$

Normal–tangential variance decomposition. Let $p = \Phi(-a)$. Then

$$m = -\frac{\Delta}{\sigma}\phi(a)n, \ \ \mathrm{Cov}(G_K) = v_\parallel nn^\top + v_\perp(I_D - nn^\top), \tag{23}$$

where

$$v_\perp = \frac{\Delta^2}{\sigma^2(K-1)}p(1-p), \ v_\parallel = \frac{\Delta^2}{\sigma^2 K}\left[a\phi(a)(1-2p) + \frac{K}{K-1}p(1-p) - \frac{K-2}{K-1}\phi(a)^2\right]. \tag{24}$$

For fixed K, as $a \to \infty$,

$$v_\parallel = \frac{\Delta^2}{\sigma^2 K}a^2 p[1 + o(1)]. \tag{25}$$

Proof. Decompose $\xi_i = Z_i n + \xi_{i,\perp}$, where $\xi_{i,\perp}$ is independent of Z_i and has mean zero. Expanding the sample mean gives

$$\mathbb{E}\left[\sum_{i=1}^{K}(B_i - \bar{B})\xi_i\right] = (K-1)\mathbb{E}[B_i\xi_i].$$

Moreover,

$$\begin{aligned}\mathbb{E}[B_i\xi_i] &= \mathbb{E}[Z_i\mathbf{1}\{Z_i < -a\}]n \\ &= \left(\int_{-\infty}^{-a} z\phi(z)\,\mathrm{d}z\right)n \\ &= -\phi(a)n.\end{aligned}$$

Substitution into equation 22 gives the stated population gradient.

The joint law is invariant under any orthogonal transformation that fixes n. Its covariance must therefore have the form in equation 23. To compute $v_\perp$, fix a unit vector $e \perp n$ and let $W_i = e^\top \xi_i$. Conditional on $B_1, \ldots, B_K$, the W_i are independent standard normal variables, and

$$e^\top G_K = \frac{\Delta}{\sigma(K-1)}\sum_{i=1}^{K}(B_i - \bar{B})W_i.$$

Its conditional mean is zero and its conditional variance is

$$\frac{\Delta^2}{\sigma^2(K-1)^2}\sum_{i=1}^{K}(B_i - \bar{B})^2.$$

For independent Bernoulli(p) variables,

$$\mathbb{E}\left[\sum_{i}(B_i - \bar{B})^2\right] = (K-1)p(1-p).$$

Taking expectations proves equation 24.

For the normal direction, set $X_i = B_i - p$ and $T_i = Z_i$. Then

$$n^\top G_K = \frac{\Delta}{\sigma}\widehat{c}_K, \;\; \widehat{c}_K = \frac{1}{K-1}\sum_{i=1}^{K}(X_i - \bar{X})(T_i - \bar{T}),$$

so $\widehat{c}_K$ is the unbiased sample covariance of X_i and T_i. Define

$$h_{ij} = \frac{1}{2}(X_i - X_j)(T_i - T_j).$$

The identity

$$\sum_{i<j}(X_i - X_j)(T_i - T_j) = K\sum_{i}(X_i - \bar{X})(T_i - \bar{T})$$

gives

$$\widehat{c}_K = \binom{K}{2}^{-1}\sum_{i<j} h_{ij}.$$

Let $\mu_{rs} = \mathbb{E}[X^r T^s]$. Direct expansion yields

$$\mathbb{E}[h_{12}] = \mu_{11}, \;\; \mathrm{Var}(h_{12}) = \frac{1}{2}(\mu_{22} + \mu_{20}\mu_{02}).$$

Pairs with disjoint indices are independent. For pairs sharing one index,

$$\mathbb{E}[h_{12} \mid X_1, T_1] = \frac{1}{2}(X_1 T_1 + \mu_{11}),$$

and hence

$$\mathrm{Cov}(h_{12}, h_{13}) = \frac{1}{4}(\mu_{22} - \mu_{11}^2).$$

There are $C_K = \binom{K}{2}$ pair terms and $K\binom{K-1}{2}$ unordered pairs of terms sharing one index. Therefore

$$\mathrm{Var}(\widehat{c}_K) = \frac{1}{K}\left[\mu_{22} + \frac{\mu_{20}\mu_{02}}{K-1} - \frac{K-2}{K-1}\mu_{11}^2\right]. \tag{26}$$

For $B \sim \mathrm{Bernoulli}(p)$ and $T \sim \mathcal{N}(0,1)$,

$$\mu_{20} = p(1-p), \;\; \mu_{02} = 1, \;\; \mu_{11} = \mathbb{E}[BT] = -\phi(a).$$

Integration by parts gives $\mathbb{E}[T^2\mathbf{1}\{T < -a\}] = p + a\phi(a)$, so

$$\begin{aligned}\mu_{22} &= \mathbb{E}[(B-p)^2T^2] \\ &= (1-p)^2[p + a\phi(a)] + p^2[1 - p - a\phi(a)] \\ &= p(1-p) + a\phi(a)(1-2p).\end{aligned}$$

Substituting these moments into equation 26 and multiplying by Δ^2/σ^2 proves equation 24.

Finally, Mills' ratio gives $\phi(a) = ap[1 + o(1)]$. Thus $a\phi(a) = a^2p[1 + o(1)]$, while the other terms in equation 24 are lower order than a^2p. This proves equation 25. □

Effective dimension. Let P_r, $1 \le r \le D$, be a rank-r orthogonal projector with $P_r n = n$. It retains the normal direction and $r-1$ tangential directions. Then

$$\mathrm{SNR}_r := \frac{\|m\|^2}{\mathbb{E}\|P_r(G_K - m)\|^2} = \frac{\Delta^2\phi(a)^2/\sigma^2}{v_\parallel + (r-1)v_\perp}. \tag{27}$$

For fixed K and r, as $a \to \infty$,

$$\mathrm{SNR}_r = \frac{K(K-1)a^2p}{(K-1)a^2 + K(r-1)}[1 + o(1)]. \tag{28}$$

Proof. The covariance decomposition gives

$$\mathbb{E}\|P_r(G_K - m)\|^2 = v_\parallel + (r-1)v_\perp,$$

which proves equation 27. In the rare-boundary limit,

$$\begin{aligned}\phi(a)^2 &= a^2p^2[1 + o(1)], \\ v_\parallel &= \frac{\Delta^2a^2p}{\sigma^2K}[1 + o(1)], \\ v_\perp &= \frac{\Delta^2p}{\sigma^2(K-1)}[1 + o(1)].\end{aligned}$$

Substitution into equation 27 yields equation 28. □

B.5 Budgeted Gradient Recovery

Let

$$S = \|m\|^2, \;\; V_K = \mathbb{E}\|G_K - m\|^2. \tag{29}$$

For $V_K > 0$, define the budgeted gradient resolvability

$$\Lambda_B = M_B\frac{S}{V_K}. \tag{30}$$

This definition also covers $S = 0$, in which case $\Lambda_B = 0$. Relative error is considered only when $S > 0$.

Relative mean squared error. For $M_B \geq 1$ independent groups,

$$\mathbb{E}\|\bar{G}_B - m\|^2 = \frac{V_K}{M_B}. \tag{31}$$

Consequently, if $S > 0$ and $0 < V_K < \infty$,

$$\frac{\mathbb{E}\|\bar{G}_B - m\|^2}{\|m\|^2} = \frac{1}{\Lambda_B}. \tag{32}$$

For any tolerated relative root mean squared error $\varepsilon > 0$,

$$\mathbb{E}\|\bar{G}_B - m\|^2 \leq \varepsilon^2 \|m\|^2 \quad \Longleftrightarrow \quad \Lambda_B \geq \varepsilon^{-2}. \tag{33}$$

Proof. By equation 12,

$$\bar{G}_B - m = \frac{1}{M_B} \sum_{b=1}^{M_B} (G_K^{(b)} - m).$$

The centered group estimates are independent, so every cross inner product has zero expectation. Therefore

$$\mathbb{E}\|\bar{G}_B - m\|^2 = \frac{1}{M_B^2} \sum_{b=1}^{M_B} \mathbb{E}\|G_K^{(b)} - m\|^2 = \frac{V_K}{M_B}.$$

Dividing by $S = \|m\|^2$ proves equation 32; comparing equation 31 with $\varepsilon^2 S$ proves equation 33. □

The same identity gives a probability bound. Markov's inequality implies

$$\Pr\big(\|\bar{G}_B - m\| \geq \varepsilon \|m\|\big) \leq \frac{1}{\varepsilon^2 \Lambda_B}. \tag{34}$$

In particular, $\langle \bar{G}_B, m \rangle \leq 0$ implies $\|\bar{G}_B - m\| \geq \|m\|$, and hence

$$\Pr(\langle \bar{G}_B, m \rangle \leq 0) \leq \min\{1, \Lambda_B^{-1}\}. \tag{35}$$

Activation limits resolvability. If $m \neq 0$, G_K is square integrable, and $q_K < 1$, then

$$\Lambda_B \leq \frac{M_B q_K}{1 - q_K}. \tag{36}$$

Proof. We first prove a general inequality. Let G be square integrable with $m = \mathbb{E}[G] \neq 0$, and let $q_G = \Pr(G \neq 0)$. Conditional Jensen's inequality gives

$$\mathbb{E}\|G\|^2 = q_G \mathbb{E}[\|G\|^2 \mid G \neq 0] \geq q_G \|\mathbb{E}[G \mid G \neq 0]\|^2.$$

Since $G = 0$ on the complementary event, $m = q_G \mathbb{E}[G \mid G \neq 0]$. Thus

$$\mathbb{E}\|G\|^2 \geq \frac{\|m\|^2}{q_G}, \;\; \mathbb{E}\|G - m\|^2 \geq \frac{1 - q_G}{q_G} \|m\|^2.$$

An inactive group has $G_K = 0$, so $q_G = \Pr(G_K \neq 0) \leq q_K$. Because $(1 - q)/q$ decreases on $(0, 1)$,

$$V_K \geq \frac{1 - q_K}{q_K} \|m\|^2.$$

Taking reciprocals and multiplying by M_B proves equation 36. □

If $\Lambda_B \geq \lambda > 0$, then $q_K \geq \lambda/(M_B + \lambda)$. Since $A_B = -M_B \log(1 - q_K)$ increases with q_K,

$$A_B \geq M_B \log\left(1 + \frac{\lambda}{M_B}\right). \tag{37}$$

Thus a resolvable gradient necessarily requires some activation evidence, while frequent activation alone does not guarantee a low-variance direction.

Threshold-equivalent margin. Let

$$\Lambda_{\mathrm{req}} = \varepsilon^{-2}. \tag{38}$$

Define

$$D_B = M_B S - \Lambda_{\mathrm{req}} V_K. \tag{39}$$

Whenever $V_K > 0$,

$$D_B \geq 0 \quad \Longleftrightarrow \quad \Lambda_B \geq \Lambda_{\mathrm{req}}. \tag{40}$$

The margin remains well defined when $V_K = 0$, which avoids a ratio at a zero or nearly zero variance.

B.6 Pairwise Representation and Pilot Estimation

The following pairwise form is useful for estimating $S = \|m\|^2$. Let $X_i = (\xi_i, Y_i)$, and for two samples define

$$h(X_i, X_j) = \frac{(Y_i - Y_j)(\xi_i - \xi_j)}{2\sigma}. \tag{41}$$

Then

$$G_K = \binom{K}{2}^{-1} \sum_{i<j} h(X_i, X_j), \;\; \mathbb{E}[h(X, X')] = m, \tag{42}$$

where X and X' are independent copies.

Proof. Expanding the pairwise sum gives

$$\sum_{i<j} (Y_i - Y_j)(\xi_i - \xi_j) = K \sum_{i=1}^{K} (Y_i - \bar{Y})\xi_i.$$

Since $2\sigma\binom{K}{2} = \sigma K(K-1)$, the first identity follows from equation 11.

For independent $X = (\xi, Y)$ and $X' = (\xi', Y')$,

$$\mathbb{E}[h(X, X')] = \frac{1}{2\sigma} \mathbb{E}[Y\xi - Y\xi' - Y'\xi + Y'\xi'].$$

Independence and $\mathbb{E}[\xi] = \mathbb{E}[\xi'] = 0$ eliminate the cross terms. The remaining terms are equal, so $\mathbb{E}[h(X, X')] = \sigma^{-1}\mathbb{E}[Y\xi] = m$. □

Pilot estimators. Let $G^{(1)}, \ldots, G^{(M_0)}$, $M_0 \geq 2$, be independent pilot group estimates drawn at the same fixed configuration (μ, σ, K) as the future groups. The pilot is separate from the future budget. Let I_b be the indicator that pilot group b is active, and define

$$\widehat{q}_K = \frac{1}{M_0} \sum_{b=1}^{M_0} I_b. \tag{43}$$

Also define

$$\bar{G}_0 = \frac{1}{M_0} \sum_{b=1}^{M_0} G^{(b)}, \;\; \widehat{T} = \frac{1}{M_0} \sum_{b=1}^{M_0} \|G^{(b)}\|^2. \tag{44}$$

The signal and one-group variance estimators are

$$\begin{aligned} \widehat{S} &= \frac{1}{M_0(M_0-1)} \sum_{b \neq c} \langle G^{(b)}, G^{(c)} \rangle \\ &= \frac{M_0 \|\bar{G}_0\|^2 - \widehat{T}}{M_0 - 1}, \end{aligned} \tag{45}$$

$$\widehat{V} = \frac{1}{M_0 - 1} \sum_{b=1}^{M_0} \|G^{(b)} - \bar{G}_0\|^2 = \widehat{T} - \widehat{S}. \tag{46}$$

For a future budget B, define

$$\widehat{A}_B = -M_B \log(1 - \widehat{q}_K), \tag{47}$$

$$\widehat{\Lambda}_B = M_B \frac{\widehat{S}}{\widehat{V}} \quad (\widehat{V} > 0), \tag{48}$$

$$\widehat{D}_B = M_B \widehat{S} - \Lambda_{\mathrm{req}} \widehat{V}. \tag{49}$$

We set $\widehat{A}_B = +\infty$ when $\widehat{q}_K = 1$. The plug-in ratio $\widehat{\Lambda}_B$ is used for continuous prediction; categorical diagnosis uses $\widehat{D}_B$ and its confidence bounds.

Unbiasedness and consistency. If $\mathbb{E}\|G_K\|^2 < \infty$, then

$$\begin{aligned} \mathbb{E}[\widehat{q}_K] &= q_K, & \mathbb{E}[\widehat{S}] &= S, \\ \mathbb{E}[\widehat{V}] &= V_K, & \mathbb{E}[\widehat{D}_B] &= D_B. \end{aligned} \tag{50}$$

All four estimators converge almost surely to their population values as $M_0 \to \infty$. If $q_K < 1$, then $\widehat{A}_B \to A_B$ almost surely. If additionally $V_K > 0$, then $\widehat{\Lambda}_B \to \Lambda_B$ almost surely.

Proof. The indicators I_b are independent Bernoulli variables with success probability q_K, proving the first expectation and the almost-sure convergence of $\widehat{q}_K$.

For $b \neq c$, independence gives

$$\mathbb{E}\langle G^{(b)}, G^{(c)} \rangle = \langle \mathbb{E}G^{(b)}, \mathbb{E}G^{(c)} \rangle = \|m\|^2 = S.$$

Hence $\mathbb{E}[\widehat{S}] = S$. The identity

$$\left\| \sum_{b=1}^{M_0} G^{(b)} \right\|^2 = \sum_{b=1}^{M_0} \|G^{(b)}\|^2 + \sum_{b \neq c} \langle G^{(b)}, G^{(c)} \rangle$$

gives the second expression in equation 45. The usual sample-variance identity gives equation 46, and thus $\mathbb{E}[\widehat{V}] = V_K$. Linearity then gives $\mathbb{E}[\widehat{D}_B] = D_B$.

The moment condition implies $\mathbb{E}\|G_K\| < \infty$. The vector and scalar strong laws yield

$$\bar{G}_0 \to m, \;\; \widehat{T} \to \mathbb{E}\|G_K\|^2 \quad \text{almost surely.}$$

Since

$$\widehat{S} = \frac{M_0}{M_0 - 1} \|\bar{G}_0\|^2 - \frac{\widehat{T}}{M_0 - 1},$$

we have $\widehat{S} \to S$, followed by $\widehat{V} \to V_K$ and $\widehat{D}_B \to D_B$. The remaining convergences follow from continuity of the logarithm and ratio on their stated domains. □

The unbiased estimator $\widehat{S}$ is not constrained to be nonnegative and may be negative for a finite pilot. It is kept unmodified when constructing confidence bounds. A negative value indicates that the pilot has not separated population signal from sampling noise.

C Finite-Pilot Inference

The statistical unit in this section is one complete pilot group $G^{(b)}$, not an individual perturbation within that group. All bounds are constructed for one fixed configuration (μ, σ, K, B). Simultaneous claims across several configurations require an additional multiplicity correction.

Activation admits an exact binomial interval. Signal, variance, and the diagnostic margin use a group-level multiplier procedure. These multiplier bounds are pointwise and asymptotic in the number of pilot groups.

C.1 Exact Activation Bounds

Let

$$X = \sum_{b=1}^{M_0} I_b$$

be the number of active pilot groups. Then $X \sim \mathrm{Binomial}(M_0, q_K)$. Write $\mathrm{Beta}^{-1}(\gamma; a, b)$ for the γ-quantile of a $\mathrm{Beta}(a, b)$ distribution.

For a two-sided Clopper–Pearson interval with coverage $1 - \alpha_A$, define

$$q_L = \mathrm{Beta}^{-1}\Big(\frac{\alpha_A}{2}; X, M_0 - X + 1\Big), \tag{51}$$

$$q_U = \mathrm{Beta}^{-1}\Big(1 - \frac{\alpha_A}{2}; X + 1, M_0 - X\Big), \tag{52}$$

when $0 < X < M_0$. At the endpoints, set $q_L = 0$ when $X = 0$ and $q_U = 1$ when $X = M_0$; the other endpoint is given by its applicable Beta quantile.

For a one-sided bound with coverage $1 - \alpha_A$, replace $\alpha_A/2$ by α_A in the lower bound, or replace $1 - \alpha_A/2$ by $1 - \alpha_A$ in the upper bound. Since $A_B(q) = -M_B \log(1 - q)$ is increasing,

$$\mathrm{LCB}(A_B) = A_B(q_L), \;\; \mathrm{UCB}(A_B) = A_B(q_U), \tag{53}$$

with $A_B(1) = +\infty$.

C.2 Multiplier Bounds for Signal, Variance, and Margin

The signal estimator is the second-order U-statistic

$$\widehat{S} = \frac{1}{M_0(M_0 - 1)} \sum_{b \neq c} H_{bc}, \;\; H_{bc} = \langle G^{(b)}, G^{(c)} \rangle.$$

Its first-order empirical projection is

$$\widehat{h}_{1,b} = \frac{1}{M_0 - 1} \sum_{c \neq b} H_{bc} - \widehat{S}, \tag{54}$$

and the remaining second-order component is

$$\widehat{h}_{2,bc} = H_{bc} - \widehat{S} - \widehat{h}_{1,b} - \widehat{h}_{1,c}. \tag{55}$$

The first component controls the usual regime with a detectable population gradient. The multiplier construction includes both empirical projection terms.

For multiplier replicate $r = 1, \ldots, R_{\mathrm{mult}}$, draw independent

$$\zeta_1^{(r)}, \ldots, \zeta_{M_0}^{(r)} \sim \mathcal{N}(0, 1)$$

and compute

$$Z_S^{(r)} = \frac{2}{M_0} \sum_{b=1}^{M_0} \zeta_b^{(r)} \widehat{h}_{1,b} + \frac{1}{M_0(M_0 - 1)} \sum_{b \neq c} \zeta_b^{(r)} \zeta_c^{(r)} \widehat{h}_{2,bc}, \tag{56}$$

$$Z_T^{(r)} = \frac{1}{M_0} \sum_{b=1}^{M_0} \zeta_b^{(r)} \left(\|G^{(b)}\|^2 - \widehat{T} \right), \tag{57}$$

$$Z_V^{(r)} = Z_T^{(r)} - Z_S^{(r)}, \tag{58}$$

$$Z_D^{(r)} = M_B Z_S^{(r)} - \Lambda_{\mathrm{req}} Z_V^{(r)}. \tag{59}$$

The same multipliers are used in all four quantities, preserving their estimated dependence.

Let $\widehat{\tau}_S$, $\widehat{\tau}_V$, and $\widehat{\tau}_D$ be the empirical standard deviations of the corresponding multiplier replicates. A two-sided joint calibration uses

$$T^{(r)}_{\text{two}} = \max\left\{ \frac{|Z^{(r)}_S|}{\widehat{\tau}_S}, \frac{|Z^{(r)}_V|}{\widehat{\tau}_V}, \frac{|Z^{(r)}_D|}{\widehat{\tau}_D} \right\}. \tag{60}$$

Any coordinate with zero multiplier standard deviation is omitted from the maximum. If such a coordinate is required for a diagnosis and is not known deterministically, that diagnosis is left unresolved.

Let c_{two} be the empirical $(1-\alpha_{SV})$-quantile of $T^{(r)}_{\text{two}}$. The resulting simultaneous bounds are

$$S \in [\widehat{S} - c_{\text{two}}\widehat{\tau}_S, \widehat{S} + c_{\text{two}}\widehat{\tau}_S] \cap [0, \infty), \tag{61}$$

$$V_K \in [\widehat{V} - c_{\text{two}}\widehat{\tau}_V, \widehat{V} + c_{\text{two}}\widehat{\tau}_V] \cap [0, \infty), \tag{62}$$

$$D_B \in [\widehat{D}_B - c_{\text{two}}\widehat{\tau}_D, \widehat{D}_B + c_{\text{two}}\widehat{\tau}_D]. \tag{63}$$

The point estimate $\widehat{S}$ is not clipped before the multiplier calculation. If either nonnegative intersection is empty, the corresponding case is unresolved.

Joint bounds used by the four-case diagnosis. The complete diagnostic rule may need a lower bound on S, a lower bound on D_B, or an upper bound on D_B, depending on the observed pilot. We therefore calibrate these three directions together:

$$T^{(r)}_{\text{diag}} = \max\left\{ \frac{Z^{(r)}_S}{\widehat{\tau}_S}, \frac{|Z^{(r)}_D|}{\widehat{\tau}_D} \right\}. \tag{64}$$

If c_{diag} is its empirical $(1-\alpha_{SV})$-quantile, then

$$\text{LCB}(S) = \max\{0, \widehat{S} - c_{\text{diag}}\widehat{\tau}_S\}, \tag{65}$$

$$\text{LCB}(D_B) = \widehat{D}_B - c_{\text{diag}}\widehat{\tau}_D, \tag{66}$$

$$\text{UCB}(D_B) = \widehat{D}_B + c_{\text{diag}}\widehat{\tau}_D. \tag{67}$$

If an analysis prespecifies only one directional conclusion, the unused directions may be removed from equation 64. In particular, recovery uses

$$T^{(r)}_{\text{res}} = \max\left\{ \frac{Z^{(r)}_S}{\widehat{\tau}_S}, \frac{Z^{(r)}_D}{\widehat{\tau}_D} \right\},$$

whereas insufficient resolution uses

$$T^{(r)}_{\text{low}} = -\frac{Z^{(r)}_D}{\widehat{\tau}_D}.$$

Their quantiles replace c_{diag} in the applicable lower or upper bound. The joint statistic in equation 64 provides one simultaneous statement over all moment directions. The directional versions provide a pointwise bound for the corresponding issued conclusion; treating both directions as one simultaneous family requires an additional error allocation.

Optional bounds for continuous resolvability. The categorical rule uses D_B, while Λ_B remains useful as a continuous measure. Given any joint confidence region $\mathcal{C}_{SV} \subseteq [0,\infty)^2$ for (S, V_K), define

$$\text{LCB}(\Lambda_B) = \inf_{\substack{(s,v)\in\mathcal{C}_{SV} \\ v>0}} M_B \frac{s}{v}, \tag{68}$$

$$\text{UCB}(\Lambda_B) = \sup_{\substack{(s,v)\in\mathcal{C}_{SV} \\ v>0}} M_B \frac{s}{v}. \tag{69}$$

If the region reaches $v = 0$ while containing some $s > 0$, the upper bound is $+\infty$. These ratio bounds are reported descriptively and are not used for categorical diagnosis.

C.3 What the Confidence Statements Guarantee

The activation interval above has exact finite-sample coverage. The multiplier bounds approximate the sampling distribution of the pilot estimators. Their required validity condition can be stated directly.

Multiplier validity condition. For the fixed configuration under study, assume that the conditional distribution of (Z_S, Z_V, Z_D) given the pilot consistently approximates the joint sampling distribution of

$$(\widehat{S} - S, \widehat{V} - V_K, \widehat{D}_B - D_B)$$

in the standardized directions used by the selected bound.

Coverage consequence. Under the multiplier validity condition and continuity at the calibrated quantile, the bounds in the preceding subsection have pointwise asymptotic coverage at least $1-\alpha_{SV}$ as $M_0 \to \infty$. If activation and multiplier error levels satisfy

$$\alpha_A + \alpha_{SV} \le \alpha, \tag{70}$$

then the activation and signal–margin statements used in one diagnosis hold jointly with probability at least $1 - \alpha - o(1)$.

Proof. Conditional multiplier consistency transfers the calibrated multiplier quantile to the corresponding sampling error statistic, giving coverage $1 - \alpha_{SV} - o(1)$. The Clopper–Pearson bound fails with probability at most α_A. The union bound and equation 70 give the joint statement. □

This guarantee is pointwise in (μ, σ, K, B). When several budgets or configurations are inspected and each receives a simultaneous claim, the nominal error level is adjusted across that collection.

C.4 Confidence-Aware Diagnosis

Let $1 - \delta_A$ be the required probability of observing at least one active future group, and let ε be the tolerated relative root mean squared error. The required levels are

$$A_{\mathrm{req}} = -\log \delta_A, \;\; \Lambda_{\mathrm{req}} = \varepsilon^{-2}. \tag{71}$$

The diagnosis for a future budget B has four cases:

1. Starved if $\mathrm{UCB}(A_B) \le A_{\mathrm{req}}$. The future budget does not meet the prescribed activation requirement.
2. Low-resolution if $\mathrm{LCB}(A_B) > A_{\mathrm{req}}$ and $\mathrm{UCB}(D_B) < 0$. Comparative feedback is likely to occur, but the budget does not support recovery at the required accuracy.
3. Gradient-resolvable if $\mathrm{LCB}(A_B) > A_{\mathrm{req}}$, $\mathrm{LCB}(S) > 0$, and $\mathrm{LCB}(D_B) \ge 0$. The budget supports recovery of a nonzero smoothed gradient at the required accuracy.
4. Unresolved otherwise. The pilot does not support a definite conclusion.

A low-resolution case with $\mathrm{LCB}(S) > 0$ may be reported more specifically as Noise-dominated: the pilot supports a nonzero gradient but rules out the required signal-to-variance ratio. This label refers to total estimation variance; it does not by itself identify whether that variance is tangential, normal, or caused by multiple competing interfaces.

The confidence level, one- or two-sided calibration, multiplier count, and error allocation are fixed before examining the diagnostic outcomes.

C.5 Floating-Point Feedback

The theoretical activation event uses exact equality of mathematical feedback values. For a deterministic evaluator, implementations should store a canonical scalar loss for each returned decision. If numerical evaluation introduces floating-point variation, an equality tolerance is fixed before examining diagnostic outcomes and held constant across compared configurations.

D Experimental Protocols and Additional Results

This appendix is a report of the experimental constructions, evaluation procedures, and results used in Section 5. It first defines the grouped estimator, pilot statistics, confidence rule, independent future target, reported metrics, and query accounting. It then reports six experiments: two controlled tests of boundary starvation and tangential noise, two stress tests of geometric approximations, and two application studies on SEMO Energy-Knapsack and UCI HAR. Each experiment states its purpose, problem, design, metrics, and observed results. Complete grids are reported so that no definition or result required to interpret the experiments must be recovered from the main text.

The appendix uses the same primary thresholds as the main text,

$$A_{\mathrm{req}} = -\log(0.05), \;\; \Lambda_{\mathrm{req}} = 1.$$

Unless stated otherwise, the application experiments use 90% overall pointwise confidence, 95% component bounds under the Bonferroni allocation, and mutually independent pilot, reference, and future perturbation streams.

D.1 Common Evaluation Protocol

Sampling model, grouped estimator, and call accounting. At a fixed sampling state (μ, σ), each group draws $\xi_1, \ldots, \xi_K \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_D)$, forms candidates $U_i = \mu + \sigma\xi_i$, and obtains scalar feedback $Y_i = L(U_i)$ from the downstream evaluator. Every grouped experiment uses

$$G_K = \frac{1}{(K-1)\sigma} \sum_{i=1}^{K} (Y_i - \bar{Y})\xi_i, \;\; \bar{Y} = \frac{1}{K}\sum_{i=1}^{K} Y_i. \tag{72}$$

At the fixed configuration, $m = \mathbb{E}[G_K]$ denotes the population smoothed gradient estimated by one group. Every requested downstream evaluation is charged, including repeated decoded decisions. A group is loss-active when its largest and smallest scalar feedback values differ by more than the fixed numerical tolerance. We record decision diversity separately because distinct decisions can still have equal scalar feedback.

Application-space transformations. SEMO first forms the score in the 48-dimensional predicted-value space and maps it to predictor parameters through the exact Jacobian of the shared linear predictor. If A denotes this fixed linear map, then $H = AG_K$ is unbiased for Am and all application diagnostics are computed in the same parameter space used by the learner. HAR has a different exact invariance: Top-30 decoding is unchanged by $u \mapsto u + c\mathbf{1}$. Its fixed identifiable-space projector is therefore

$$P = I - \frac{1}{561}\mathbf{1}\mathbf{1}^{\top},$$

and every HAR group gradient is replaced by PG_K. Neither transformation is estimated from the pilot.

Pilot estimators. For n independent pilot groups $G^{(1)}, \ldots, G^{(n)}$, define

$$\widehat{S} = \frac{1}{n(n-1)} \sum_{b \neq c} \langle G^{(b)}, G^{(c)} \rangle, \tag{73}$$

$$\widehat{T} = \frac{1}{n}\sum_{b=1}^{n} \|G^{(b)}\|^2, \;\; \widehat{V} = \widehat{T} - \widehat{S}. \tag{74}$$

The cross-group expression is unbiased for $S = \|m\|^2$ and is not truncated when negative. The population one-group variance is $V_K = \mathbb{E}\|G_K - m\|^2$, and $\widehat{V}$ is its unbiased sample estimator. For future query budget B, $M_B = \lfloor B/K \rfloor$ independent groups are available. All tested budgets are divisible by their corresponding group sizes, so $M_B = B/K$ on the reported grids. At the population level, gradient resolvability and its threshold-equivalent margin are

$$\Lambda_B = M_B \frac{S}{V_K}, \;\; D_B = M_B S - \Lambda_{\mathrm{req}} V_K.$$

The pilot estimators for a future budget are

$$\widehat{\Lambda}_B = M_B \frac{\widehat{S}}{\widehat{V}}, \;\; \widehat{D}_B = M_B \widehat{S} - \Lambda_{\mathrm{req}} \widehat{V}, \;\; \Lambda_{\mathrm{req}} = \varepsilon^{-2} = 1. \tag{75}$$

At the population level, $D_B < 0$ is equivalent to $\Lambda_B < \Lambda_{\mathrm{req}}$ whenever $V_K > 0$. We use $1/\widehat{\Lambda}_B$ as the continuous predictor of future relative MSE and the threshold-equivalent margin $\widehat{D}_B$ for certification.

Activation estimator and bounds. Let q_K denote the probability that a group contains at least two distinct scalar feedback values. The future activation evidence is

$$A_B = -M_B \log(1 - q_K), \;\; A_{\mathrm{req}} = -\log(0.05), \tag{76}$$

so $1 - e^{-A_B}$ is the probability of at least one active group in the future budget. If I_b is the active-group indicator for pilot group b, then $\widehat{q}_K = n^{-1} \sum_{b=1}^{n} I_b$ and $\widehat{A}_B = -M_B \log(1 - \widehat{q}_K)$, with $\widehat{A}_B = +\infty$ when $\widehat{q}_K = 1$. The code uses exact one-sided Clopper–Pearson bounds on q_K, then maps them monotonically to bounds on A_B.

Multiplier bounds for signal and variance. The signal statistic is a second-order U-statistic. Write $h_{bc} = \langle G^{(b)}, G^{(c)} \rangle$ and define

$$\widehat{h}_{1,b} = \frac{1}{n-1} \sum_{c \neq b} h_{bc} - \widehat{S}, \tag{77}$$

$$\widehat{h}_{2,bc} = h_{bc} - \widehat{S} - \widehat{h}_{1,b} - \widehat{h}_{1,c}. \tag{78}$$

For i.i.d. standard Gaussian multipliers ζ_b, each replicate uses

$$Z_S = \frac{2}{n} \sum_b \zeta_b \widehat{h}_{1,b} + \frac{1}{n(n-1)} \sum_{b \neq c} \zeta_b \zeta_c \widehat{h}_{2,bc}, \tag{79}$$

$$Z_T = \frac{1}{n} \sum_b \zeta_b \left(\|G^{(b)}\|^2 - \widehat{T} \right), \;\; Z_V = Z_T - Z_S. \tag{80}$$

The quadratic term is evaluated using Gram identities without explicitly forming all $O(n^2)$ pair terms for every multiplier replicate. All primary application results use 1000 multiplier replicates.

Pointwise confidence allocation and diagnostic rule. The primary application setting is 90% *overall pointwise* confidence. For an overall level c, Bonferroni splitting between the activation and moment-inference families gives

$$c_{\mathrm{comp}} = 1 - \frac{1 - c}{2}.$$

Thus $c = 0.90$ uses 95% component bounds. The multiplier construction is asymptotic and pointwise, so the resulting intervals apply separately to each configuration rather than simultaneously over the full hyperparameter grid.

For the margin, let $Z_D = M_B Z_S - \Lambda_{\mathrm{req}} Z_V$. A one-sided upper bound on D_B certifies insufficient resolution. To certify recovery, a joint one-sided multiplier statistic for (S, D_B) yields lower bounds on both quantities. We write LCB and UCB for lower and upper confidence bounds. The ordered operational rule is

1. STARVED if $\mathrm{UCB}(A_B) \leq A_{\mathrm{req}}$;
2. UNRESOLVED if $\mathrm{LCB}(A_B) \leq A_{\mathrm{req}}$;
3. LOW-RESOLUTION if $\mathrm{UCB}(D_B) < 0$;
4. GRADIENT-RESOLVABLE if $\mathrm{LCB}(S) > 0$ and $\mathrm{LCB}(D_B) \geq 0$;
5. UNRESOLVED otherwise.

The negative operational label does not require a positive-signal certificate. Separately, the code reports a stronger NOISE-DOMINATED mechanistic certificate only when the pilot jointly establishes $S > 0$ and $D_B < 0$.

For comparison, the ratio baseline uses the same pilot groups and the same Hoeffding multiplier, but certifies the threshold through one-sided ratio-specific bounds on Λ_B. The lower-ratio direction protects against S decreasing and V increasing; the upper-ratio direction protects against S increasing and V decreasing. This baseline differs only in the inference geometry, not in data or oracle calls.

Independent future target. Pilot, reference, and future perturbation streams are mutually disjoint. The reference pool contains M_{ref} groups. Its mean $\widehat{m}_{\mathrm{ref}}$, signal estimate $\widehat{S}_{\mathrm{ref}}$, and variance estimate $\widehat{V}_{\mathrm{ref}}$ are computed using the same moment formulas as the pilot. Future groups are partitioned into W non-overlapping windows of M_B groups with means $\bar{G}_w^{\mathrm{fut}}$. We report

$$\widehat{\mathrm{RelMSE}}_{\mathrm{fut}} = \frac{\left[W^{-1}\sum_w \|\bar{G}_w^{\mathrm{fut}} - \widehat{m}_{\mathrm{ref}}\|^2 - \widehat{V}_{\mathrm{ref}}/M_{\mathrm{ref}}\right]_+}{\widehat{S}_{\mathrm{ref}}}, \tag{81}$$

when $\widehat{S}_{\mathrm{ref}} > 0$, where $[x]_+ = \max\{x, 0\}$. A future configuration is STARVED when its independently estimated $A_B \le A_{\mathrm{req}}$; otherwise it is LOW-RESOLUTION if corrected relative MSE exceeds one and GRADIENT-RESOLVABLE if relative MSE is at most one. Undefined directional outcomes remain UNRESOLVED. For the target-stability audit, a reference label is constructed from the reference pool alone: it is STARVED when the reference activation evidence is at most A_{req} and, otherwise, uses $\widehat{V}_{\mathrm{ref}}/(M_B\widehat{S}_{\mathrm{ref}})$ with the same unit threshold to distinguish LOW-RESOLUTION from GRADIENT-RESOLVABLE; nonpositive reference signal is UNRESOLVED. No pilot interval is used to define either evaluation label.

Reported metrics. Mechanism experiments compare empirical activation probabilities, gradient moments, and signal-to-noise ratios (SNRs) with their analytic values and report absolute or relative errors as specified in each table. Application experiments report three types of quantities. Activation calibration uses the mean absolute error (MAE) and Pearson correlation between pilot and future active-group probabilities and the Pearson/Spearman correlations between $\log \widehat{A}_B$ and $\log A_B^{\mathrm{fut}}$ on finite positive pairs. Gradient calibration compares $1/\widehat{\Lambda}_B$ with independent future relative mean squared error (MSE) using log-scale Pearson/Spearman correlations, the median future-to-predicted ratio, and the fraction of pairs within a factor of two. State diagnosis reports coverage, selective accuracy, and class recall. Coverage is the fraction of future-resolved configurations receiving a non-UNRESOLVED pilot label; selective accuracy is computed only over issued labels; class recall is reported separately for STARVED, LOW-RESOLUTION, and GRADIENT-RESOLVABLE. Reference–future agreement measures the stability of the independently constructed evaluation target. Cosine similarity is reported as a descriptive directional metric and does not define any state.

D.2 CONTROLLED MECHANISM EXPERIMENTS

D.2.1 BOUNDARY-STARVATION CONSTRUCTION

Purpose. This experiment tests whether boundary distance predicts the frequency and query cost of informative groups when every candidate is decoded by a complete combinatorial solver.

Problem. Let $\mathcal{P}$ be the set of right/down source–target paths on a 16×16 grid. For candidate node costs $u \in \mathbb{R}^{256}$, the decoder returns $\mathcal{D}(u) \in \arg\min_{z\in\mathcal{P}}\langle u, z\rangle$ using deterministic tie breaking, and the evaluator returns the fixed scalar loss of the selected path. For an isolated interface crossed with probability p, the binary prediction is $q_K = 1 - (1-p)^K - p^K$, and every inactive centered group has $G_K = 0$.

Design. Each coordinate is a node cost, giving $D = 256$ and $\binom{30}{15} = 155{,}117{,}520$ right/down source–target paths. Four independent base instances are generated. For each, the code finds two adjacent optimal paths, bisects their tie to numerical precision, and computes the nearest non-pair switching margin by dynamic programming. The perturbation scale is 2.5% of the base-cost standard deviation, and the retained interface is required to keep all other nominal-versus-competing

boundaries at least 5σ farther away over the tested range. The maximum pair-tie residual is 3.55×10^{-15}.

The tested grid is

$$a \in \{0.5, 1, 1.5, 2, 2.5, 3, 3.5\}, \quad K \in \{2, 4, 8, 16, 32, 64\},$$

with 65,536 perturbations per $(\text{instance}, a)$ pair. Hence the CSV contains 168 instance–distance–group-size rows.

Metrics. We record the observed loss-active probability, its absolute error relative to the exact finite-K binary law and the plug-in two-feedback occupancy law, the number of distinct decisions and feedback values, the number of inactive groups with nonzero centered updates, and requested solver calls per active group. The exact query comparator is K/q_K^{binary}; the rare-transition scale $1/\Phi(-a)$ is reported separately, where Φ is the standard normal CDF.

Results. All 168 rows contain exactly two observed decisions and exactly two scalar feedback values; the decision/loss activation gap is zero. Every loss-inactive probe group produces an exactly zero centered group update, with zero violations in the full audit.

Across the 168 rows, the loss-activation probability has MAE 2.83×10^{-3} against the exact local binary finite-K law and 1.51×10^{-3} against the plug-in two-feedback occupancy law. The mean relative error of observed calls per active group against the finite-K prediction is 5.38%.

Table 2: Boundary-starvation error summary by group size. For this compact table, the four instance estimates are first averaged at fixed (a, K) and the error is then averaged over a; this aggregation is therefore smaller than the all-168-row MAE reported in the text. Activation error is relative to the exact finite-K binary law and query error compares observed calls per active group with K/q_K^{binary}.

K	Activation MAE	Mean relative query error
2	4.4640e-04	2.53%
4	8.1033e-04	2.55%
8	1.1384e-03	2.53%
16	1.2974e-03	2.48%
32	2.2820e-03	2.68%
64	4.5091e-03	2.92%

Table 3: Observed loss-activation probability, averaged over the four certified shortest-path instances. Each cell is computed from the full solver.

a	$K=2$	$K=4$	$K=8$	$K=16$	$K=32$	$K=64$
0.5	0.4269	0.7634	0.9474	0.9975	1.0000	1.0000
1	0.2684	0.5001	0.7515	0.9362	0.9960	1.0000
1.5	0.1258	0.2436	0.4281	0.6719	0.8894	0.9868
2	0.0445	0.0880	0.1676	0.3055	0.5128	0.7542
2.5	0.0122	0.0242	0.0480	0.0935	0.1775	0.3218
3	0.0026	0.0052	0.0103	0.0206	0.0406	0.0791
3.5	0.0006	0.0011	0.0022	0.0045	0.0089	0.0178

At the reference setting $a = 2.5, K = 8$, the four instances yield active-group rates between 0.0458 and 0.0508 and calls per active group between 157.5 and 174.8. Their mean is 166.9, compared with the exact finite-K prediction 164.6 and the rare-transition scale $1/\Phi(-2.5) = 161.0$. At $a = 3.5, K = 8$, only $O(10)$ active groups occur per instance, so the finite-sample call estimate has larger sampling variation; its mean is 3.88×10^3 calls.

D.2.2 Tangential-Noise Construction

Purpose. This experiment tests whether restoring directions tangent to a loss-changing interface increases estimator variance and decreases SNR while leaving the population signal unchanged.

Table 4: Exact finite-K binary-interface activation probability $1-(1-p)^K-p^K$ for the same shortest-path grid, with $p=\Phi(-a)$.

a	$K=2$	$K=4$	$K=8$	$K=16$	$K=32$	$K=64$
0.5	0.4267	0.7623	0.9477	0.9973	1.0000	1.0000
1	0.2670	0.4983	0.7489	0.9370	0.9960	1.0000
1.5	0.1247	0.2416	0.4249	0.6692	0.8906	0.9880
2	0.0445	0.0879	0.1682	0.3080	0.5212	0.7707
2.5	0.0123	0.0246	0.0486	0.0949	0.1807	0.3288
3	0.0027	0.0054	0.0107	0.0214	0.0423	0.0828
3.5	0.0005	0.0009	0.0019	0.0037	0.0074	0.0148

Table 5: Observed requested solver calls per loss-active group, averaged over the four shortest-path instances. Calls are counted as K/q_K using the empirical active-group rate.

a	$K=2$	$K=4$	$K=8$	$K=16$	$K=32$	$K=64$
0.5	4.69	5.24	8.44	16.04	32.00	64.00
1	7.45	8.00	10.65	17.09	32.13	64.00
1.5	15.89	16.42	18.69	23.82	35.98	64.86
2	44.98	45.47	47.74	52.39	62.43	84.87
2.5	164.30	165.43	166.91	171.41	180.63	199.08
3	779.44	779.44	781.63	781.63	793.16	815.00
3.5	3879.95	3879.95	3879.95	3879.95	3879.95	3879.95

Problem. The decoder and path problem are the same as in Appendix D.2.1. A Householder rotation maps the certified interface normal to the first coordinate. In this basis, the population mean has only a normal component, while the remaining coordinates are tangent to the interface. A retained rank r keeps the normal and $r-1$ tangent coordinates.

Design. The same certified shortest-path construction is reused with an independent perturbation bank. Four instances are evaluated at

$$a \in \{0.5, 1.5, 2.5, 3.0\}, \quad K \in \{2, 4, 8, 16, 32\},$$

using 131,072 perturbations per instance and distance. A Householder rotation places the certified interface normal in coordinate one. Retained ranks

$$r \in \{1, 2, 4, 8, 16, 32, 64, 128, 256\}$$

keep the normal plus $r-1$ tangent coordinates. No other-boundary contamination is observed in any row.

Metrics. For every (a, K, r) setting, we estimate the normal mean, normal variance, per-direction tangent variance, tangent-coordinate coefficient of variation, and gradient SNR. Means and variances are compared with the exact finite-K binary-interface formulas. The retained-rank sweep reports the mean and standard error of the mean (SEM) over four independent instances.

Results. Across the complete grid, mean absolute relative errors are 3.65% for the normal mean, 4.02% for normal variance, and 3.66% for per-direction tangent variance. The tangent-coordinate coefficient of variation is approximately 0.107 in the $a=3$ rows and at most 0.054 in the $a \le 2.5$ rows.

Reference setting. At $a=2.5, K=8$, SNR decreases from 0.05014 at $r=1$ to 0.001355 at $r=256$, while the exact finite-K values are 0.04909 and 0.001330. The decline is produced by accumulated tangent variance, while the estimated population mean has only a normal component.

Table 6: Complete (a, K) tangential-noise audit, averaged over four instances. Ratios compare empirical quantities with the exact finite-K binary-interface formulas. $\mathrm{CV}_\perp$ is the coefficient of variation across tangent coordinates.

a	K	$\hat{m}_\parallel/m_\parallel$	$\hat{v}_\parallel/v_\parallel$	$\hat{v}_\perp/v_\perp$	$\mathrm{CV}_\perp$	$\mathrm{SNR}_{r=1}$	$\mathrm{SNR}_{r=256}$
0.5	2	0.999	0.999	1.000	0.008	0.5025	0.002268
0.5	4	0.999	0.998	1.001	0.009	1.846	0.006806
0.5	8	0.999	0.995	1.001	0.011	4.863	0.01588
0.5	16	1.000	0.994	1.002	0.016	11.14	0.034
0.5	32	1.001	0.985	1.003	0.022	24.03	0.07023
1.5	2	1.006	1.006	1.006	0.015	0.1138	0.00104
1.5	4	1.007	1.003	1.006	0.016	0.2784	0.00311
1.5	8	1.005	0.996	1.006	0.018	0.5983	0.007257
1.5	16	1.003	1.001	1.005	0.021	1.223	0.01554
1.5	32	1.003	0.999	1.006	0.025	2.483	0.03211
2.5	2	0.983	0.981	0.983	0.049	0.01128	0.0001953
2.5	4	0.983	0.981	0.983	0.050	0.02443	0.0005819
2.5	8	0.982	0.980	0.982	0.050	0.05014	0.001355
2.5	16	0.984	0.983	0.982	0.051	0.1009	0.0029
2.5	32	0.986	0.996	0.981	0.054	0.2001	0.005987
3	2	0.969	0.988	0.951	0.109	0.002529	5.957e-05
3	4	0.952	0.959	0.949	0.107	0.005524	0.0001774
3	8	0.950	0.948	0.953	0.107	0.01139	0.0004107
3	16	0.945	0.939	0.952	0.107	0.02315	0.0008793
3	32	0.947	0.939	0.951	0.107	0.04635	0.001817

Table 7: Full retained-rank sweep at the reference setting $a = 2.5, K = 8$. Values are mean $\pm$ SEM over four independent shortest-path instances.

Rank r	Empirical SNR	SEM	Finite-K theory
1	0.050141	0.000614	0.049094
2	0.043727	0.000517	0.043032
4	0.034995	0.000413	0.034510
8	0.025054	0.000307	0.024719
16	0.015964	0.000222	0.015770
32	0.009276	0.000140	0.009148
64	0.005039	0.000073	0.004972
128	0.002647	0.000038	0.002599
256	0.001355	0.000020	0.001330

D.3 Geometric Stress Tests

D.3.1 Finite-Pilot Subspace Recovery

Purpose. This experiment measures the finite-sample behavior of a pilot-estimated low-rank subspace when a rarely visited decision region has a large loss change. It records both short-budget variance and long-budget error.

Problem. The $D = 8$ piecewise-constant evaluator contains a base region, two commonly visited regions, and one rare region. Their scalar losses are 0, 1, 0.8, and 30, respectively. The two common switching directions are at standardized distance 0.8, and the rare direction is at distance 2.878. Four estimators are evaluated: the unprojected estimator (Full), a hard rank-two PCA projection learned from the structural pilot (Hard rank-2), a soft rule whose orthogonal-complement coefficient is calibrated from 64 disjoint groups (RCSP), and the same soft rule with its coefficient selected from an independent 4096-group reference bank (Oracle soft).

Design. The exact rare-region probability is 0.0020010, so a 256-perturbation structural pilot misses it with probability 0.598831. Thirty-two pilots are generated before examining their outcomes. Seventeen miss the rare region and fifteen observe it.

Every reported MSE is evaluated on a fresh bank untouched by the structural pilot, RCSP calibration, and oracle reference. The rare-missed and rare-seen strata are both retained.

Metrics. The structural-pilot audit reports common-subspace retention, rare-direction retention, and retained population signal. For each evaluation-group size $K \in \{4, 8, \ldots, 4096\}$, fresh evaluation groups estimate the gradient MSE of all four estimators. The calibrated and oracle soft-retention coefficients are also reported.

Table 8: Structural-pilot diagnostics for finite-pilot subspace recovery. The rare-missed and rare-seen strata contain all generated pilots.

Stratum	Pilots	Mean rare count	Common-subspace retention	Rare-dir. retention	Signal retention
rare seen	15	1.47	0.500	0.729	0.219
rare missed	17	0.00	0.944	0.020	0.729

Results. When the rare region is missed, the learned rank-two subspace retains on average 94.4% of the two common directions but only 2.0% of the rare direction. When it is observed, the rank-two fit changes substantially: rare-direction retention rises to 72.9% while mean retention of the original two-dimensional common subspace falls to about 50%.

Table 9: Complete MSE curve conditional on the rare region being absent from the 256-sample structural pilot. Evaluation groups are fresh; K is the evaluation-group size rather than a matched total-query budget.

K	Full	Hard rank-2	RCSP	Oracle soft
4	9.04275	1.33152	1.33363	1.33137
8	3.96320	0.60874	0.60912	0.60842
16	1.98876	0.30306	0.30414	0.30241
32	0.98053	0.16242	0.16136	0.16110
64	0.50265	0.09619	0.09381	0.09354
128	0.27509	0.07014	0.06533	0.06515
256	0.12107	0.05069	0.04232	0.04206
512	0.06312	0.04216	0.02885	0.02869
1024	0.03005	0.03833	0.01881	0.01872
2048	0.01519	0.03628	0.01158	0.01154
4096	0.00781	0.03525	0.00675	0.00674

Table 10: Complete MSE curve conditional on the rare region being observed by the structural pilot. The same independent evaluation protocol is used in both pilot strata.

K	Full	Hard rank-2	RCSP	Oracle soft
4	9.04275	4.22979	4.27888	4.22781
8	3.96320	2.00611	2.09226	2.00157
16	1.98876	1.05072	1.04267	1.04199
32	0.98053	0.57574	0.56364	0.55946
64	0.50265	0.34514	0.31908	0.31721
128	0.27509	0.23598	0.19436	0.19400
256	0.12107	0.15773	0.09682	0.09641
512	0.06312	0.13046	0.05527	0.05520
1024	0.03005	0.11330	0.02783	0.02781
2048	0.01519	0.10606	0.01456	0.01456
4096	0.00781	0.10255	0.00766	0.00766

The full/hard mean MSE ordering in the rare-missed stratum first reverses at $K = 1024$. At $K = 4096$, Full and Hard rank-2 have MSEs 0.00781 and 0.03525 in the rare-missed stratum and 0.00781 and 0.10255 in the rare-seen stratum.

At $K = 4096$, calibrated-soft and oracle-soft MSEs are 0.00675 and 0.00674 in the rare-missed stratum and both are 0.00766 in the rare-seen stratum. The complete coefficients are reported in Table 11.

Table 11: Soft-retention coefficients in the finite-pilot subspace experiment. RCSP is calibrated from 64 groups disjoint from both the structural pilot and evaluation bank; oracle coefficients use a separate 4096-group reference bank.

K	RCSP (missed)	Oracle (missed)	RCSP (seen)	Oracle (seen)
4	0.015	0.004	0.082	0.019
8	0.012	0.010	0.179	0.045
16	0.030	0.019	0.085	0.086
32	0.027	0.039	0.167	0.166
64	0.062	0.068	0.307	0.268
128	0.141	0.138	0.434	0.440
256	0.264	0.246	0.594	0.613
512	0.393	0.383	0.756	0.752
1024	0.546	0.553	0.855	0.857
2048	0.714	0.712	0.923	0.924
4096	0.833	0.828	0.961	0.960

D.3.2 Curved Switching Boundaries

Purpose. This experiment tests whether a nearest flat-interface approximation predicts activation and gradient moments when the true switching boundary is curved.

Problem. We compare two actions $J_0(u) = 0$ and $J_1(u) = \|u\|^2 - R^2$, so action 1 is selected inside a ball.

Design. The smoothing scale is $\sigma = 1$ and the nearest standardized distance is fixed at $a = 1.5$ while

$$R/\sigma \in \{0.5, 0.75, 1, 1.5, 2, 4, 8, 16, 32\}.$$

Thus nearest distance is held fixed and only curvature changes.

Metrics. A direct Gaussian-weighted surface integral on 200,001 quadrature points provides the numerical reference. Eight independent Monte Carlo repetitions use 2^{20} samples each. The maximum relative change under the nested quadrature check is 6.62×10^{-8}. Averaged over the nine curvatures, the score estimator has 0.604% relative vector error and the finite-difference (FD) estimator with common random numbers (CRN) has 2.12% relative vector error.

Table 12: Population-gradient checks for curved switching boundaries. The boundary integral is the numerical reference; score and CRN finite difference use Monte Carlo, and the final column reports the nearest-plane approximation relative to the surface value.

R/σ	Surface $m_\|$	Score $m_\|$	Score rel. err.	Score cosine	FD rel. err.	Plane/surface
0.5	-0.008764	-0.008825	1.62%	0.999998	6.02%	14.778
0.75	-0.019942	-0.019971	0.48%	0.999999	3.51%	6.495
1	-0.031573	-0.031553	0.53%	0.999999	2.69%	4.102
1.5	-0.050378	-0.050496	0.70%	0.999997	1.35%	2.571
2	-0.063437	-0.063393	0.63%	0.999994	1.76%	2.042
4	-0.089913	-0.089987	0.31%	0.999998	1.43%	1.440
8	-0.107632	-0.107861	0.43%	0.999997	1.04%	1.203
16	-0.117993	-0.118186	0.40%	0.999997	0.50%	1.098
32	-0.123603	-0.123613	0.34%	0.999997	0.75%	1.048

Population-gradient results. Across the nine curvature values, score-estimator relative error ranges from 0.31% to 1.62%, and its cosine with the surface-integral reference is at least 0.999994. The nearest-plane gradient magnitude is $14.78\times$ the surface-integral value at $R/\sigma = 0.5$ and $1.0479\times$ at $R/\sigma = 32$.

Grouped design and metrics. The grouped score experiments use $K \in \{4, 8, 16, 32\}$ and six independent repetitions. Each repetition contains approximately 262,144 perturbations, subject to a minimum of 4096 groups. The exact spherical transition probability is computed from a noncentral χ^2 distribution; no Monte Carlo quantity is fitted back into the prediction. We report observed and exact activation, exact signal ratios, and empirical normal-variance, tangent-variance, and SNR ratios relative to the flat-interface formulas.

Grouped results. The mean absolute activation error across all 36 $(R/\sigma, K)$ settings is 2.28×10^{-3}.

Table 13: Activation error against the exact spherical occupancy law, averaged over the nine curvature values.

K	Mean absolute activation error
4	7.8371e-04
8	1.7810e-03
16	2.6534e-03
32	3.9102e-03

Table 14: Complete grouped curvature results for $K \in \{4, 8\}$. Exact spherical activation is computed from a noncentral χ^2 law. $S_{\mathrm{sph}}/S_{\mathrm{flat}}$ uses the exact population gradients; the normal variance, tangential variance, and SNR ratios use empirical grouped-score moments normalized by the flat-interface formulas.

K	R/σ	q_{obs}	q_{sphere}	$S_{\mathrm{sph}}/S_{\mathrm{flat}}$	$v_{\parallel}/v_{\parallel}^{\mathrm{flat}}$	$v_{\perp}/v_{\perp}^{\mathrm{flat}}$	$\mathrm{SNR}/\mathrm{SNR}_{\mathrm{flat}}$
4	0.5	0.0182	0.0183	0.005	0.076	0.021	0.086
4	0.75	0.0392	0.0389	0.024	0.182	0.053	0.183
4	1	0.0593	0.0591	0.059	0.294	0.092	0.282
4	1.5	0.0913	0.0912	0.151	0.470	0.177	0.432
4	2	0.1138	0.1140	0.240	0.579	0.259	0.536
4	4	0.1624	0.1621	0.482	0.781	0.487	0.729
4	8	0.1963	0.1963	0.691	0.891	0.682	0.857
4	16	0.2175	0.2173	0.830	0.945	0.834	0.923
4	32	0.2287	0.2290	0.911	0.974	0.910	0.961
8	0.5	0.0361	0.0362	0.005	0.076	0.012	0.089
8	0.75	0.0765	0.0762	0.024	0.184	0.035	0.188
8	1	0.1143	0.1147	0.059	0.300	0.067	0.283
8	1.5	0.1764	0.1742	0.151	0.483	0.147	0.429
8	2	0.2160	0.2150	0.240	0.592	0.225	0.534
8	4	0.2970	0.2980	0.482	0.792	0.455	0.729
8	8	0.3536	0.3541	0.691	0.895	0.667	0.856
8	16	0.3873	0.3873	0.830	0.954	0.812	0.923
8	32	0.4048	0.4055	0.911	0.979	0.898	0.961

Curvature distorts activation, signal, normal variance, tangent variance, and SNR at different rates. For example, at $R/\sigma = 8, K = 8$ normal variance is already 89.5% of the flat prediction while per-direction tangent variance is 66.7% of the flat prediction; the corresponding SNR ratio is 85.6%.

Dimension-sweep design and results. The dimension sweep fixes $a = 1.5, K = 8$ and evaluates $D \in \{3, 8\}$ using analytic noncentral-χ^2 probabilities only.

At $R/\sigma = 2$, the spherical/flat activation ratio is 0.506 in $D = 3$ but 0.0423 in $D = 8$. Even at $R/\sigma = 32$, the corresponding ratios are 0.954 and 0.845. Hence the nearest normalized distance does not determine finite-budget activation independently of boundary geometry and ambient dimension.

Result summary. Tables 9–16 report the complete projection and curvature sweeps, including both structural-pilot strata, all group sizes, all curvature values, and both tested dimensions.

Table 15: Complete grouped curvature results for $K \in \{16, 32\}$. The flat model uses the same nearest standardized distance $a = 1.5$ in every row. As above, the signal ratio is an exact population quantity, while variance and SNR ratios are empirical grouped estimates.

K	R/σ	q_{obs}	q_{sphere}	$S_{\text{sph}}/S_{\text{flat}}$	$v_{\parallel}/v_{\parallel}^{\text{flat}}$	$v_{\perp}/v_{\perp}^{\text{flat}}$	$\text{SNR}/\text{SNR}_{\text{flat}}$
16	0.5	0.0723	0.0712	0.005	0.080	0.008	0.087
16	0.75	0.1468	0.1467	0.024	0.186	0.026	0.188
16	1	0.2148	0.2163	0.059	0.297	0.054	0.289
16	1.5	0.3220	0.3180	0.151	0.489	0.130	0.428
16	2	0.3832	0.3837	0.240	0.601	0.208	0.531
16	4	0.5092	0.5072	0.482	0.790	0.437	0.734
16	8	0.5835	0.5828	0.691	0.898	0.661	0.854
16	16	0.6285	0.6247	0.830	0.956	0.809	0.922
16	32	0.6473	0.6466	0.911	0.972	0.891	0.967
32	0.5	0.1384	0.1373	0.005	0.078	0.006	0.090
32	0.75	0.2735	0.2718	0.024	0.187	0.022	0.189
32	1	0.3906	0.3858	0.059	0.309	0.049	0.280
32	1.5	0.5365	0.5349	0.151	0.482	0.121	0.436
32	2	0.6200	0.6202	0.240	0.606	0.199	0.528
32	4	0.7592	0.7571	0.482	0.799	0.432	0.728
32	8	0.8270	0.8260	0.691	0.897	0.648	0.859
32	16	0.8585	0.8591	0.830	0.956	0.806	0.922
32	32	0.8772	0.8751	0.911	0.973	0.899	0.964

Table 16: Exact dimension sweep at $a = 1.5, K = 8$. All probabilities are analytic noncentral-χ^2 values; no Monte Carlo is used in this table.

D	R/σ	$p_{\text{sphere}}/p_{\text{plane}}$	q_{sphere}	$q_{\text{sphere}}/q_{\text{plane}}$
3	0.5	0.0689	0.0362	0.0853
3	1	0.2263	0.1147	0.2700
3	2	0.4461	0.2150	0.5060
3	4	0.6475	0.2980	0.7013
3	8	0.7959	0.3541	0.8335
3	16	0.8892	0.3873	0.9117
3	32	0.9421	0.4055	0.9545
8	0.5	0.0000	0.0000	0.0000
8	1	0.0016	0.0008	0.0020
8	2	0.0339	0.0180	0.0423
8	4	0.1832	0.0938	0.2208
8	8	0.4279	0.2071	0.4874
8	16	0.6542	0.3006	0.7075
8	32	0.8089	0.3588	0.8445

D.4 SEMO ENERGY-KNAPSACK

Purpose. This experiment tests whether pilot activation evidence and gradient resolvability predict an independent future gradient across smoothing scales, training states, seeds, and query budgets in a discrete optimization problem.

Problem. SEMO uses the official PredOpt Energy-Knapsack benchmark and verifies the archived dataset hash before use. Each example contains 48 items and capacity 120. A candidate $u \in \mathbb{R}^{48}$ supplies the item coefficients used by the decoder; batched dynamic programming returns an exact feasible knapsack solution, and the evaluator returns the benchmark downstream loss of that solution.

Design. The fixed data split contains 550 training, 100 validation, and 139 held-out examples. A shared linear predictor is warm-started for 100 MSE epochs and then trained by Full-score updates for 100,000 requested solver calls with $K = 8$, $\sigma = 0.25$, Adam learning rate 10^{-3}, batch size 64, and gradient clipping at 5. Five checkpoints at fractions $\{0, 0.25, 0.5, 0.75, 1\}$ are frozen for each

of three seeds. Full-score denotes the unprojected grouped estimator in equation 72 after the fixed SEMO parameter-space map.

The prospective grid uses 12 fixed validation contexts,

$$\sigma \in \{0.001, 0.003, 0.01, 0.03, 0.1, 0.3, 0.8\}, \quad K = 8,$$
$$B \in \{128, 512, 2048\},$$

for $3 \times 5 \times 7 \times 3 = 315$ primary configurations. Each checkpoint–σ cell uses 256 pilot groups, 1024 reference groups, and 4096 future groups; at the largest budget this gives 16 non-overlapping future windows. Pilot, reference, and future streams are disjoint and the same pilot moments are reused across candidate budgets.

Metrics. The experiment reports activation calibration, gradient calibration, state diagnosis, and agreement between reference and future outcomes as defined in Appendix D.1. It also reports median future relative MSE and cosine similarity within each future state, sensitivity to the confidence level, regime counts by σ, B, and checkpoint, coverage as pilot size changes, and the smallest certified future budget.

Results: independent target and continuous calibration. The independent future grid contains 34 STARVED, 118 LOW-RESOLUTION, 118 GRADIENT-RESOLVABLE, and 45 UNRESOLVED configurations. The high-budget reference and future outcomes agree on 266 of the 270 configurations for which both labels are defined.

Table 17: Reference-versus-independent-future labels for the 270 SEMO configurations for which both are resolved. The agreement is $266/270 = 98.5\%$.

Reference \ future	Starved	Low-resolution	Gradient-resolvable
Starved	33	0	0
Low-resolution	1	115	0
Gradient-resolvable	0	3	118

Table 18: Independent future gradient quality by SEMO outcome class. Relative MSE is undefined when the reference signal estimate is nonpositive, hence the smaller finite count in the Starved row.

Future class	Count	Finite RelMSE	Median RelMSE	Median cosine
Starved	34	7	42.784	0.000
Low-resolution	118	118	5.207	0.313
Gradient-resolvable	118	118	0.216	0.907

Pilot loss-activation probability has MAE 0.0118 and Pearson correlation 0.9987 against the future group-pool estimate. Among 177 finite positive activation-evidence pairs, the Pearson/Spearman correlations of $\log \widehat{A}_B$ with $\log A_B^{\mathrm{fut}}$ are 0.994/0.994.

For gradient calibration there are two reporting sets. Across all 213 finite positive pairs of $1/\widehat{\Lambda}_B$ and future relative MSE, the Pearson/Spearman log correlations are 0.916/0.946; the median future/predicted ratio is 1.075 and 78.9% lie within a factor of two. Excluding the four future-Starved pairs leaves 209 operationally relevant non-starved points, for which the correlations are 0.923/0.945 and 79.9% lie within a factor of two. Neither set is restricted to configurations on which the confidence rule issues a label.

Results: confidence-aware state diagnosis. Coverage below is defined on the 270 independently future-resolved cases, whereas the unconditional issue rate over all 315 rows is lower because the 45 future-Unresolved rows remain in the denominator. The primary 90% margin rule issues 178 resolved-case diagnoses and makes no observed cross-state error.

The threshold-specific margin adds 38 issued diagnoses, or 14.1 percentage points of resolved-case coverage, without an observed cross-state error. Only six of the 315 primary configurations also receive the stronger NOISE-DOMINATED mechanistic certificate; all six have future Low-resolution outcomes.

Table 19: SEMO confusion counts for the primary 90% overall margin rule. Pilot UNRESOLVED outputs are reported in the final column.

Future \ pilot	Starved	Low-resolution	Gradient-resolvable	Unresolved
Starved	28	0	0	6
Low-resolution	0	67	0	51
Gradient-resolvable	0	0	83	35

Table 20: SEMO confusion counts for the ratio-confidence baseline at the same 90% overall pointwise level and on the same pilot samples.

Future \ pilot	Starved	Low-resolution	Gradient-resolvable	Unresolved
Starved	28	0	0	6
Low-resolution	0	37	0	81
Gradient-resolvable	0	0	75	43

Results: confidence sensitivity. The 85%, 90%, and 95% rows below are *overall* pointwise operating levels. Their corresponding activation/moment component levels are 92.5%, 95%, and 97.5%.

At 85%, two future GRADIENT-RESOLVABLE configurations are certified as LOW-RESOLUTION; no error occurs in the reverse direction. No issued error is observed at the primary 90% level or at 95%.

Results: regime breakdowns. The following tables report future regimes and certification results by smoothing scale, future budget, and trajectory checkpoint.

The future state shifts systematically with both smoothing and future budget. For example, $\sigma = 0.001$ produces no future Gradient-resolvable rows, while all 45 rows at $\sigma = 0.8$ are future Gradient-resolvable. Increasing B changes the Starved/Low-resolution/Gradient-resolvable counts from $21/56/19$ at $B = 128$ to $5/22/59$ at $B = 2048$.

Results: pilot-size ablation. A nested-prefix ablation reuses the same underlying pilot stream at the pilot sizes

$$M_0 \in \{32, 64, 128, 256, 512, 1024\}.$$

The stored ablation uses 500 multiplier replicates rather than the primary 1000, so its $M_0 = 256$ row is a cost–coverage ablation rather than an exact duplicate of the primary row.

Coverage rises steadily with pilot information while observed selective accuracy remains at least 99.4%. The ratio baseline improves with pilot size as well but has lower coverage than the margin rule at every stored pilot size.

Results: conditional budget selection. One pilot can be reused across the three candidate future budgets because $\widehat{S}$ and $\widehat{V}$ do not change with B. Choosing the smallest budget certified as GRADIENT-RESOLVABLE yields a certified budget in 38 of 105 checkpoint–σ cells.

All 38 selected budgets are independently future Gradient-resolvable. The mean selected budget is 623.2 calls, a 69.6% reduction relative to always using the largest candidate future budget 2048. This conditional comparison reports the selected future-update budget; pilot calls are accounted for separately.

D.5 UCI HAR FEATURE SELECTION

Purpose. This experiment applies the fixed pilot diagnosis to a different decoder, objective, dimension, group size, and budget grid. It evaluates prospective prediction and state diagnosis without changing the thresholds or confidence construction.

Problem. We use the official UCI HAR data, whose canonical training and test array shapes are $(7352, 561)$ and $(2947, 561)$. A candidate $u \in \mathbb{R}^{561}$ is decoded into the indices of its 30 largest

Table 21: SEMO certification summary at the 90% overall pointwise level. Margin and ratio inference use identical pilot samples and 1000 multiplier replicates.

Rule	Issued	Coverage	Sel. acc.	Starved rec.	Low-res. rec.	Grad-res. rec.
Ratio confidence	140	51.9	100.0	82.4	31.4	63.6
Margin D_B	178	65.9	100.0	82.4	56.8	70.3

Table 22: Complete SEMO sensitivity to the declared *overall pointwise* confidence level. Component confidence is $1-(1-c)/2$ by Bonferroni.

Overall	Component	Issued	Coverage	Sel. acc.	Starved rec.	Low-res. rec.	Grad-res. rec.
85%	92.5%	190	70.4	98.9	85.3	60.2	74.6
90%	95.0%	178	65.9	100.0	82.4	56.8	70.3
95%	97.5%	158	58.5	100.0	64.7	49.2	66.1

coordinates. The evaluator trains no model: it returns the validation log loss of a fixed probabilistic 7-nearest-neighbor classifier using only those features, with Laplace smoothing 0.5.

Design. A fixed stratified split takes 200 training and 100 validation examples from each of six classes. Standardization statistics are computed only from this training subset, and the official test partition is not queried during diagnosis. The vectorized GPU oracle is audited against an independent NumPy implementation. Because Top-30 decoding is invariant to a common score translation, every group gradient is projected onto $\mathbf{1}^{\perp}$ using the fixed projector in Appendix D.1.

Frozen states come from a Full-score trajectory, using the unprojected grouped estimator in equation 72, with 32,768 requested calls, $K = 16$, $\sigma = 0.50$, and Adam learning rate 3×10^{-2}. For each of three seeds, checkpoints are stored at fractions $\{0, 0.25, 0.5, 0.75, 1\}$. The diagnostic grid is

$$\sigma \in \{0.05, 0.10, 0.25, 0.50\},$$
$$K \in \{8, 16\},$$
$$B \in \{128, 512, 2048, 4096\},$$

for 480 primary rows. The primary pilot has 256 groups; nested ablations use 32, 64, 128, 256, 512, and 1024 pilot groups. Every cell uses 1024 reference and 4096 future groups.

Metrics. The same activation calibration, gradient calibration, state diagnosis, and agreement between reference and future outcomes defined in Appendix D.1 are reported. Additional tables give future-state quality, confidence sensitivity, counts by σ, K, B, and checkpoint, pilot-size ablations, and the smallest certified future budget.

Results: independent target and continuous calibration. All 480 primary HAR configurations receive an independently defined future state: 68 are STARVED, 321 are LOW-RESOLUTION, and 91 are GRADIENT-RESOLVABLE. The high-budget reference and independent future labels agree in 476 of 480 rows (99.2%).

The MAE between pilot and future group-level activation probabilities is 0.00449 and their Pearson correlation is 0.9993. Among the 112 rows for which both activation evidences are finite and positive, the Pearson/Spearman correlations of $\log \widehat{A}_B$ with $\log A_B^{\mathrm{fut}}$ are 0.991/0.994.

For gradient-quality prediction, 412 future Low-resolution or Gradient-resolvable rows have finite positive $1/\widehat{\Lambda}_B$ and future relative MSE. Their log Pearson/Spearman correlations are 0.976/0.981, the median future/predicted ratio is 0.989, and 94.9% lie within a factor of two. This calibration uses all finite non-starved pairs, not only rows on which the confidence rule issues a label.

Results: confidence-aware state diagnosis. At the pre-specified 90% overall pointwise level, the margin rule issues 425 of 480 diagnoses, for 88.5% coverage, and all 425 issued labels agree with the independent future state. The three recalls are 86.8%, 93.8%, and 71.4% for Starved, Low-

Table 23: SEMO primary margin diagnosis by smoothing scale. "S/L/G/U" are independent future counts for Starved, Low-resolution, Gradient-resolvable, and Unresolved. Coverage and accuracy use future-resolved configurations only.

σ	S/L/G/U	Coverage	Sel. acc.	Starved rec.	Low-res. rec.	Grad-res. rec.	Resolved	Total
0.001	25/4/0/16	75.9	100.0	84.0	25.0	–	29	45
0.003	9/19/0/17	53.6	100.0	77.8	42.1	–	28	45
0.01	0/32/4/9	58.3	100.0	–	65.6	0.0	36	45
0.03	0/32/10/3	47.6	100.0	–	62.5	0.0	42	45
0.1	0/20/25/0	55.6	100.0	–	65.0	48.0	45	45
0.3	0/11/34/0	71.1	100.0	–	36.4	82.4	45	45
0.8	0/0/45/0	95.6	100.0	–	–	95.6	45	45

Table 24: SEMO primary margin diagnosis and independent future-state counts by future query budget.

B	S/L/G/U	Coverage	Sel. acc.	Starved rec.	Low-res. rec.	Grad-res. rec.	Resolved
128	21/56/19/9	78.1	100.0	85.7	78.6	68.4	96
512	8/40/40/17	68.2	100.0	75.0	55.0	80.0	88
2048	5/22/59/19	50.0	100.0	80.0	4.5	64.4	86

resolution, and Gradient-resolvable, respectively. The ratio-confidence baseline uses the same pilot groups and multiplier replicates but covers only 82.9%.

The threshold-specific margin adds 27 issued diagnoses over ratio inference, or 5.6 percentage points of coverage, without an observed cross-state error. The stricter NOISE-DOMINATED mechanistic certificate is available in 241 primary rows; all 241 independently have future Low-resolution outcomes.

Results: confidence sensitivity. The same 85%, 90%, and 95% overall levels used in SEMO are evaluated without changing the rule.

Results: regime breakdowns. Unlike SEMO, HAR does not exhibit a monotone transition with smoothing scale: for example the future Gradient-resolvable count is 20, 33, 30, and 8 as σ increases from 0.05 to 0.50. The selective accuracy is 100% in each σ block.

Results: pilot-size ablation. The nested-prefix ablation uses the pilot sizes shown below. The primary $M_0 = 256$ row uses 1000 multiplier replicates; the non-primary ablation rows use 500.

Results: conditional budget selection. Because the same pilot moments are reused across future budgets, one can choose the smallest B certified as Gradient-resolvable within a fixed checkpoint–σ–K cell. A certified budget exists in 53 of 120 cells: 12 choose $B = 2048$ and 41 choose $B = 4096$. All 53 selected rows are independently future Gradient-resolvable. The mean selected budget is 3632.3 calls, an 11.3% reduction relative to always using 4096. As in SEMO, this conditional comparison reports the selected future-update budget and accounts for pilot cost separately.

Result summary. Across the 480 configurations, the log Pearson/Spearman correlations between $1/\widehat{\Lambda}_B$ and future relative MSE are $0.976/0.981$. The primary margin rule issues 425 diagnoses, corresponding to 88.5% coverage and 100.0% observed selective accuracy.

Table 25: SEMO primary margin diagnosis by frozen training checkpoint. The same diagnostic rule is applied without checkpoint-specific tuning.

Checkpoint	S/L/G/U	Coverage	Sel. acc.	Starved rec.	Low-res. rec.	Grad-res. rec.
0.00	2/22/26/13	64.0	100.0	0.0	54.5	76.9
0.25	3/26/24/10	60.4	100.0	100.0	46.2	70.8
0.50	7/21/25/10	67.9	100.0	71.4	57.1	76.0
0.75	8/24/20/11	69.2	100.0	87.5	62.5	70.0
1.00	14/25/23/1	67.7	100.0	92.9	64.0	56.5

Table 26: SEMO pilot-size ablation on the same 270 future-resolved configurations. Margin and ratio rules use identical nested pilot prefixes; the stored ablation uses 500 multiplier replicates.

M_0	Calls	Margin cov.	Margin acc.	Ratio cov.	Starved rec.	Low-res. rec.	Grad-res. rec.	Ratio acc.
32	256	17.4	100.0	10.4	29.4	15.3	16.1	100.0
64	512	35.9	100.0	22.6	61.8	30.5	33.9	100.0
128	1024	50.4	100.0	34.4	61.8	43.2	54.2	100.0
256	2048	66.3	99.4	52.2	82.4	57.6	69.5	100.0
512	4096	76.7	99.5	64.8	85.3	71.2	78.8	100.0
1024	8192	84.8	100.0	75.6	91.2	82.2	85.6	100.0

Table 27: Smallest certified SEMO future budget under the primary margin rule. A certified budget exists in 38 of 105 checkpoint–σ cells; all 38 selected rows are independently Gradient-resolvable. The comparison reports candidate future-update budgets and accounts for pilot cost separately.

Selected B	Cells	Future-resolvable	Share of selected
128	13	13	34.2%
512	19	19	50.0%
2048	6	6	15.8%
Total	38	38	100.0%

Table 28: Pre-specified HAR evaluation protocol.

Quantity	Value
Features / selected features	561 / 30
Black-box model	probabilistic 7-NN
Training / validation per class	200 / 100
Seeds / checkpoints	3 / 5
Diagnostic σ	0.05, 0.10, 0.25, 0.50
Group sizes K	8, 16
Future budgets B	128, 512, 2048, 4096
Primary pilot groups	256
Reference / future groups	1024 / 4096
Primary overall / component confidence	90% / 95%

Table 29: Reference-versus-independent-future labels on the 480 HAR primary configurations. Agreement is $476/480 = 99.2\%$.

Reference \ future	Starved	Low-resolution	Gradient-resolvable
Starved	68	2	0
Low-resolution	0	318	1
Gradient-resolvable	0	1	90

Table 30: Independent future gradient quality by HAR outcome class. Relative MSE is undefined for 48 of the Starved rows because the reference signal estimate is nonpositive.

Future class	Count	Finite RelMSE	Median RelMSE	Median cosine
Starved	68	20	105.566	0.000
Low-resolution	321	321	6.264	0.304
Gradient-resolvable	91	91	0.654	0.702

Table 31: HAR certification summary at the frozen 90% overall pointwise level. Margin and ratio inference use identical pilot samples and 1000 multiplier replicates.

Rule	Issued	Coverage	Sel. acc.	Starved rec.	Low-res. rec.	Grad-res. rec.
Ratio confidence	398	82.9	100.0	86.8	87.2	64.8
Margin D_B	425	88.5	100.0	86.8	93.8	71.4

Table 32: HAR confusion counts for the primary frozen margin rule. Pilot UNRESOLVED outputs are reported in the final column.

Future \ pilot	Starved	Low-resolution	Gradient-resolvable	Unresolved
Starved	59	0	0	9
Low-resolution	0	301	0	20
Gradient-resolvable	0	0	65	26

Table 33: HAR sensitivity to the pre-specified overall pointwise confidence level. Component confidence is $1 - (1 - c)/2$ by Bonferroni.

Overall	Component	Issued	Coverage	Sel. acc.	Starved rec.	Low-res. rec.	Grad-res. rec.
85%	92.5%	429	89.4	100.0	86.8	94.1	74.7
90%	95.0%	425	88.5	100.0	86.8	93.8	71.4
95%	97.5%	401	83.5	100.0	66.2	91.9	67.0

Table 34: HAR primary margin diagnosis by smoothing scale. "S/L/G/U" are independent future counts for Starved, Low-resolution, Gradient-resolvable, and Unresolved.

σ	S/L/G/U	Coverage	Sel. acc.	Starved rec.	Low-res. rec.	Grad-res. rec.	Total
0.05	40/60/20/0	90.8	100.0	85.0	93.3	95.0	120
0.10	28/59/33/0	84.2	100.0	89.3	89.8	69.7	120
0.25	0/90/30/0	85.0	100.0	–	92.2	63.3	120
0.50	0/112/8/0	94.2	100.0	–	97.3	50.0	120

Table 35: HAR primary margin diagnosis by group size.

K	S/L/G/U	Coverage	Sel. acc.	Starved rec.	Low-res. rec.	Grad-res. rec.	Total
8	34/163/43/0	83.8	100.0	73.5	92.6	58.1	240
16	34/158/48/0	93.3	100.0	100.0	94.9	83.3	240

Table 36: HAR primary margin diagnosis and independent future-state counts by future query budget.

B	S/L/G/U	Coverage	Sel. acc.	Starved rec.	Low-res. rec.	Grad-res. rec.	Total
128	22/98/0/0	97.5	100.0	95.5	98.0	–	120
512	18/102/0/0	99.2	100.0	94.4	100.0	–	120
2048	14/83/23/0	81.7	100.0	100.0	86.7	52.2	120
4096	14/38/68/0	75.8	100.0	50.0	81.6	77.9	120

Table 37: HAR primary margin diagnosis by frozen Full-score trajectory checkpoint.

Checkpoint	S/L/G/U	Coverage	Sel. acc.	Starved rec.	Low-res. rec.	Grad-res. rec.	Total
0.00	0/72/24/0	90.6	100.0	–	100.0	62.5	96
0.25	0/72/24/0	83.3	100.0	–	87.5	70.8	96
0.50	2/65/29/0	93.8	100.0	50.0	98.5	86.2	96
0.75	30/59/7/0	89.6	100.0	90.0	93.2	57.1	96
1.00	36/53/7/0	85.4	100.0	86.1	88.7	57.1	96

Table 38: HAR pilot-size ablation on the same 480 future-resolved configurations. Margin and ratio rules use identical nested pilot prefixes.

M_0	Margin cov.	Margin acc.	Ratio cov.	Ratio acc.	Starved rec.	Low-res. rec.	Grad-res. rec.	Calls ($K = 8/16$)
32	50.8	100.0	45.4	100.0	39.7	63.2	15.4	256/512
64	64.0	99.7	54.6	100.0	51.5	75.4	31.9	512/1024
128	75.8	99.7	67.1	100.0	63.2	85.7	49.5	1024/2048
256	88.5	100.0	82.9	100.0	86.8	93.8	71.4	2048/4096
512	92.9	100.0	88.1	100.0	97.1	95.3	81.3	4096/8192
1024	96.7	99.6	92.5	100.0	97.1	97.2	92.3	8192/16384