PARETOTRACE: MINIMAX RECOVERY OF MGDA DIRECTIONS FROM PARTIAL GRADIENT GEOMETRY

Anonymous authors
Paper under double-blind review

ABSTRACT

Multiple gradient descent (MGDA) finds a direction that descends on every objective at once. It needs the Gram matrix of task gradients, and building that matrix exactly is the bottleneck as objectives multiply, so a prevailing response is to substitute cheaper geometry. We ask what the observed geometry determines, and answer exactly. Observing Gram-vector products leaves a family of consistent Grams, and we characterize which quantities are identical across it. The MGDA objective is not: it can only be narrowed to a computable range. The worst-task descent margin is, at any weighting the observations already span, and reading it costs no action beyond those already taken. Two methods follow. ParetoTrace-Direct contracts every parameter block, leaving one consistent Gram and returning the exact direction with no queries. It is $14.9\times$ faster than serial task gradients at 32 objectives, where four of six external implementations exhaust an NVIDIA A100 40 GB card. ParetoTrace-Stream buys products one at a time and stops once the range certifies. It is $1.97\times$ faster than contraction where materializing the Gram exhausts the card, and has lower worst-task loss and higher held-out accuracy than the strongest cheap baseline on all three seeds. That baseline lets one task run away to a loss of 41.0 while ours stays under 8.1. In replay on a training run it drives, the sharp endpoint certifies 37 of 100 updates against the cheaper rule's 27, never later, while a numerical guard limits the live run to 19, and a line search on its exact responses keeps every task loss non-increasing on its minibatch at 60 of 60 updates. Used as an instrument, the characterization shows that MGDA-UB (Sener & Koltun, 2018), the standard cheap substitute, increases at least one task loss at forty CelebA objectives on every trunk and batch size we tested.

1 INTRODUCTION

Deep networks are increasingly trained against many objectives at once. When those objectives disagree, descending their average loss can raise an individual objective's loss. Multiple gradient descent (Fliege & Svaiter, 2000; Désidéri, 2012) prevents this. At any nonstationary point it returns a direction along which *every* objective descends, and it is now a standard tool in deep multitask learning (Sener & Koltun, 2018).

What makes MGDA expensive is that it needs the geometry of the gradients, not just the gradients. Write $g_i = \nabla L_i(\theta) \in \mathbb{R}^P$ for M objectives, $G = [g_1, \dots, g_M]$, and $\Delta_M = \{\lambda \geq 0 : \mathbf{1}^\top \lambda = 1\}$. MGDA solves

$$q^\star = \min_{\lambda \in \Delta_M} q(\lambda), \qquad q(\lambda) = \lambda^\top K \lambda, \qquad K = G^\top G, \tag{1}$$

and steps along $-G\lambda^\star$. Forming K exactly means materializing a $P \times M$ table of per-objective gradients, and P is the size of the model. At forty objectives on a modest trunk that table is larger than the network it describes.

The field's response is to substitute cheaper geometry, and to validate the substitution by downstream accuracy. Downstream accuracy is the right end goal, but it cannot report geometric fidelity, so success on one benchmark does not predict the next task set. We ask the question that precedes the substitution. Given the gradient geometry one has actually observed, what is *determined* about the MGDA direction?

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.

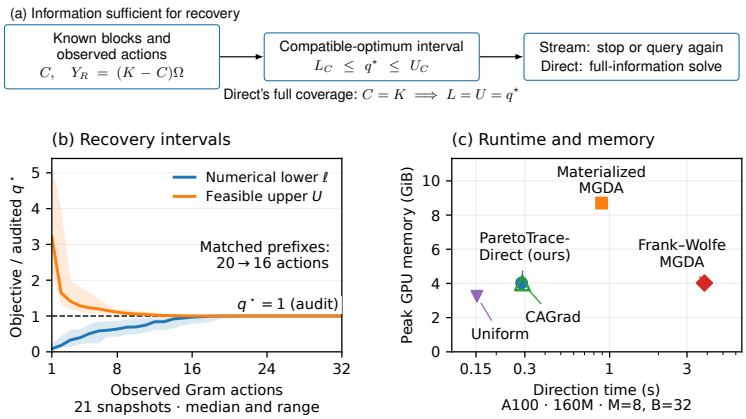

Figure 1: **From observed geometry to an MGDA direction.** (a) Actions and contracted blocks feed one recovery model, and full block coverage gives ParetoTrace-Direct's exact endpoint. (b) Completion intervals on 21 dependent clean-answer snapshots, median and full range, normalized by the full-Gram audit value. At the 5% target median stopping falls from 20 (FW lower) to 16 actions (completion lower) on identical uniform-first paths, and the pure-coordinate path gives 21 and 18 (Appendix B.4). U is the achieved feasible objective, so $U \ge U_{\mathrm{sk}}$. The panel's ℓ is the plotted numerical lower endpoint, not the composite ℓ_{c} of Table 1. (c) Complete-direction cost ablation on one held-out batch (Pythia-160M/A100, $M = 8$), medians of three warm repetitions on one A100-SXM4-40GB.

The question has an exact, constructive answer. Observing Gram-vector products leaves a family of consistent Grams, and the quantities that take one value across it split the problem in two. The MGDA objective can only be narrowed to a computable range. The worst-task descent margin, at any weighting the observations already span, is known exactly and costs no action beyond those already taken, though re-reading it after a parameter update costs one. Two methods follow from that one fact, and Table 1 reports what each buys.

Contributions.

- **An exact characterization.** The consistent Grams take the classical shorted-operator form $K_{\mathrm{sk}} + \{B \succeq 0 : B\Omega = 0\}$, their optima fill $[L_{\mathrm{sk}}, U_{\mathrm{sk}})$ when $L_{\mathrm{sk}} < U_{\mathrm{sk}}$ and the single point $q^\star$ otherwise, and that width is the exact conditional minimax regret (Theorem 1), and it stays sharp under a measured trace budget (Proposition 4). A functional is determined iff its off-span block vanishes (Theorem 2).
- **The width belongs to the transcript.** Any algorithm correct at every consistent Gram and reading only this transcript inherits the same width (Theorem 3), so it is not an artifact of our stopping rule. It is an indistinguishability argument, not a query lower bound (Proposition 2).
- **Two methods and a rule for choosing between them.** ParetoTrace-Direct is exact and, at $M = 32$, $14.9\times$ faster than the fastest external exact route that still fits the card, four of six having run out of memory. At $M = 8$, where all six fit, it leads by $1.49\times$. ParetoTrace-Stream certifies at a median 50 actions of 2048, beating contraction by $1.97\times$ where materialization runs out of memory.
- **A margin that costs no new observation.** Any weighting can be audited with one extra action, so the advantage is not that others cannot check their direction. It is that for an in-span weighting the margin is already determined, and a certified step at an admissible size then raises no task loss (Proposition 1).
- **An audit of deployed practice.** MGDA-UB increases a task loss at forty CelebA objectives on $3/3$ trunks at $3/3$ batch sizes (Appendix C.9; Table 17 tabulates batch 512).

Table 1: ParetoTrace against baselines. Each block is one protocol, stated in its header, and we never compare across blocks. Bold marks the best result in its block. Blocks 2 and 4 run both ParetoTrace methods under one protocol. Block 1's exact arm is MGDA on forty separate backward passes, not ParetoTrace-Direct's contraction, and its Stream row is an action count from the 400-step geometry study, not the 1200-step run. Block 6 runs every arm on one protocol, and only block 7 stops on L_{sk}. Block 2's trained arms stop on the cheaper $\ell_{\mathrm{c}} = \max\{0, 2m - U, m_+^2/U\}$ under a margin gate, which Corollary 1 shows is dominated by L_{sk} (Appendix B.5). On block 2's task set uniform and CAGrad reach lower test cross-entropy than either ParetoTrace method, and uniform also reaches lower worst-task cross-entropy on all three paired seeds, though only uniform fails the descent guarantee (Appendix C.1). ParetoTrace-Direct needs no queries, so its action count is zero by construction.

Method	Guarantee	Result
(1) CelebA-40, M=40: worst-attribute accuracy after 1200 steps		
Uniform	none	0.6541
PCGrad	none	0.6547
CAGrad	none	0.6702
MGDA-UB	none	0.6734
ParetoTrace-Stream	5% target	certifies at 33 of 40 actions
Exact MGDA	**exact**	**0.6778**
(2) Pythia-160M, 500 updates: worst-task gain of the returned direction, not a trained outcome		
Uniform	none	-0.420 (raises a task loss)
CAGrad	none	0.108
ParetoTrace-Stream	$U/\ell_{\mathrm{c}} \le 1.05$	0.983
ParetoTrace-Direct	**exact**	**0.9999**
(3) Pythia-160M, M=32: cost of one exact MGDA direction		
Serial task gradients	exact	12.803 s, 30.87 GiB
Chunked vmap	exact	13.527 s, 31.96 GiB
TorchJD	exact	out of memory
ParetoTrace-Direct	**exact**	**0.857 s, 14.21 GiB**
(4) Synthetic panels, M=2048: both methods, one protocol		
Materialized	exact	out of memory
ParetoTrace-Direct	exact	8.364 s, 6,993 MB, no queries
ParetoTrace-Stream	$\mathbf{U/\ell_{\mathrm{c}} \le 1.05}$, **9/9**	**4.252 s, 6,998 MB**, 50 actions (median)
(5) Pythia-160M, M=32, 100 updates: range over three seeds, ours better on 3/3		
MGDA-FA (Zhang et al., 2025)	none	worst loss 9.7–41.0, acc 0.668, 0.167 s
Stream, periodic	**rule not recorded**	**worst loss 4.7–8.0, acc 0.691**, 0.206 s
(6) Pythia-160M, M=16, 60 equal updates, one seed: held-out worst-task loss		
Uniform	none	17.38
FAMO	none	7.60
ParetoTrace-Direct	exact	15.22
ParetoTrace-Stream	5% certificate	6.40
ParetoTrace-Stream, certified step	**no task loss rises, 60/60**	77.48
(7) 103 recorded Grams, projected-gradient queries: actions to certify a 5% objective target		
Frank–Wolfe gap	$U/(U - g_{\mathrm{FW}}) \le 1.05$	22 actions
ParetoTrace-Stream	$\mathbf{U/L_{\mathrm{sk}} \le 1.05}$	**11 actions**
ParetoTrace-Direct	exact	0 actions, full contraction

2 BACKGROUND

A *Gram action* is a single matrix-vector product

$$Kv = G^\top (Gv), \tag{2}$$

which costs one weighted-gradient sweep and one task-loss JVP sweep. It never forms G. Suppose we have spent k of them, collected as queries Ω and responses $Y = K\Omega$. We call the pair (Ω, Y) an *action sketch*.

Many positive semidefinite Grams reproduce a given sketch, and they need not share an MGDA optimum. Write $\mathcal{S} = \mathrm{range}(\Omega)$ and $W = \Omega^\top Y$. The *shorted completion* and the two simplex

programs we use throughout are

$$K_{\mathrm{sk}} = YW^{\dagger}Y^{\top}, \tag{3}$$

$$L_{\mathrm{sk}} = \min_{\mu \in \Delta_M} \mu^{\top} K_{\mathrm{sk}} \mu, \tag{4}$$

$$U_{\mathrm{sk}} = \min_{\mu \in \Delta_M \cap \mathcal{S}} \mu^{\top} K_{\mathrm{sk}} \mu. \tag{5}$$

Equation (3) is the shorted operator of Anderson & Trapp (1975). Positive semidefinite consistency gives $W \succeq 0$ and $\ker W \subseteq \ker Y$. For a candidate Gram H we write $q_H^{\star} = \min_{\lambda \in \Delta_M} \lambda^{\top} H \lambda$.

Throughout, *effective rank* means the participation ratio $(\sum_i \lambda_i)^2 / \sum_i \lambda_i^2$ of the Gram spectrum. It is 1 for a rank-one Gram and M for a flat one, and it is how we index geometry from agreeing to conflicting.

We condition on the current step's data and on the queries already issued. Queries may depend on earlier responses and on randomness drawn independently of the Gram. Lemma 1 shows that every compatible Gram then reproduces the same transcript, which is what makes conditioning on it legitimate. This is the step that separates our setting from worst-case matrix-vector query bounds.

3 Related Work

Cheaper geometry for MGDA. MGDA-UB replaces the shared-trunk Gram with gradients taken with respect to the shared representation (Sener & Koltun, 2018), which is the substitution this paper adjudicates. PSMGD reuses weights across updates (Xu et al., 2025), and Zhang et al. (2025) analyze convergence under generalized smoothness. Each trades geometric fidelity for cost without reporting how much was given up, which our characterization measures.

Alternative update criteria. PCGrad projects conflicting gradients, CAGrad constrains the average descent, and Nash-MTL, FairGrad and FAMO pursue bargaining, fairness and loss-balancing criteria (Yu et al., 2020; Liu et al., 2021; Navon et al., 2022; Ban & Ji, 2024; Liu et al., 2023). These solve different problems rather than approximate MGDA, so they are baselines here rather than targets of the audit. Unitary scalarization remains a strong control (Kurin et al., 2022; Xin et al., 2022) and we include it.

Efficient Gram and Jacobian products. Structured NTK contractions (Novak et al., 2022) and ghost-gradient products (Wang et al., 2025) compute Gram entries without materializing per-objective gradients, and ParetoTrace-Direct uses that arithmetic. Our contribution is not the contraction but the model that makes full coverage an endpoint of the same recovery problem.

Stopping rules for matrix-free MGDA. Frank–Wolfe supplies the duality gap that matrix-vector MGDA solvers conventionally stop on (Jaggi, 2013), with Wolfe's minimum-norm-point algorithm and nonnegative conjugate-gradient variants as the standard alternatives (Wolfe, 1976; Schmelzer & Stoll, 2026; Liang, 2020). Our lower endpoint dominates that gap pointwise, so it substitutes wherever the gap is used and can only stop earlier, which is what block 7 of Table 1 measures.

What is classical here. Our ingredients are standard: the shorted-operator completion (Anderson & Trapp, 1975; Halko et al., 2011), positive semidefinite matrix completion (Laurent, 2009), and optimal recovery (Micchelli et al., 1976; Traub et al., 1988; Raginsky & Rakhlin, 2011). The matrix-vector query literature bounds what any query scheme can learn (Sun et al., 2019; Braverman et al., 2020; Meyer et al., 2021; Yeon et al., 2025), and probabilistic numerics treats the induced posterior (Hennig et al., 2015). Those bounds are worst-case over a class of matrices, whereas ours is conditional on the sketch actually realized, which turns it from a complexity statement into a stopping test. Prior work approximates K and solves. We keep the consistent family and return the weighting with least regret over it.

4 What an Action Sketch Determines

4.1 The compatible Grams, and the range they admit

Our first result describes the whole set of Grams consistent with a sketch, and the range of MGDA optima inside it.

Theorem 1 (Minimax recovery conditional on a realized sketch)**.** *Suppose* $\mathcal{V}(\Omega, Y) = \{H \succeq 0 : H\Omega = Y\}$ *is nonempty. Then*

$$\mathcal{V}(\Omega, Y) = \{K_{\mathrm{sk}} + B : B \succeq 0,\ B\Omega = 0\}, \tag{6}$$

and consequently

$$\inf_{H \in \mathcal{V}(\Omega,Y)} q_H^\star = L_{\mathrm{sk}}, \qquad \sup_{H \in \mathcal{V}(\Omega,Y)} q_H^\star = U_{\mathrm{sk}}, \tag{7}$$

with $U_{\mathrm{sk}} = +\infty$ *when* $\Delta_M \cap \mathcal{S}$ *is empty. If* $\Delta_M \cap \mathcal{S} \neq \varnothing$*, then any*

$$\lambda_{\mathrm{sk}} \in \operatorname*{arg\,min}_{\lambda \in \Delta_M \cap \mathcal{S}} \lambda^\top K_{\mathrm{sk}} \lambda \tag{8}$$

has objective U_{sk} *under every consistent Gram, and attains the minimax additive regret*

$$\inf_{\lambda \in \Delta_M} \sup_{H \in \mathcal{V}(\Omega,Y)} \left(\lambda^\top H \lambda - q_H^\star\right) = U_{\mathrm{sk}} - L_{\mathrm{sk}}. \tag{9}$$

If $L_{\mathrm{sk}} > 0$ *it is also minimax for the relative factor, which equals*

$$\inf_{\lambda \in \Delta_M} \sup_{H \in \mathcal{V}(\Omega,Y)} \frac{\lambda^\top H \lambda}{q_H^\star} = \frac{U_{\mathrm{sk}}}{L_{\mathrm{sk}}}. \tag{10}$$

The proof is in Appendix A.2. One consequence is easily inverted. The unrestricted minimizer of the lower program is *not* the weighting we return. We return the restricted one. A weighting outside $\mathcal{S}$ has a component we have never observed, and a positive semidefinite residual can inflate that component's objective without changing any response we have seen. A fully corrective Frank–Wolfe method computes the same minimizer.

4.2 Which quantities the sketch pins down

Theorem 1 pins one number to a range. It does not say which other quantities are pinned exactly, and the answer drives the rest of the paper. Write $\Pi = P_{\mathcal{S}^\perp}$ for the projector onto the unobserved subspace, distinct from the parameter count P, and $\langle A, H \rangle = \mathrm{tr}(AH)$.

Theorem 2 (Determined functionals)**.** *For symmetric* A*, the value* $\langle A, H \rangle$ *is the same at every* $H \in \mathcal{V}(\Omega, Y)$ *if and only if* $\Pi A \Pi = 0$*, and then equals* $\langle A, K_{\mathrm{sk}} \rangle$*. Hence for an incomplete sketch the response* $H\lambda$ *and the form* $\lambda^\top H \lambda$ *are exact for* $\lambda \in \mathcal{S}$ *and for no other* λ*, while* $\mathrm{tr}(H)$ *is exact for no incomplete sketch at all.*

The proof takes three lines (Appendix A.5). The forward direction is linearity. The converse is a limit, not an observation: no estimator recovers a functional with a nonzero off-span block, because two compatible Grams disagree on it.

Two consequences orient everything that follows. First, the trace is not determined, so the one extra scalar we use in Section 5 has to be bought. Second, at an in-span weighting $\lambda = \Omega c$ every compatible Gram returns the same Yc. The worst-task margin $m(\lambda) = \min_i (Yc)_i$ is therefore *exact*, while $q^\star$ is known only to lie in a range. This is one fact, not two: $q^\star$ escapes only because its minimizer may leave $\mathcal{S}$.

4.3 The width is a property of the transcript, not of our rule

When $\Delta_M \cap \mathcal{S} \neq \varnothing$, a weighting with regret at most η at every compatible Gram exists if and only if $U_{\mathrm{sk}} - L_{\mathrm{sk}} \leq \eta$. The converse is a statement about every algorithm, not a property of ours.

Algorithm 1 ParetoTrace. Lines 1–4 are ParetoTrace-Direct, lines 5–17 and the across-update rule are ParetoTrace-Stream. Inputs: objectives $L_1..L_M$, parameters θ, target ϵ, action budget $k_{\max}$.

1: **if** all parameter blocks can be contracted **then**
2: $C \leftarrow$ contract every block from activation–adjoint factors
3: **return** $\arg\min_{\lambda \in \Delta_M} \lambda^\top C \lambda$ ▷ ParetoTrace-Direct: exact, no action spent
4: **end if**
5: $\Omega \leftarrow [\mathbf{1}/M], Y \leftarrow K\Omega$ ▷ one action, keeps $\Delta_M \cap \mathcal{S}$ nonempty
6: **for** $k = 1$ to $k_{\max}$ **do**
7: $K_{\mathrm{sk}} \leftarrow YW^\dagger Y^\top$ with $W = \Omega^\top Y$
8: $L_{\mathrm{sk}} \leftarrow \min_{\Delta_M} \mu^\top K_{\mathrm{sk}} \mu$; $U_{\mathrm{sk}}, \lambda_{\mathrm{sk}} \leftarrow \min_{\Delta_M \cap \mathcal{S}} \mu^\top K_{\mathrm{sk}} \mu$
9: **if** $\min_i (Yc)_i > 0$ and descent is the target **then**
10: **return** λ_{sk} ▷ margin is exact by Thm. 2
11: **end if**
12: **if** $L_{\mathrm{sk}} > 0$ and $U_{\mathrm{sk}}/L_{\mathrm{sk}} \le 1 + \epsilon$ **then**
13: **return** λ_{sk} ▷ certified by Cor. 1
14: **end if**
15: append a new query to Ω and its response to Y
16: **end for**
17: **return** NOT CERTIFIED ▷ route to ParetoTrace-Direct

Across updates. Hold the returned λ. At each later update spend one action on $K\lambda$ and read $m = \min_i (K\lambda)_i$, which is exact by Thm. 2. Re-enter the loop above only when m falls below the target or the refresh cadence elapses.

Theorem 3 (Transcript-conditional necessity)**.** *Fix a query rule and let $\mathcal{W}$ be any family of positive semidefinite Grams that all produce the transcript (Ω, Y) under it. Let $\mathcal{A}$ map a transcript to a weighting, reading the queries, their responses, and any side information constant on $\mathcal{W}$, but nothing that separates two of its members. Suppose $\mathcal{A}$ is* correct at every instance*: for each $H \in \mathcal{W}$ it returns some $\lambda_H \in \Delta_M$ with $q_H(\lambda_H) - q_H^\star \le \eta$. Then*

$$\min_{\lambda \in \Delta_M} \sup_{H \in \mathcal{W}} \left(q_H(\lambda) - q_H^\star \right) \le \eta. \tag{11}$$

The hypothesis gives a separate guarantee at each instance, not one uniform guarantee. Indistinguishability collapses them onto a single weighting: every $H \in \mathcal{W}$ produces the same transcript, so $\mathcal{A}$ returns one λ for all of them, and the per-instance guarantees hold together there. Lemma 1 certifies that our spaces have that property under adaptive querying.

Two limits travel with this theorem. It prices certification, not accuracy, so a method that is merely accurate is not bound by it. That is why our own hindsight oracle stops $1.46\times$ sooner without contradiction. And it says nothing about which actions to buy. No theorem can: Proposition 2 shows the single query $\Omega = [\lambda^\star]$ drives both endpoints to $q^\star$, so no bound on certified stopping is a function of K alone.

5 PARETOTRACE

Algorithm 1 states both methods. They differ only in whether geometry is contracted or acquired, and they share one recovery model.

ParetoTrace-Stream: acquire, stop on the range, then hold and audit. Within an update, Stream issues actions and stops when the range certifies. Corollary 1 makes that stop valid: the recovered weighting has objective exactly U_{sk} under the true Gram, a relative stop $U_{\mathrm{sk}}/L_{\mathrm{sk}} \le 1 + \epsilon$ certifies $U_{\mathrm{sk}}/q^\star \le 1 + \epsilon$ when $L_{\mathrm{sk}} > 0$, and nested sketches give $L_s \le L_t \le U_t \le U_s$, so the certificate never moves backwards.

Across updates Stream holds the weighting rather than recomputing it, and Theorem 2 makes this safe. After parameters move, the held weighting's response costs one action to re-read, but the reading is *exact*: a single action settles whether the direction still descends, with no completion solve and no interval. Stream refreshes the weighting only when that check fails or a cadence elapses. Holding is therefore one action per update against a full recovery, not an approximation traded for speed.

The test is also never worse than the one MGDA solvers already run:

$$0 \;\le\; q_K(\lambda_{\mathrm{sk}}) - q_K^\star \;\le\; U_{\mathrm{sk}} - L_{\mathrm{sk}} \;\le\; g_{\mathrm{FW}}(\lambda_{\mathrm{sk}}), \qquad g_{\mathrm{FW}}(\lambda) = 2\big(q_K(\lambda) - \min_i (K\lambda)_i\big). \tag{12}$$

The sharp endpoint can replace the Frank–Wolfe gap wherever that gap is used to stop, and it can only stop earlier. Section 6 measures how much earlier.

ParetoTrace-Direct: contract every block. Contraction is the same information model under a different observation. Because $K = G^\top G$ and the rows of G split along a parameter partition, the Gram decomposes as $K = \sum_b K_b$ with every term $K_b \succeq 0$. Contracting a subset exactly gives $K = C + R$, where $R = \sum_{b \notin C} K_b$ is positive semidefinite *as a fact about the decomposition*, not as an assumption. That is the hypothesis Theorem 1 needs, so it applies to R and gives

$$\mathcal{V}_C = \{A + B : B \succeq 0,\ B\Omega = 0\}, \qquad A = C + Y_R(\Omega^\top Y_R)^\dagger Y_R^\top, \tag{13}$$

with residual responses $Y_R = R\Omega$ (Theorem 4). Retaining C can only raise the lower endpoint and never moves the upper one, so contracting never hurts. And when *every* block is contracted no residual term remains, so $R = 0$ is known rather than merely consistent: the version space is the singleton $\{C\}$, the range collapses to $[q^\star, q^\star]$, and the solve returns the exact direction with no query. That endpoint is ParetoTrace-Direct, and it is why a refusal routes there.

One measured scalar makes the model compact. Theorem 1 lets unseen curvature be arbitrarily large, which is what forces the span restriction, and buying the one functional Theorem 2 says is missing removes it. A norm sweep fixes $\mathrm{tr}(K)$, every compatible residual is then trace-bounded, and the compatible optima fill exactly $[L_{\mathrm{sk}}, U_\tau]$, one convex QP whose minimizer may leave the span (Proposition 4). On the 103 recorded Grams the exact interval is only a median 1.7% narrower, with off-span mass 0.0003, and the median stop does not change (Theorem 5 in Appendix A, measured in Appendix C.6).

Which method to use, and when. The range decides, and the routing test costs the Gram diagonal and one dense action rather than a full recovery (Appendix F). When the activation–adjoint factors fit, use ParetoTrace-Direct: at $M = 32$ it is $14.9\times$ faster than a serial per-objective backward loop (Appendix C.3), and on real CIFAR-10 geometry $1.41\times$ faster than materializing the Gram at $13.6\times$ less memory (Appendix C.7). When contracting every block costs more than a short sketch, use ParetoTrace-Stream, which at $M = 2048$ is $1.97\times$ faster than contraction where materialization exhausts the card, and amortize the weighting when it can serve several updates. When the geometry conflicts, which effective rank above roughly 3 identifies, Stream refuses the objective target: certify descent instead, at 7.5 to 10.7% of a full Gram, or fall back to Direct. At two or three objectives certify nothing. Every Direct output is checked against an independently materialized Gram to relative objective error 10^{-3}, and the factor arithmetic is Novak et al. (2022)'s.

Two targets to certify. Objective fidelity asks that q be within ϵ of $q^\star$. Local gain asks only that the direction descend on every task, and by Theorem 2 the recovered response $u = Yc$ answers it with no extra action: $m = \min_i u_i > 0$ establishes common descent, and $m \ge \gamma U$ retains a fraction γ of MGDA's worst-task slope. We report the margin normalized as $\min_i (K\lambda)_i / q(\lambda)$, which never exceeds 1 because $q(\lambda)$ is a convex combination of the $(K\lambda)_i$, equals 1 at exact MGDA where simplex stationarity gives $\min_i (K\lambda^\star)_i = q^\star$, and is negative exactly when the update raises some task's loss to first order. Experiments report both targets. Acquisition carries only the weakest guarantee: under a policy each of whose queries leaves the current span, each failure adds a direction and a full-rank sketch closes exactly (Corollary 1), so at most M can fail. A policy that re-queries a spanned prefix has no such bound, which is what Appendix C.7 measures. Savings below that are measured, not predicted.

The margin is worth stating as a guarantee on training, since it is the one thing a substituted geometry cannot supply.

Proposition 1 (A certified step raises no task loss)**.** *Let λ have margin $m = \min_i (K\lambda)_i > 0$ under the true Gram, and step along $d = -G\lambda$ with $q = \lambda^\top K\lambda$. If each L_i is β_i-smooth then*

$$L_i(\theta + \eta d) \le L_i(\theta) - \eta m + \tfrac{\eta^2 \beta_i}{2} q \quad \textit{for every } i, \tag{14}$$

so task i's loss does not rise for $0 \le \eta \le 2m/(\beta_i q)$. If every step of a run is certified and taken at $\eta \le 2m/(q \max_i \beta_i)$, admissible for all tasks at once, then $\max_i L_i(\theta_T) \le \max_i L_i(\theta_0)$.

The proof is two lines of the descent lemma (Appendix A.5.8). Three things separate it from validating a direction downstream. The margin is exact from the actions already taken, needing no completion and no lower endpoint (Theorem 2). It is evaluated under the true Gram, so it holds at the step rather than on average over a benchmark. Along certified steps it constrains each step on that step's minibatch, not one fixed objective: it says where a task cannot go, not where it lands. A method substituting an approximate Gram must spend an extra action to test m.

Queries share a fixed minibatch, so one positive semidefinite Gram explains all responses, and the certificate is conditional on that Gram rather than on the population one.

6 EXPERIMENTS

Three questions organize the seven blocks of Table 1. Does the cheap substitution survive more objectives, is our verdict on a sketch right, and what does acting on it buy? No single harness covers every regime, and we never combine numbers across studies. Three achievements are also distinct and we keep them apart: recovering the MGDA direction, certifying that a step descends on every task, and lowering held-out worst-task loss. The evidence here is strong for the first, conditional on hypotheses we do not enforce for the second, and mixed for the third.

6.1 DOES THE CHEAP SUBSTITUTION SURVIVE MORE OBJECTIVES?

MGDA-UB is the substitution this literature deploys (Sener & Koltun, 2018). It replaces the shared-trunk Gram with gradients taken with respect to the shared representation. We score it against the exact trunk Gram, alongside ParetoTrace-Direct. At two objectives MGDA-UB is near exact. At forty it returns a direction that *increases* at least one task loss, on every trunk and every batch size we ran. ParetoTrace-Direct descends on all forty, because full coverage leaves one consistent Gram and it solves the exact program on that.

The last column gives the mechanism. As objectives multiply the min-norm point collapses toward Pareto-stationarity, with $q^\star/\overline{K_{ii}}$ falling from 0.406 to 0.016. A surrogate that is merely close in angle can then land on the wrong side of a task. The sign is robust and reproduces on a fourth, independently trained trunk. The magnitude is not, ranging from -0.39 to -1.36, so we report the range rather than a point.

We report no result on CityScapes or QM9 and only a pilot on NYUv2. Over 398 SegNet updates on NYUv2 the certificate spends a median two actions of three, so its ceiling is one backward pass in three, and it quadruples the cost of the step it protects (Appendix C.10). Certification is not worth its price at three objectives.

This is one many-objective benchmark, and the min-norm collapse is what makes the loss of fidelity measurable, so a task set whose min-norm point stays away from zero need not show it. Block 1 of Table 1 is the downstream check: under five weighting rules sharing a learning rate, schedule and data order, exact MGDA is the most accurate arm on worst-attribute accuracy, though last on mean accuracy, by 2.4 points over uniform and roughly four times the run-to-run spread. Its 0.44-point margin over MGDA-UB sits inside MGDA-UB's own spread, so separating those two needs of order twenty seeds per arm (Appendix C.9.2).

6.2 IS THE VERDICT RIGHT?

The regime boundary transfers out of sample. A boundary fitted on synthetic data predicts real-data behaviour without refitting. A label-noise sweep at $M = 1024$ finds the block schedule certifying to effective rank 2.41 and failing from 3.15 (Appendix E), and real CIFAR-10 groups sit at 57.8 to 63.0, where the block schedule returns *not certified* on three seeds of three. Theorem 3 conditions on the realized transcript, so nothing could have stopped on *that* one, which is weaker than the information being absent: a coordinate order closes it at 227 actions of 512 (Appendix C.7).

The guarantee is conservative but not vacuous. The direction bound holds on all 103 Grams and overstates the realized squared error by a median factor of 33, so a 5% stop lands within a median $7.8 \times 10^{-4} q^\star$ of the exact direction (Appendix D). Contracted blocks narrow the range at fixed action count, as Theorem 4 requires.

The endpoint and the query policy are complements. In the full factorial of query policy against stopping rule, the sharp endpoint improves from 14 actions to 11 when queries move from coordinate to projected gradient, and from 22 to 11 against the Frank–Wolfe gap, improving 92 of 103 Grams. It is worth more once the policy is good, and it overturns our own worst case: the conflicting ResNet Gram stays open at 100 of 100 actions under a coordinate order and closes at 28 under projected gradient, a saving in actions that inverts in time because dense queries do not batch (Appendix C.8).

6.3 What does acting on the verdict buy?

Where acquisition pays outright. Block 4 runs both methods under one protocol. Near rank two, ParetoTrace-Stream certifies on all nine runs at $M = 512$, 1024 and 2048 at a median 50 actions rather than M, though one seed of three needs 244 over four passes and loses to contraction at $M = 1024$. Direct is the best route at $M = 512$, and beyond it materialization exhausts the card while Stream pulls ahead because contraction touches all M objectives (Appendix C.2).

Where objectives conflict, certify descent rather than the objective. On the real CIFAR-10 groups the objective closes at 227 actions of 512 while descent certifies at 10.7% and 7.5% of a full Gram, and uniform weights *raise* a task loss on all six. Retargeting acquisition onto the margin cuts the price to 5.5% and 4.5% (Appendix C.6.5), while on the agreeing panels uniform already descends on 101 of 102. Live on this geometry with acquisition fixed, completion stopping needs 54 against 79 actions for the squared-margin bound, better on 14 of 15 states, and descent recovery beats Direct on all 15 states at $M = 512$ but not at 1024 (Appendix G.3).

Against the strongest cheap baseline. MGDA-FA estimates MGDA weights from successive forward losses at constant gradient cost (Zhang et al., 2025). On Pythia-160M at 32 objectives over 100 updates, Stream's periodic variant has lower worst-task loss and higher token accuracy on all three seeds. The spread is the point: ours varies by $1.69\times$ across seeds and MGDA-FA's by $4.25\times$. MGDA-FA never checks its direction, so a task can ascend while the estimate is wrong. A positive margin is what guarantees a safe step: it bounds every task's directional derivative away from zero, so some step size lowers all of them, whereas with a negative margin no small step can (Proposition 1). These runs show the outcome of that asymmetry rather than the theorem, since they refresh every five updates and enforce no step-size bound. Appendix G reports the run that does. Our cadence was tuned on a development seed while MGDA-FA ran at its default, at 0.206 s per update against 0.167.

Which endpoint each experiment uses. The sharp endpoint drives block 7 in replay, and blocks 2 and 4 stop on the cheaper gated test (Table 3). It also drives one live training run, where a numerical guard limits it to 19 early stops: 100 updates of Pythia-160M at $M{=}16$, no false stop, $q/q^\star \le 1.0002$. Replaying all three rules on that run's own transcripts at identical queries, L_{sk} certifies 37 of 100 steps at a median 11 actions against the cheap rule's 27 at 15 and is strictly earlier on all 27 they share. Taking the step it licenses, no task loss rises on its minibatch at any of 60 updates against 48 of 60 at the tuned step, and audited uniform certifies as often at one action while reaching worst task 246.3 against 65.9, so what recovery buys is progress under the guarantee rather than the guarantee itself (Appendix G.1). Replay certificates survive a provable FP64 rounding radius (Appendix A.5.9).

7 Conclusion and Limitations

A functional of the Gram is determined precisely when its block on the unobserved subspace vanishes, and that one criterion places both methods. On identical transcripts L_{sk} certifies earlier than the cheap rule and, live, saves actions and time against it, though a numerical guard still limits its live stops. The 3.15 effective-rank boundary constrains one acquisition schedule, not certification, which closes at effective ranks 17 to 22 on CelebA. Open are standard benchmarks, the Nash-MTL, FairGrad, Aligned-MTL, MoCo and ghost-norm baselines, and evidence that fidelity improves downstream metrics.

REFERENCES

W. N. Jr. Anderson and G. E. Trapp. Shorted operators. II. *SIAM Journal on Applied Mathematics*, 28(1):60–71, 1975. doi: 10.1137/0128007.

Hao Ban and Kaiyi Ji. Fair resource allocation in multi-task learning. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 2715–2731, 2024. URL https://proceedings.mlr.press/v235/ban24a.html.

Mark Braverman, Elad Hazan, Max Simchowitz, and Blake Woodworth. The gradient complexity of linear regression. In *Conference on Learning Theory*, 2020.

Yifan Chen, Ethan N. Epperly, Joel A. Tropp, and Robert J. Webber. Randomly pivoted cholesky: Practical approximation of a kernel matrix with few entry evaluations. *Communications on Pure and Applied Mathematics*, 78(5):995–1041, 2025. doi: 10.1002/cpa.22234.

Jean-Antoine Désidéri. Multiple-gradient descent algorithm (MGDA) for multiobjective optimization. *Comptes Rendus Mathematique*, 350(5–6):313–318, 2012.

J. Fliege and B. F. Svaiter. Steepest descent methods for multicriteria optimization. *Mathematical Methods of Operations Research*, 51:479–494, 2000.

Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53 (2):217–288, 2011.

Philipp Hennig, Michael A. Osborne, and Mark Girolami. Probabilistic numerics and uncertainty in computations. *Proceedings of the Royal Society A*, 471, 2015.

Martin Jaggi. Revisiting Frank–Wolfe: Projection-free sparse convex optimization. In *International Conference on Machine Learning*, volume 28, pp. 427–435, 2013.

Vitaly Kurin, Alessandro De Palma, Ilya Kostrikov, Shimon Whiteson, and M. Pawan Kumar. In defense of the unitary scalarization for deep multi-task learning. In *Advances in Neural Information Processing Systems*, 2022. URL https://arxiv.org/abs/2201.04122.

Monique Laurent. Matrix completion problems. *Encyclopedia of Optimization*, 2009.

Youwei Liang. Gradient projection for solving quadratic programs with standard simplex constraints. *arXiv preprint arXiv:2006.06934*, 2020.

Bo Liu, Xingchao Liu, Xiaojie Jin, Peter Stone, and Qiang Liu. Conflict-averse gradient descent for multi-task learning. In *Advances in Neural Information Processing Systems*, 2021.

Bo Liu, Yihao Feng, Peter Stone, and Qiang Liu. FAMO: Fast adaptive multitask optimization. In *Advances in Neural Information Processing Systems*, 2023.

Raphael A. Meyer, Cameron Musco, Christopher Musco, and David P. Woodruff. Hutch++: Optimal stochastic trace estimation. In *Symposium on Simplicity in Algorithms*, 2021.

Charles A. Micchelli, Theodore J. Rivlin, and Shmuel Winograd. The optimal recovery of smooth functions. *Numerische Mathematik*, 26:191–200, 1976.

Aviv Navon, Aviv Shamsian, Idan Achituve, Haggai Maron, Kenji Kawaguchi, Gal Chechik, and Ethan Fetaya. Multi-task learning as a bargaining game. In *International Conference on Machine Learning*, 2022.

Roman Novak, Jascha Sohl-Dickstein, and Samuel S. Schoenholz. Fast finite width neural tangent kernel. In *International Conference on Machine Learning*, volume 162, pp. 17018–17044, 2022.

Maxim Raginsky and Alexander Rakhlin. Information-based complexity, feedback and dynamics in convex programming. *IEEE Transactions on Information Theory*, 57(10):7036–7056, 2011. doi: 10.1109/TIT.2011.2165796.

Thomas Schmelzer and Martin Stoll. Non-negative conjugate gradients. *arXiv preprint arXiv:2607.22121*, 2026.

Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *Advances in Neural Information Processing Systems*, 2018.

Xiaoming Sun, David P. Woodruff, Guang Yang, and Jialin Zhang. Querying a matrix through matrix-vector products. In *International Colloquium on Automata, Languages, and Programming*, 2019.

Joseph F. Traub, Grzegorz W. Wasilkowski, and Henryk Woźniakowski. *Information-Based Complexity*. Academic Press, 1988.

Jiachen T. Wang, Prateek Mittal, Dawn Song, and Ruoxi Jia. Data shapley in one training run. In *International Conference on Learning Representations*, 2025. URL https://arxiv.org/abs/2406.11011.

Philip Wolfe. Finding the nearest point in a polytope. *Mathematical Programming*, 11(1):128–149, 1976.

Derrick Xin, Behrooz Ghorbani, Justin Gilmer, Ankush Garg, and Orhan Firat. Do current multi-task optimization methods in deep learning even help? In *Advances in Neural Information Processing Systems*, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/580c4ec4738ff61d5862a122cdf139b6-Abstract-Conference.html.

Mingjing Xu, Peizhong Ju, Jia Liu, and Haibo Yang. PSMGD: Periodic stochastic multi-gradient descent for fast multi-objective optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 21770–21778, 2025. doi: 10.1609/aaai.v39i20.35482.

Kingsley Yeon and Mihai Anitescu. Beyond low rank: Fast low-rank + diagonal decomposition with a spectral approach. *arXiv preprint arXiv:2512.17120*, 2025.

Kingsley Yeon, Promit Ghosal, and Mihai Anitescu. BOLT: Block-orthonormal Lanczos for trace estimation of matrix functions. *arXiv preprint arXiv:2505.12289*, 2025.

Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. In *Advances in Neural Information Processing Systems*, 2020.

Qi Zhang, Peiyao Xiao, Shaofeng Zou, and Kaiyi Ji. MGDA converges under generalized smoothness, provably. In *International Conference on Learning Representations*, 2025.

Reproducibility statement. The appendices and artifact provide algorithms, proofs, protocols, seeds, raw results, timing procedures, stored audit Grams, and machine-checkable recomputation of the reported numerical-reference intervals. The replay rule and its consistency diagnostics are stated in the replay protocol. The bounded-completion replay of Appendix A.3.1 and the fidelity–outcome reduction behind the fidelity–outcome table are hashed scripts over the stored records and perform no new model runs.

AI use statement. AI tools assisted with proof drafts, literature discovery, code, numerical audits, and manuscript editing. The authors are responsible for the final arguments, attribution, implementation, and data provenance.

A PROOFS AND SUPPORTING THEORY

In this section we collect every proof omitted from the main text, together with the supporting results they rely on. Appendix B gives protocols and solver settings, and Appendix C reports the experiments referenced but not tabulated above.

A.1 STRUCTURED-INFORMATION RECOVERY: PROOF AND EXECUTION CONTRACT

Full statement.

Theorem 4 (Recovery with contracted parameter geometry)**.** *Define the residual version space* $\mathcal{V}_C = \{C + R : R \succeq 0,\ R\Omega = Y_R\}$. *Assume it is nonempty and put* $A = C + Y_R(\Omega^\top Y_R)^\dagger Y_R^\top$ *and* $\mathcal{S} = \mathrm{range}(\Omega)$. *Then*

$$\mathcal{V}_C = \{A + B : B \succeq 0,\ B\Omega = 0\}.$$

If $\Delta_M \cap \mathcal{S} \neq \varnothing$, *define* $L_C = \min_{\Delta_M} \lambda^\top A\lambda$ *and* $U_C = \min_{\Delta_M \cap \mathcal{S}} \lambda^\top A\lambda$. *A minimizer* λ_C *of the latter problem has exact conditional minimax additive regret* $U_C - L_C$ *and relative factor* U_C/L_C *when* $L_C > 0$. *Discarding* C *while retaining the same total actions* $K\Omega = C\Omega + Y_R$ *leaves* U_C *unchanged and can only decrease the lower endpoint. For every compatible true gradient matrix, its recovered direction satisfies* $\|G\lambda_C - G\lambda^\star\|^2 \leq U_C - L_C$. *When all parameter blocks are contracted, the coverage observation instead forces the singleton version space* $\{C\}$. *An unrestricted simplex solve then gives the full-information interval* $[q^\star, q^\star]$ *without residual actions.*

A.1.1 PROOF OF THEOREM 4

The observed C is a sum of Gram contributions from disjoint parameter coordinates, not an arbitrary low-rank fit, so the unobserved contribution is PSD. That condition is essential: subtracting an arbitrary diagonal or approximate factor model from K does not ensure a PSD residual. The information model retains C, the residual responses Y_R and the identity of fully covered blocks, and asserts sharpness only for that declared PSD space, not over every additional constraint available in cached neural factors.

Apply the PSD completion identity of Theorem 1 to R:

$$R = R_0 + B, \qquad R_0 = Y_R(\Omega^\top Y_R)^\dagger Y_R^\top, \qquad B \succeq 0, \quad B\Omega = 0.$$

Thus $H = A + B$, where $A = C + R_0 \succeq 0$, is exactly the unrestricted residual version space. In particular A is its smallest element in Loewner order, and $\inf_H q_H^\star = \min_{\Delta_M} \lambda^\top A\lambda = L_C$. For every $\lambda \in \mathcal{S}$, $B\lambda = 0$, so its objective is fixed. The restricted minimizer λ_C has objective U_C for every completion.

Let Π project onto $\mathcal{S}^\perp$ and take $H_t = A + t\Pi$. Its simplex optimum increases to U_C. To see this, choose minimizers μ_t and use $q_{H_t}(\mu_t) \leq U_C$ to obtain $t\|\Pi\mu_t\|^2 \leq U_C$. Compactness of the simplex gives subsequential limits in $\mathcal{S} \cap \Delta_M$; continuity of the A quadratic gives the limiting lower bound U_C. If the intersection is empty, its positive distance from $\mathcal{S}$ instead makes the optima diverge.

For a decision in $\mathcal{S}$, the worst additive regret is its fixed objective minus L_C, attained at $H = A$. This is minimized by λ_C. For a decision outside $\mathcal{S}$, its objective under H_t diverges while $q_{H_t}^\star \leq U_C$, so its worst regret is infinite. Dividing by $q_H^\star$ gives the identical minimax argument when $L_C > 0$: in-span worst ratios are attained at A, while off-span ratios diverge. This proves the sharp additive and relative statements, including singular residual cores.

If only total actions $Y_K = C\Omega + Y_R$ are retained, the admissible space is $\mathcal{V}(\Omega, Y_K)$. It contains $\mathcal{V}_C$. Its smallest completion K_{sk} therefore satisfies $K_{\mathrm{sk}} \preceq A$, and its simplex minimum is no larger than L_C. On $\mathcal{S}$, both matrices have the same quadratic $\lambda^\top K\lambda$, so their upper endpoints and restricted recovery decisions agree. The extra geometry sharpens the uncertainty about the true optimum; it need not change the restricted weights.

For any compatible G, write $d = G\lambda_C$ and $d^\star = G\lambda^\star$. Projection of zero onto the gradient convex hull implies $\langle d^\star, d - d^\star\rangle \geq 0$. Hence

$$\|d - d^\star\|^2 \leq \|d\|^2 - \|d^\star\|^2 \leq U_C - L_C.$$

Finally, full coverage is additional structural information: no parameter coordinates remain, so $R = 0$ is known. The feasible space is the singleton $\{C\}$, and ParetoTrace-Direct solves $\min_{\Delta_M} \lambda^\top C\lambda$ without any residual query. This holds even when C is singular or $q^\star = 0$. A zero residual response by itself does *not* establish full coverage.

A.1.2 A STRICT SEPARATION INDEPENDENT OF RESIDUAL RANK AND SCALE

Retaining C can change a stopping decision at arbitrarily high residual rank and arbitrary residual scale. With C alternating 1 and $3/2$ on the diagonal, $R = \tau(I - \mathbf{1}\mathbf{1}^\top/M)$ and the single action $\Omega = \mathbf{1}/M$, one has $L_C = 6/5M$, $U_C = 5/4M$ and $U_C/L_C = 25/24 < 1.05$, while discarding C gives $L_{\mathrm{sk}} = 4/5M$ and a factor $25/16 > 1.05$. The separation is independent of τ and of $\operatorname{rank} R = M - 1$. The residual is chosen orthogonal to the query to isolate the mechanism. A generic residual reveals part of itself through the action and gives a smaller separation.

A.1.3 ALGORITHMIC AND IMPLEMENTATION CORRESPONDENCE

ParetoTrace stores a parameter-coverage ledger and task-space information. Factor contractions add exact disjoint-block contributions to C; residual actions add columns to (Ω, Y_R). If new block geometry D is contracted after actions were collected, update $C \leftarrow C + D$ and $Y_R \leftarrow Y_R - D\Omega$, which preserves the total observed responses and restricts the version space to completions also consistent with the new known contribution, so the information-optimal lower endpoint cannot decrease. ParetoTrace-Direct contracts all blocks, ParetoTrace-Stream uses $C = 0$ and acquires actions, and a hybrid retains both.

The released implementation covers the known-geometry completion, the numerical endpoint solves and the full-coverage case, with tests for valid random PSD completions, preservation of observed responses, interval dominance, direction error, the high-rank separation and the singular and stationary full-coverage cases. These are numerical checks, not floating-point interval arithmetic. The general hybrid acquisition scheduler has not been benchmarked on A100: the reported ParetoTrace-Direct timings are for its already executed full-coverage path, and the structured theorem does not turn those into hybrid training results.

A.2 PROOF OF THEOREM 1

Take any $H \in \mathcal{V}(\Omega, Y)$ and let $A = H^{1/2}\Omega$. Then $Y = H^{1/2}A$ and $W = A^\top A$. The projector identity gives

$$YW^\dagger Y^\top = H^{1/2} P_{\mathrm{range}(A)} H^{1/2} \preceq H. \tag{15}$$

It also gives $K_{\mathrm{sk}}\Omega = Y$. Thus $B = H - K_{\mathrm{sk}} \succeq 0$ and $B\Omega = 0$. Conversely every $K_{\mathrm{sk}} + B$ with these properties is consistent. This proves Equation (6). It also proves the lower identity because K_{sk} is itself consistent.

Let $\mathcal{S} = \mathrm{range}(\Omega)$. If $\lambda \in \mathcal{S}$, then $B\lambda = 0$ for every consistent completion. Thus its objective is known exactly. Every consistent optimum is at most U_{sk}. To show sharpness, define

$$H_\eta = K_{\mathrm{sk}} + \eta P_{\mathcal{S}^\perp}.$$

Each H_η is consistent and

$$q^\star_{H_\eta} = \min_{\lambda \in \Delta_M} \left\{ \lambda^\top K_{\mathrm{sk}} \lambda + \eta \operatorname{dist}(\lambda, \mathcal{S})^2 \right\}.$$

Compactness shows that this value increases to U_{sk}. If the simplex and $\mathcal{S}$ are disjoint, their distance is positive and the value tends to infinity. This proves the upper identity.

When $L_{\mathrm{sk}} < U_{\mathrm{sk}} < \infty$, no finite completion attains the upper endpoint. If $H = K_{\mathrm{sk}} + B$ attained it, a minimizer $\bar{\lambda}$ of the known objective over $\Delta_M \cap \mathcal{S}$ would be a global minimizer for H. Since $B\bar{\lambda} = 0$, the gradients of the two quadratic forms agree at $\bar{\lambda}$. The simplex KKT conditions would make $\bar{\lambda}$ a global minimizer for K_{sk}, which forces $L_{\mathrm{sk}} = U_{\mathrm{sk}}$. Continuity of $q^\star_{H_\eta}$ in η then gives every value in $[L_{\mathrm{sk}}, U_{\mathrm{sk}})$. The same argument gives $[L_{\mathrm{sk}}, \infty)$ when the intersection is empty.

Now suppose that $\Delta_M \cap \mathcal{S}$ is nonempty. For any candidate $\lambda \in \mathcal{S}$, its objective is fixed across the version space, and K_{sk} attains the smallest compatible optimum L_{sk}. Its worst-case additive regret is

therefore $\lambda^\top K_{\mathrm{sk}}\lambda - L_{\mathrm{sk}}$. Minimizing this quantity over $\Delta_M \cap \mathcal{S}$ gives $U_{\mathrm{sk}} - L_{\mathrm{sk}}$. For any candidate outside $\mathcal{S}$, the objective under H_η diverges, while the optimum is at most U_{sk} because an in-span recovery point remains feasible. The worst-case additive regret is therefore infinite. This proves Equation (9).

Suppose that $L_{\mathrm{sk}} > 0$ and $\Delta_M \cap \mathcal{S}$ is nonempty. For any candidate in $\mathcal{S}$, its worst-case ratio equals its known objective divided by L_{sk} because K_{sk} attains the denominator. For any candidate outside $\mathcal{S}$, the numerator under H_η diverges while the optimum is at most U_{sk}. Its worst-case ratio is therefore infinite. Minimizing over all simplex candidates therefore selects a member of Equation (8) and gives Equation (10)'s exact factor $U_{\mathrm{sk}}/L_{\mathrm{sk}}$.

A.3 Bounded-completion extension

Let $P_{\mathcal{S}^\perp}$ denote the orthogonal projector onto the complement of the query span.

Theorem 5 (Additive recovery under bounded completions)**.** *For $\tau \geq 0$, define the bounded version space*

$$\mathcal{V}_\tau = \{K_{\mathrm{sk}} + B : 0 \preceq B \preceq \tau P_{\mathcal{S}^\perp},\ B\Omega = 0\}. \tag{16}$$

For any $\lambda \in \Delta_M$ and every $H \in \mathcal{V}_\tau$,

$$0 \leq q_H(\lambda) - q_H^\star \leq R_\tau(\lambda) := \lambda^\top K_{\mathrm{sk}}\lambda - L_{\mathrm{sk}} + \tau\|P_{\mathcal{S}^\perp}\lambda\|_2^2. \tag{17}$$

Hence

$$\lambda_\tau \in \operatorname*{arg\,min}_{\lambda\in\Delta_M} R_\tau(\lambda) \tag{18}$$

returns a finite additive guarantee for every nonempty sketch. If $R_\tau(\lambda_\tau) \leq \varepsilon_{\mathrm{abs}}$, then $q_H(\lambda_\tau) \leq q_H^\star + \varepsilon_{\mathrm{abs}}$ for all $H \in \mathcal{V}_\tau$. When $L_{\mathrm{sk}} > 0$, it also gives

$$\frac{q_H(\lambda_\tau)}{q_H^\star} \leq 1 + \frac{R_\tau(\lambda_\tau)}{L_{\mathrm{sk}}}. \tag{19}$$

Unlike Equation (8), the bounded decision may leave the query span when the reduction in the observed quadratic offsets the uncertainty penalty. A trace budget $\mathrm{tr}(B) \leq \rho$ implies the valid, more conservative choice $\tau = \rho$. Equation (17) is a computable uniform bound, not an exact minimax regret formula.

A.3.1 Measured value of the cheapest informative prior

Since $B \succeq 0$ and $B\Omega = 0$ force $B \preceq \mathrm{tr}(B)P_{\mathcal{S}^\perp}$, with $\mathrm{tr}(B) = \mathrm{tr}(K) - \mathrm{tr}(K_{\mathrm{sk}})$ and $\mathrm{tr}(K) = \|G\|_F^2$ available as one scalar from the same factor sweep, the exact residual trace is the tightest budget a practitioner can buy for one extra number. Replaying Theorem 5 with that τ on all 103 recorded Grams along both frozen paths, the bounded guarantee is never worse than $U_{\mathrm{sk}} - L_{\mathrm{sk}}$, and the bounded decision leaves the span on 51.8% of the 6,728 prefixes yet the median ratio of $R_\tau(\lambda_\tau)$ to $U_{\mathrm{sk}} - L_{\mathrm{sk}}$ is 0.9917. It leaves the span often and barely.

That comparison alone would not settle it, because R_τ is conservative for a reason the Loewner step does not cause: the step is tight for this functional, since $\rho\, xx^\top/\|x\|^2$ at $x = P_{\mathcal{S}^\perp}\lambda$ is feasible with trace ρ and attains what the relaxation gives. The looseness is the decoupling in Equation (17), which pairs $\sup_H q_H(\lambda)$ with $\inf_H q_H^\star$ while a residual inflating the numerator also raises the optimum. Measuring instead through the endpoint the prior moves, $U_\rho = \sup\{q_H^\star : B \succeq 0,\ B\Omega = 0,\ \mathrm{tr}\, B \leq \rho\}$, the inner minimum is concave in B, so projected supergradient ascent gives a rigorous lower bound on U_ρ while $U_\rho \leq U_{\mathrm{sk}}$ holds a priori. Over the 3,062 prefixes of positive width, those residuals already recover a median 87.1% of the unrestricted width, so the trace prior removes at most about 13% of it and changes the median action count on none of the eight panel–path cells. Evaluated exactly through Proposition 4 and Appendix C.6.1, rather than bounded, it removes a median 1.7% along the uniform-first coordinate paths. Once the residual is bounded, off-span decisions become admissible and the minimax regret is no longer the width of the compatible optima, so this bounds what the trace prior buys for this stopping rule rather than what an optimal bounded-residual decision could achieve.

Three numbers for the trace appear in this paper and they measure three different things. As the inequality $\mathrm{tr}\, B \leq \rho$ it lowers the upper endpoint and removes at most about 13% of the interval

width, which is the figure just given. As the equality $\operatorname{tr} B = \tau$ it leaves the lower endpoint essentially untouched, lifting L_{sk} by 10^{-17} to within float64 resolution, which on these normalized Grams is a measured zero rather than a number. And it shrinks the left side of Equation (11) by under 0.6% (Corollary 3), a smaller move than the width figure because that left side is a minimax regret over decisions free to leave the span rather than the width of the compatible optima. So the trace is informative about the upper endpoint alone, and on none of the three measures does it change a median action count.

A.3.2 Proof of Theorem 5

Fix $H = K_{\mathrm{sk}} + B \in \mathcal{V}_\tau$. Since $B \succeq 0$ and K_{sk} is the Loewner-smallest consistent Gram,

$$q_H^\star \geq \min_{\mu \in \Delta_M} \mu^\top K_{\mathrm{sk}} \mu = L_{\mathrm{sk}}.$$

The spectral completion bound gives, for every $\lambda \in \Delta_M$,

$$\lambda^\top B \lambda \leq \tau \lambda^\top P_{\mathcal{S}^\perp} \lambda = \tau \|P_{\mathcal{S}^\perp} \lambda\|_2^2.$$

Subtracting the lower bound on $q_H^\star$ from the resulting upper bound on $q_H(\lambda)$ proves Equation (17). The objective in Equation (18) is a positive semidefinite quadratic plus a constant. The simplex is compact, so a minimizer exists. The additive stopping statement follows directly. If $L_{\mathrm{sk}} > 0$, divide $q_H(\lambda) - q_H^\star \leq R_\tau(\lambda)$ by $q_H^\star \geq L_{\mathrm{sk}}$ to obtain Equation (19).

A.4 Recovery and monotone stopping

Corollary 1 (Recovery and monotone stopping)**.** *Let λ_{sk} be defined by Equation* (8)*. For the true Gram K, $q_K(\lambda_{\mathrm{sk}}) = U_{\mathrm{sk}}$. If $U_{\mathrm{sk}} = 0$, then it is exactly optimal. If $L_{\mathrm{sk}} > 0$ and $U_{\mathrm{sk}}/L_{\mathrm{sk}} \leq 1+\epsilon$, then $1 \leq U_{\mathrm{sk}}/q^\star \leq 1+\epsilon$. No finite multiplicative conclusion follows from the sketch when $L_{\mathrm{sk}} = 0 < U_{\mathrm{sk}}$.*

The additive regret remains bounded without dividing by L_{sk}:

$$0 \leq q_K(\lambda_{\mathrm{sk}}) - q_K^\star \leq U_{\mathrm{sk}} - L_{\mathrm{sk}} \leq g_{\mathrm{FW}}(\lambda_{\mathrm{sk}}), \tag{20}$$

where $g_{\mathrm{FW}}(\lambda) = 2(q_K(\lambda) - \min_i (K\lambda)_i)$.

L_{sk} also dominates the rule that matrix-free MGDA solvers stop on. Writing $m = \min_i (K\lambda_{\mathrm{sk}})_i$,

$$L_{\mathrm{sk}} \;\geq\; \max\{0,\; 2m - U_{\mathrm{sk}}\} \;=\; U_{\mathrm{sk}} - g_{\mathrm{FW}}(\lambda_{\mathrm{sk}}) \vee 0, \tag{21}$$

with no side condition, so a run stopping on the Frank–Wolfe bound stops no earlier than one stopping on L_{sk}. A Cauchy–Schwarz refinement adds m_+^2/U_{sk} when the margin is nonnegative, and knowing exactly when a certificate is licensed is the point of this paper: without that gate the refinement is false, and Appendix A gives a positive definite 2×2 counterexample where it would report an exact certificate at five times the optimum. We therefore state Equation (21) *in the form that needs no hypothesis, and Appendix G finds the refinement never binds on a real trajectory.*

For nested sketches generated from the same Gram, $L_s \leq L_t \leq U_t \leq U_s$, and if $\operatorname{rank}(\Omega_t) = M$ then $L_t = U_t = q_K^\star$.

Proof. The recovered vector is feasible with value U_{sk} under the true Gram, and Theorem 1 gives $L_{\mathrm{sk}} \leq q^\star$. If $U_{\mathrm{sk}} = 0$, positive semidefiniteness gives exact optimality, and under the second premise $1 \leq U_{\mathrm{sk}}/q^\star \leq U_{\mathrm{sk}}/L_{\mathrm{sk}} \leq 1 + \epsilon$. For the additive statement, $q_K(\lambda_{\mathrm{sk}}) = U_{\mathrm{sk}}$ and $q_K^\star \geq L_{\mathrm{sk}}$ give the first bound, while convexity of the K_{sk} quadratic gives $L_{\mathrm{sk}} \geq q_{K_{\mathrm{sk}}}(\lambda_{\mathrm{sk}}) + \min_{\mu \in \Delta_M} 2(K_{\mathrm{sk}}\lambda_{\mathrm{sk}})^\top(\mu - \lambda_{\mathrm{sk}}) = U_{\mathrm{sk}} - g_{\mathrm{FW}}(\lambda_{\mathrm{sk}})$. Because $K_{\mathrm{sk}}\lambda_{\mathrm{sk}} = K\lambda_{\mathrm{sk}}$, that quantity is available from the recovered response. For nested spans, appending actions shrinks the version space, so the lower endpoint can only increase and the upper only decrease, and at $\operatorname{rank}(\Omega) = M$ the sketch identifies K.

A.5 What a sketch determines, and the stopping bound

We prove the determined-functional characterization, the necessity bound stated in Section 4, and the impossibility of removing its conditioning on the transcript.

Throughout, $\mathcal{S} = \mathrm{range}(\Omega)$, $P = P_{\mathcal{S}^\perp}$, and the sketch is *incomplete* when $P \neq 0$. Two facts drive everything below. The version space is $\mathcal{V}(\Omega, Y) = K_{\mathrm{sk}} + \mathcal{B}$ with $\mathcal{B} = \{B \succeq 0 : B\Omega = 0\}$, because K_{sk} is its Loewner-least element; and any $\mathcal{K}$-monotone functional is therefore minimized at K_{sk} with no search, which covers trace, log-determinant, largest eigenvalue and every optimization value here. Both hold verbatim for any closed convex cone whose feasible set has a minimum element, which is why the statements below are about the semidefinite case without depending on it.

A.5.1 The semidefinite case, exactly

We prove Theorem 2 and enumerate the functionals it rules in and out.

Specializing the version-space decomposition to $\mathcal{K} = \mathbb{S}^M_+$ and the measurement $H \mapsto H\Omega$ gives Theorem 2, whose proof we give here.

Proof of Theorem 2. By Equation (6), $\mathcal{V}(\Omega, Y) = K_{\mathrm{sk}} + \mathcal{B}$ with $\mathcal{B} = \{B \succeq 0 : B\Omega = 0\}$. For symmetric B, $B\Omega = 0$ says $\mathrm{range}(B) \subseteq \mathcal{S}^\perp$, which for a symmetric matrix is the same as $B = \Pi B \Pi$, so $\mathcal{B}$ is exactly the positive semidefinite cone of the subspace $\mathrm{Sym}(\mathcal{S}^\perp)$. Then $\langle A, K_{\mathrm{sk}} + B\rangle - \langle A, K_{\mathrm{sk}}\rangle = \mathrm{tr}(A\Pi B\Pi) = \langle \Pi A \Pi, B\rangle$, so $\langle A, \cdot\rangle$ is constant on $\mathcal{V}$ if and only if $\langle \Pi A\Pi, B\rangle = 0$ for every $B \in \mathcal{B}$. That cone is full dimensional in $\mathrm{Sym}(\mathcal{S}^\perp)$ and $\Pi A \Pi$ lies in that space, so a symmetric matrix orthogonal to all of it is zero. Hence constancy holds if and only if $\Pi A \Pi = 0$, and the common value is $\langle A, K_{\mathrm{sk}}\rangle$. □

Corollary 2 (What this rules in and out)**.** *Let the sketch be incomplete. Then*

(a) the response $H\lambda$ is determined if and only if $\lambda \in \mathcal{S}$, and then $H\lambda = Yc$ for any c with $\lambda = \Omega c$;

(b) the quadratic form $\lambda^\top H \lambda$ is determined if and only if $\lambda \in \mathcal{S}$, and then equals $c^\top W c$ with $W = \Omega^\top Y$;

(c) the entry H_{ij} is determined if and only if Pe_i and Pe_j are not both nonzero;

(d) $\mathrm{tr}(H)$ is determined for no incomplete sketch.

Proof. (a) $(H\lambda)_i = \langle A_i, H\rangle$ with $A_i = \frac{1}{2}(e_i\lambda^\top + \lambda e_i^\top)$ and $PA_iP = \frac{1}{2}\big((Pe_i)(P\lambda)^\top + (P\lambda)(Pe_i)^\top\big)$. If $P\lambda = 0$ this vanishes for every i. Conversely, if $P\lambda \neq 0$ pick i with $Pe_i \neq 0$, which exists because $P \neq 0$, and then $PA_iP \neq 0$, since a sum $uv^\top + vu^\top$ with $u, v \neq 0$ is nonzero. Apply Theorem 2. The closed form is $K_{\mathrm{sk}}\Omega c = YW^\dagger W c = Yc$, using $\ker W \subseteq \ker Y$. (b) Take $A = \lambda\lambda^\top$, so $\Pi A \Pi = (\Pi\lambda)(\Pi\lambda)^\top$, which vanishes exactly when $P\lambda = 0$, and then $\lambda^\top K_{\mathrm{sk}}\lambda = (\Omega c)^\top Yc = c^\top W c$. (c) Take $A = \frac{1}{2}(e_ie_j^\top + e_je_i^\top)$. (d) Take $A = I$, so $\Pi A \Pi = \Pi \neq 0$. □

The forward direction is linearity. The converse is the limit. No estimator recovers a functional with a nonzero off-span block, because two compatible Grams disagree on it. That bounds the nulls of Appendix C.6 from one side without explaining their size, since every prior we test leaves $\mathcal{B}$ full dimensional. Part (d) earns its keep separately: the trace is determined by no incomplete sketch, so the scalar Proposition 4 needs must be bought. The impossibility is prior-free, which is both its strength and its scope, since a structural prior on the residual lies outside $\{B \succeq 0 : B\Omega = 0\}$.

A.5.2 The interval bounds the stopping time of every correct algorithm

We restate the proof of Theorem 3, stated in Section 4.

Proof. $\mathcal{W}$ is by hypothesis a set of Grams each producing the transcript (Ω, Y), and $\mathcal{A}$ reads only that transcript, so $\lambda_H = \mathcal{A}(\Omega, Y) =: \lambda$ does not depend on H. The per-instance guarantees, which separately concern different H, therefore hold simultaneously at that single λ, giving $\sup_{H \in \mathcal{W}}(q_H(\lambda) - q_H^\star) \leq \eta$, and the minimum over Δ_M is no larger. □

Indistinguishability is assumed rather than established, and Lemma 1 certifies it for the sets we use. We avoid the name "instance-optimal", since Proposition 2 shows genuine instance optimality is degenerate. The bound also covers randomized output: if $\mathcal{A}$ draws from a distribution μ depending

only on the transcript and is correct at every instance in expectation, convexity of $\lambda \mapsto q_H(\lambda)$ gives $q_H(\mathbb{E}_\mu[\lambda]) \le \mathbb{E}_\mu\, q_H(\lambda)$ by Jensen, and the mean weighting is a single deterministic decision to which the argument applies. We claim nothing for guarantees holding only with probability $1-\delta$. The proof never uses unboundedness, so it applies to any version space carrying a further prior a deployment can justify.

Corollary 3 (The two spaces this paper uses)**.** *With $\mathcal{W} = \mathcal{V}(\Omega, Y)$ and $\Delta_M \cap \mathcal{S} \neq \emptyset$, Theorem 1 evaluates the left side of Equation* (11) *in closed form, so $U_{\mathrm{sk}} - L_{\mathrm{sk}} \le \eta$ and the interval is itself the stopping time. With $\mathcal{W} = \mathcal{V}_{\mathrm{eq}}$, the trace-equality space the measurement buys, the left side is no larger.*

On the recorded Grams that reduction is under 0.6%. That figure is measured and is not part of the corollary.

The difference is computability, not validity. The hypothesis $\Delta_M \cap \mathcal{S} \neq \emptyset$ is needed only for the closed form: with $M = 2$, one query $\Omega = (1, -1)^\top$ and true $K = I$, we get $U_{\mathrm{sk}} = +\infty$ and $L_{\mathrm{sk}} = 0$ while $\lambda = (\frac{1}{2}, \frac{1}{2})$ has regret zero, and Theorem 3 is unaffected because its left side is zero there. Our recovery initializes with a simplex vector, so every transcript we report satisfies it. Consequently the stopping rule of Corollary 1 fires, on a given transcript, when correctness first becomes possible, and the refusals Section 6.2 reports are statements about the actions bought rather than about the rule reading them.

A.5.3 Why the bound must be conditional on the transcript

We show that no bound on certified stopping can be written as a function of the Gram alone.

The natural repair to the conditioning is a query-complexity bound depending on K alone. No such bound exists.

Proposition 2 (No nontrivial policy-independent lower bound)**.** *Let $q^\star > 0$ and let $\lambda^\star$ attain it. The single query $\Omega = [\lambda^\star]$ yields $L_{\mathrm{sk}} = U_{\mathrm{sk}} = q^\star$. Hence for every $k \ge 1$ the minimum over query policies of the interval width after k actions is zero, and no lower bound on certified stopping can be a function of K alone.*

Proof. $\lambda^\star \in \Delta_M \cap \mathcal{S}$, so $U_{\mathrm{sk}} \le \lambda^{\star\top} K \lambda^\star = q^\star$, and $U_{\mathrm{sk}} \ge q^\star$ always, giving $U_{\mathrm{sk}} = q^\star$. For the lower endpoint, $K_{\mathrm{sk}} = (K\lambda^\star)(K\lambda^\star)^\top / q^\star$, so for any $\lambda \in \Delta_M$, $\lambda^\top K_{\mathrm{sk}} \lambda = (\lambda^\top K \lambda^\star)^2 / q^\star \ge q^\star$, because MGDA optimality gives $\lambda^\top K \lambda^\star \ge q^\star$ on the simplex. With $L_{\mathrm{sk}} \le q^\star$ this forces $L_{\mathrm{sk}} = q^\star$. □

The same construction shows the rank-budgeted obstruction $q^\star - \max\{\min_{\Delta_M} \lambda^\top A \lambda : 0 \preceq A \preceq K,\ \mathrm{rank}\, A \le k\}$ is identically zero for every $k \ge 1$, so that candidate is vacuous too.

What this rules out needs care. The proposition concerns policies allowed to depend on the instance, the familiar degeneracy of instance-aware policies rather than an impossibility theorem, and it leaves the minimax count $\inf_\pi \sup_K T(\pi, K)$ used by the matrix-vector query literature (Sun et al., 2019; Braverman et al., 2020) open. What it establishes is why our result is conditional: no transcript-free statement about a fixed Gram is available, so our negatives concern the transcripts our own acquisition orders produced, and Appendix C.8 measures the order effect separately.

A.5.4 Adaptive queries do not break the conditioning

Lemma 1 (Conditioning under adaptive acquisition)**.** *Let the query rule choose $\omega_t = \pi_t(\omega_1, r_1, \ldots, \omega_{t-1}, r_{t-1}; \xi)$, where r_i is the observed response and ξ is randomness drawn independently of the Gram. Run it against the true Gram K and let (Ω, Y) be the realized transcript. Then every $H \in \mathcal{V}(\Omega, Y)$ produces exactly the same transcript when the same rule is run against H with the same ξ.*

Proof. Induction on t. At $t = 1$, $\omega_1 = \pi_1(\xi)$ does not depend on the matrix, so both runs issue it, and $H\omega_1 = Y_{:,1} = K\omega_1$ because $H\Omega = Y$. Suppose both runs agree on $(\omega_i, r_i)_{i<t}$. Then ω_t is the same function of the same history and the same ξ, so both issue it, and $H\omega_t = Y_{:,t} = K\omega_t$ again by $H\Omega = Y$. □

So the transcript carries no information distinguishing members of $\mathcal{V}(\Omega, Y)$, and the adversary choosing H may be thought of as moving after the transcript is fixed. Independence of ξ from the Gram is needed, and a rule whose randomness depended on K would not be covered. The deployed rule conditions on more than the transcript: on the contracted blocks C of Theorem 4 and on the measured trace τ. The induction extends to both without change, since every member of $\mathcal{V}_C$ shares the same C and residual responses Y_R, every member of $\mathcal{V}_{\text{eq}}$ shares the same $\text{tr}(K)$, and neither quantity is a function of the queries.

A.5.5 Reading the characterization as a construction

Corollary 4 (Exact margin and certified gain)**.** *Let $\lambda \in \Delta_M \cap \mathcal{S}$, so $\lambda = \Omega c$. Then the worst-task slope $m(\lambda) = \min_i (H\lambda)_i = \min_i (Yc)_i$ and the direction norm $\|G\lambda\|_2 = \sqrt{c^\top W c}$ are the same for every $H \in \mathcal{V}(\Omega, Y)$, so both are known exactly from the sketch. If in addition $m(\lambda) \geq 0$, then*

$$\inf_{H \in \mathcal{V}} \frac{\min_i (H\lambda)_i}{\|G\lambda\|_2 \sqrt{q_H^\star}} \;=\; \frac{m(\lambda)}{\sqrt{c^\top W c \, U_{\text{sk}}}}. \tag{22}$$

Proof. Exactness of numerator and norm is Corollary 2(a) and (b). For the infimum, the numerator and the norm are constant over $\mathcal{V}$ while $q_H^\star$ ranges over $[L_{\text{sk}}, U_{\text{sk}})$ by Theorem 1, the upper end not being attained when $L_{\text{sk}} < U_{\text{sk}}$. When $m(\lambda) \geq 0$ the ratio decreases as $q_H^\star$ grows, so the infimum is taken at $\sup_H q_H^\star = U_{\text{sk}}$ and is itself not attained in that case. □

The sign hypothesis is not cosmetic. If $m(\lambda) < 0$ the ratio is negative and increases in $q_H^\star$, so the infimum sits at L_{sk}, the certified quantity becomes a worst-case *ascent* rate, and the expensive lower solve returns. Every use here is gated on $m_{\text{sk}} > 0$, a sign test on Yc that costs nothing.

Two reductions follow. Only U_{sk} appears in Equation (22), so the lower completion solve is never required, and everything there lives in Y and the $k \times k$ core W, so a round costs $O(Mk)$ against the objective certificate's $O(M^2)$ per iterate. Both identities were checked to 10^{-13}. Any weighting can be audited for ascent with one action on $K\lambda$, redundant when $\lambda \in \mathcal{S}$, so the audit is free for spanned decisions and costs one action for uniform, for MGDA-UB's output, or for any rule returning a weighting off the span.

Since $\sqrt{U_{\text{sk}}}$ does not depend on λ, maximizing Equation (22) means maximizing $m(\lambda)/\sqrt{\lambda^\top K_{\text{sk}} \lambda}$. Numerator and denominator are both homogeneous of degree one, so after rescaling this is

$$\min_c \; c^\top W c \quad \text{s.t.} \quad (Yc)_i \geq 1 \;\; \forall i, \qquad \Omega c \geq 0, \tag{23}$$

a convex quadratic program under linear constraints in k variables, not M. Its minimizer rescaled to the simplex is the gain-optimal decision and its value v gives best certified gain $1/\sqrt{v\, U_{\text{sk}}}$. Infeasibility is informative rather than a solver failure: it says the observed span contains no common-descent direction, a certified reason to acquire another action rather than return a weighting.

A.5.6 Retargeting acquisition to the margin

We replay both query paths to measure what acquisition saves when it chases the determined margin instead of the bracketed objective.

If the margin is what is determined, it is what acquisition should chase. We replayed the 103 recorded Grams and the six real CIFAR-10 Grams of Appendix C.7, comparing the objective-ordered path against one that greedily improves the certified gain of Equation (22).

Pooled over the six real Grams the median falls from 67.5 actions to 45, a factor of 1.50, with six of six improving, and the median descent certificate falls from 10.7% and 7.5% of a full Gram to 5.5% and 4.5%. These are the conflicting geometries on which Appendix C.6.5's block schedule refuses, so the retarget halves the price on the regime that was most expensive. The margin round is not uniformly cheaper in time: at $k = 12$ it costs 4.3 to 19.1 ms across $M = 128$ to 1024 against the objective round's 1.1 to 21.8 ms, which combined with the action counts is a $1.35\times$ loss at $M = 512$ and a $1.9\times$ saving at $M = 1024$. We state that as an estimate, since the per-round costs were timed at fixed k and not along the paths.

Table 2: Actions to certify common descent, objective-ordered path against margin-targeted path. Recorded rows are medians over the 103 Grams at three gain thresholds γ; the real rows are the six CIFAR-10 Grams at exact closure, reported as medians per size.

Panel	objective path	margin path	improved
Recorded, $\gamma = 0.5$	7	6	71/103
Recorded, $\gamma = 0.8$	12	10	92/103
Recorded, $\gamma = 0.9$	14	12	99/103
Real CIFAR-10, $M = 512$	55	**28**	3/3
Real CIFAR-10, $M = 1024$	77	**46**	3/3

A.5.7 Proof of Equation (21)

The linearization candidate is standard: convexity of $\mu \mapsto \mu^\top K_{\mathrm{sk}}\mu$ at λ_{sk} gives $L_{\mathrm{sk}} \geq 2\min_i(K_{\mathrm{sk}}\lambda_{\mathrm{sk}})_i - U_{\mathrm{sk}} = 2m - U_{\mathrm{sk}}$, using $K_{\mathrm{sk}}\lambda_{\mathrm{sk}} = K\lambda_{\mathrm{sk}}$ on the span. Nonnegativity of L_{sk} handles the first candidate.

A sharper candidate is available when the margin is nonnegative, and the hypothesis is exactly what licenses it. Assume $U_{\mathrm{sk}} > 0$ and $m \geq 0$. Write $\langle x, y\rangle = x^\top K_{\mathrm{sk}} y$, a positive semidefinite bilinear form. For any $\mu \in \Delta_M$,

$$\langle \mu, \lambda_{\mathrm{sk}}\rangle^2 \leq \langle \mu, \mu\rangle \, \langle \lambda_{\mathrm{sk}}, \lambda_{\mathrm{sk}}\rangle = \mu^\top K_{\mathrm{sk}}\mu \cdot U_{\mathrm{sk}},$$

so $\mu^\top K_{\mathrm{sk}}\mu \geq (\mu^\top K\lambda_{\mathrm{sk}})^2/U_{\mathrm{sk}}$ whenever $U_{\mathrm{sk}} > 0$. Since μ lies in the simplex and every coordinate of $K\lambda_{\mathrm{sk}}$ is at least $m \geq 0$, we have $\mu^\top K\lambda_{\mathrm{sk}} \geq m$, hence $\mu^\top K_{\mathrm{sk}}\mu \geq m^2/U_{\mathrm{sk}}$. Minimizing the left side over Δ_M gives $L_{\mathrm{sk}} \geq m^2/U_{\mathrm{sk}}$.

The sign condition is necessary, not technical. Take $K = \left(\begin{smallmatrix} 1 & -2 \\ -2 & 5 \end{smallmatrix}\right)$, which is positive definite, and $\Omega = (\frac{1}{2}, \frac{1}{2})^\top$. The span-simplex intersection is the single point $\lambda_{\mathrm{sk}} = (\frac{1}{2}, \frac{1}{2})$, so $U_{\mathrm{sk}} = \frac{1}{2}$ and $m = -\frac{1}{2}$, while $K_{\mathrm{sk}} = \left(\begin{smallmatrix} 1/2 & -3/2 \\ -3/2 & 9/2 \end{smallmatrix}\right)$ has $L_{\mathrm{sk}} = 0$ and the true optimum is $q^\star = \frac{1}{10}$. An ungated $m^2/U_{\mathrm{sk}} = \frac{1}{2}$ would assert $L_{\mathrm{sk}} \geq \frac{1}{2}$, report the ratio $U/\ell = 1$, and license a stop at five times the optimum. Every use of the refinement in this paper therefore carries a margin gate, recorded per experiment in Appendix B.5, and Equation (21) states the ungated form because Appendix G measures the refinement to be inert on a real trajectory: the Frank–Wolfe and composite rules stop at identical action counts on all 100 transcripts.

A.5.8 Proof of Proposition 1

Proof. Smoothness gives $L_i(\theta + \eta d) \leq L_i(\theta) + \eta\nabla L_i(\theta)^\top d + \frac{\eta^2\beta_i}{2}\|d\|_2^2$. Here $\nabla L_i(\theta)^\top d = -g_i^\top G\lambda = -(K\lambda)_i \leq -m$ by the definition of the margin, and $\|d\|_2^2 = \lambda^\top K\lambda = q$, which is Equation (14). Its right side is at most $L_i(\theta)$ exactly when $\eta\beta_i q/2 \leq m$. A step size $\eta \leq 2m/(q\max_i \beta_i)$ satisfies this for every i at once, so no task loss rises at such a step, and chaining over consecutive certified steps gives $\max_i L_i(\theta_T) \leq \max_i L_i(\theta_0)$. □

The hypothesis is what the result costs. It needs a valid β_i, and we do not estimate smoothness constants for the neural runs, so there the proposition explains the observed behaviour instead of certifying a level. It needs every step certified at an admissible size, whereas the runs of Section 6 refresh on a five-update cadence, so the chain covers the certified steps rather than all of them. What it does not need is a completion, a lower endpoint, or any appeal to the unobserved block: the margin is exact from the actions already taken. With a strictly negative response the first-order term is positive for that task, so no small step lowers it, which is the mechanism behind an uncertified method's runaway seed. A zero margin decides nothing: with $L_1(x) = \frac{1}{2}(x-1)^2$ and $L_2(x) = 1 + \cos x$ at $x = 0$ and $\lambda = (1, 0)$, one has $K\lambda = (1, 0)^\top$ and $m = 0$, yet both losses fall for every small positive step.

A.5.9 Scope and numerical checks

Transferring the certificate to inexact actions. The theorems assume exact responses $Y = K\Omega$. The following sufficient transfer states what an implementation must add.

Proposition 3 (Action-error transfer). *Suppose $K \succeq 0$, $\widehat{Y} = K\Omega + E$ and $\|E\|_2 \le \delta$. For any $A \in \mathbb{R}^{M\times k}$ define*

$$Q_A = \widehat{Y}A^\top + A\widehat{Y}^\top - A\,\mathrm{sym}(\Omega^\top \widehat{Y})A^\top, \qquad b_A = \delta(2\|A\|_2 + \|A\|_2^2\|\Omega\|_2).$$

If $\ell_A \le \min_{\lambda\in\Delta_M} \lambda^\top Q_A \lambda$, then $q_K^\star \ge L_\delta := \max(0, \ell_A - b_A)$. For a feasible recovered vector $\lambda = \Omega c \in \Delta_M$, let $\widehat{u} = \widehat{Y}c$ and $\widehat{q} = \lambda^\top \widehat{u}$. Then

$$\begin{aligned} q_K(\lambda) &\le U_\delta := \widehat{q} + \delta\|\lambda\|_2\|c\|_2, \\ \min_i (K\lambda)_i &\ge \min_i \widehat{u}_i - \delta\|c\|_2, \\ \|G\lambda - G\lambda^\star\|_2^2 &\le q_K(\lambda) - q_K^\star \le U_\delta - L_\delta. \end{aligned}$$

Thus $U_\delta \le (1+\epsilon)L_\delta$ with $L_\delta > 0$ is a sufficient relative recovery test, and $\min_i \widehat{u}_i > \delta\|c\|_2$ is a sufficient strict-descent test.

Proof. With exact responses $Y = K\Omega$, direct expansion gives

$$K - (YA^\top + AY^\top - A\Omega^\top Y A^\top) = (I - A\Omega^\top)K(I - \Omega A^\top) \succeq 0.$$

Replacing Y by $\widehat{Y}$ perturbs the parenthesized symmetric matrix by operator norm at most b_A, using $\|\,\mathrm{sym}(\Omega^\top E)\|_2 \le \|\Omega\|_2\delta$. Therefore $K \succeq Q_A - b_A I$. Since every simplex vector has squared norm at most one, the stated lower bound follows. Also $\widehat{u} - K\lambda = Ec$, so its Euclidean and coordinatewise errors are at most $\delta\|c\|_2$. This proves the upper and margin bounds. Projection of zero onto the true gradient convex hull gives the direction-error inequality as in Theorem 4. □

For a stored FP64 Gram the radius is provable: floating-point products satisfy $\|\widehat{Y} - K\Omega\|_2 \le \gamma_M \||K||\Omega|\|_2$ with $\gamma_M = Mu/(1-Mu)$ and unit roundoff $u = 2^{-53}$. With this δ and $A = YW^\dagger$, all 103 retained-Gram stops of the canonical coordinate replay remain certified, and the transfer moves their brackets by a median relative 1.7×10^{-11} and at most 1.8×10^{-8}. Replay certificates are therefore rigorous for the stored Grams. Live FP32 runs supply no such radius, so their certificates are numerical and are checked against independent native-autograd Grams instead (Appendix G).

The cone minimum exists whenever a positive semidefinite object is observed through products, covering Gauss–Newton, Fisher and neural tangent matrices, kernel and Gaussian-process inference, and covariance reached through products. We prove the general statement and measure only the multitask instance, so this is scope and not a result.

Every claim above was checked numerically. Over 40 random problems the shorted operator is feasible and Loewner-minimal, and trace, largest eigenvalue, log-determinant and the MGDA optimum are each minimized there. Over 30 problems the bracket holds on the simplex and no sampled decision beat the interval on the eleven with a nondegenerate one. Theorem 2 and all four cases of Corollary 2 agree to 10^{-13} over 300 random symmetric A.

B Experimental setup

We state the harness, solver settings and capture protocols shared by the experiments. Per-study protocols appear with their results in Appendix C.

B.1 Solvers and numerical gates

All arithmetic is NumPy float64. Simplex quadratic programs use an active-set solver checked against the release reference to 3×10^{-12}, and the solver handles singular faces: one that skips them reports spurious failures on the concavity check of Appendix A. Gram replays divide K by $\mathrm{tr}(K)/M$, so the trace is fixed by construction on replayed instances. Eigenvalues above -2×10^{-13} are clipped and renormalized.

B.2 Recorded Gram collection

The 103 recorded Grams come from about five independent captures: 60 snapshots from three language-model trajectories, 42 from one 21-checkpoint paired capture, and one ResNet Gram at $M = 100$. Snapshots within a trajectory are dependent, so pooled counts summarize roughly five independent sources rather than 103, and Appendix C reports every headline count per capture as well as pooled.

B.3 Which audit quantities gate a run

The completed-Gram audit mixes correctness checks with scale-sensitive diagnostics, and only the former abort a run. Interval containment, absence of a false 5% stop, completed rank M, basis orthogonality, candidate feasibility and the scaled minimum eigenvalue are correctness properties and are enforced. Two diagnostics are recorded instead.

The scaled Gram asymmetry $\|YQ^\top - QY^\top\|/\|YQ^\top\|$ is float32 round-off. Over 100 certified training steps its median is 2.7×10^{-4} and its maximum 3.8×10^{-3}.

The stored-response error is the same asymmetry applied to λ and divided by $\|K\lambda\|$, so it is a ratio whose denominator vanishes as the min-norm point degenerates. Over the same run $\|K\lambda\|$ falls from 5.5×10^7 to 0.64 and the ratio rises to a median 1.05, while the numerator stays at round-off scale: the correlation between $\log\|K\lambda\|$ and the log ratio is -0.885, and the global asymmetry is *smaller* on the steps where the ratio is largest (1.4×10^{-4} against 4.1×10^{-4}). Gating on it would abort runs for approaching Pareto stationarity, which is the behaviour the method is meant to detect, so it is reported rather than enforced. A small error relative to the whole Gram can still flip the sign of a near-vanishing margin. The transfer bound of Appendix A.5.9 states the radius a rigorous margin test needs, and Appendix G identifies the steps where that sign is unresolvable in float32. The direction is unaffected: $q/q^\star$ stays below 1.0003 and no step records a false stop.

B.4 Endpoint audit paths

The paired endpoint audit normalizes each Gram by its mean diagonal, charges its initial query, and acquires one further column at a time. Two initializations are run over the identical Grams: a uniform-first path, and a pure-coordinate path that starts at the fixed first coordinate without consulting the diagonal. Exact span-simplex recovery is implemented by the acquired coordinates together with the uniform vector on their complement, which spans precisely the feasible simplex section. Each prefix reuses one upper solution across all lower tests, so the endpoints differ only in the lower bound and not in the queries or the iterates.

At a 5% objective target the median actions for the Frank–Wolfe, one-action and completion endpoints are $15/15/12$, $20/20/16$ and $17/17/15$ on the historical-EOS, answer-only and paired-EOS panels under the uniform-first path, and $15/15/13$, $21/21/18$ and $17/17/16$ under the pure-coordinate path. The single ResNet Gram gives $99/99/89$ and $100/100/100$; on the coordinate path neither endpoint stops inside the budget. Figure 1(b) quotes the answer-only pair. The completion endpoint is no later than the Frank–Wolfe endpoint on every panel under both initializations, which is Corollary 1; the initialization moves the level and not the ordering.

B.5 Which endpoint drives each experiment

Theorem 1's endpoint L_{sk} and the cheap stopping test used in training are different objects, and Table 3 records which one decides each block of Table 1. The cheap test is $U \le 1.05\,\ell_{\mathrm{c}}$ with $\ell_{\mathrm{c}} = \max\{0, 2m - U, m_+^2/U\}$, the Frank–Wolfe linearization together with the one-action Cauchy–Schwarz endpoint. Both candidates are dominated by L_{sk} (Corollary 1), so a run stopping on them stops no earlier than one stopping on L_{sk} would, and its action count is an upper bound on what the sharp endpoint needs.

The Cauchy–Schwarz candidate is valid only for $m \ge 0$, so every use of ℓ_{c} must carry a margin gate. Table 3 records whether one is documented. Block 2 states $m \ge 0.95U$ in its protocol. Block 4 applies the joint gain condition $\min_i(K\lambda)_i/q(\lambda) \ge 0.95$, and that gate is checkable in the released record: on the real $M = 512$ Grams the rule is marked certified exactly on the runs whose gain

reaches 0.95, and it is refused at gains 0.750, 0.866 and 0.865. For blocks 1 and 5 the released records do not identify the lower endpoint or the gate, so we do not claim one was applied. That is a valid bound and not a demonstration. Row 8 is the exception, and it is the run of Appendix G.

Table 3: Stopping rule behind each block of Table 1, and the margin gate ℓ_c requires. Row 8 is not a block of that table but the separate end-to-end run of Appendix G. Blocks that certify nothing acquire no actions and so have no endpoint. Where the released record does not identify the rule we say so rather than infer it.

Block	Setting	Produces	Endpoint that decides
1	CelebA-40	trained model	5% target, endpoint not recorded
2	Pythia-160M, 500 updates	trained model	ℓ_c, gate $m \geq 0.95U$
3	Pythia-160M, $M{=}32$	one direction	none, Direct contracts all blocks
4	Synthetic, $M{=}2048$	one direction	ℓ_c, gate $m/q \geq 0.95$
5	Pythia-160M, 100 updates	trained model	not recorded in the artifact
6	Pythia-160M, $M{=}16$, 60 updates	trained models	one protocol, six arms
7	103 recorded Grams	replay only	L_{sk}, the sharp endpoint
8	Pythia-160M, 100 updates	trained model	L_{sk}, end to end

Row 8 is the run of Appendix G, where the sharp endpoint drives a trained model and, on that run's own transcripts, certifies earlier than the cheap rule at identical queries. A query policy that is at once pass-economical and endpoint-aware remains open, and Appendix C.5 prices it.

B.6 Hardware and timing

Wall-clock figures are medians of three warm repetitions on one A100-SXM4-40GB. Memory is reported two ways and each table says which: stored allocation is incremental above the resident state from the previous repetition, whereas Table 9 reports total allocated GPU memory. The two are not comparable across tables. Action counts and wall clock are reported separately throughout, since dense queries do not batch into passes the way a coordinate block does.

C Additional results

C.1 Clean answer-only training and implementation controls

We train Pythia-160M under four weighting rules for 500 updates and report fidelity, guarantee and test cross-entropy for each.

This appendix gives the full protocol for Section 5. We train the pinned Pythia-160M model for 500 scalar-SGD updates on eight related instruction tasks. Each update uses four examples per task. Prompt and EOS positions are excluded from the loss; the softmax retains the full vocabulary. Tokenized prompts are deduplicated across the training, validation, and test partitions. These are internal partitions of the instruction dataset, not the official test sets of its source benchmarks. Cross-entropy (CE) is averaged over supervised tokens within each example, then over examples within each task, then equally over the eight tasks. Thus neither answer length nor task partition size supplies extra weight; this teacher-forced CE is not classification accuracy.

Each method starts from identical pretrained parameters and uses three paired data-order seeds. A separate 30-update calibration chooses its learning rate from $\{10^{-4}, 3\times10^{-4}, 10^{-3}\}$ by validation cross-entropy over 30 of the 500 updates, a short horizon on a three-point grid. ParetoTrace-Stream uses the paired ParetoTrace-Direct learning rate, which also wins its own calibration (validation CE 0.723 at 3×10^{-4} against 1.276 and 2.771). Its cached action kernel is selected on a separate fixed batch and validated before training; it is not the two-pass memory-bounded path. No gradient normalization, clipping, or momentum is used. We retain all three runs and report observed seed ranges, not confidence intervals or independent replications at the update level.

The update timer includes direction recovery and the explicit scalar-SGD parameter update, with CUDA synchronization at its boundaries. Batch selection, packing and host-to-device transfer occur beforehand; there is no data-loader loop inside this interval. Recorded wall additionally includes these operations, evaluation, separately computed numerical audits and periodic checkpoint writes. It

Table 4: What the accuracy is worth. $\cos \geq$ lower-bounds the *cosine* of the angle to exact MGDA via $\sqrt{q^\star/q(\lambda)}$; the angle column inverts it. Gain is the minimum normalized worst-task slope, negative when the update raises some task's loss. Both are worst cases over six audited updates in each of three runs, 18 per arm, against an FP64 reference, rounded outward so each entry stays a valid bound. Δ is the paired held-out difference in test CE against ParetoTrace-Direct, and within-arm seed ranges span 0.010 to 0.091.

Arm	$\cos \geq$	angle $\leq$	min gain	Test CE	Δ
ParetoTrace-Direct	$1 - 6\times10^{-14}$	0.0004 mrad	0.9999	0.418	—
ParetoTrace-Stream	0.999975	7.1 mrad	0.983	0.442	$+0.024$
CAGrad	0.533	1008 mrad	0.108	**0.343**	-0.076
Uniform	0.150	1421 mrad	-0.420	0.350	-0.069

Table 5: Clean training, three paired seeds and 500 updates per run. Test CE is seed mean [minimum, maximum]; both time columns are seed means. Total update time sums all 500 measured recovery-and-SGD intervals, excluding batch selection/packing/transfer, numerical audits, evaluation and checkpoint writes. Recorded wall is run-invocation time; see the resume qualification below. Peak is the maximum allocated GPU memory over all updates and seeds.

Method	Learning rate	Test CE	Total 500-update time (s)	Recorded wall (s)	Peak (GiB)
ParetoTrace-Direct	3e-04	0.418 [0.369, 0.445]	139.1	155.9	4.04
ParetoTrace-Stream (cached actions)	3e-04	0.442 [0.413, 0.484]	407.7	427.7	4.04
Uniform	3e-04	0.350 [0.293, 0.384]	75.9	92.2	3.25
CAGrad	1e-04	0.343 [0.338, 0.349]	140.0	156.9	4.04

excludes model construction/loading, pool preparation, learning-rate calibration, and the final metrics write/cleanup. On checkpoint resume, update totals remain cumulative but the wall timer restarts. Final metric records do not identify restart counts, so recorded wall is not asserted to be cumulative experiment time. Table 5 separately reports verified cumulative update totals.

Conditions behind the fidelity–outcome table. The cosine and gain columns are worst cases over each run's six audited steps (1, 100, 200, 300, 400, 500) and three seeds, evaluated at that arm's own parameters against an FP64 full-Gram reference. The cosine bound requires a simplex argument, so the recorded ratio simplex-normalizes each arm's weights before the bound is applied; this matters only for CAGrad, whose audited weights sum to between 1.2158 and 2.0624 rather than to one; Uniform's sum to one exactly. The record stores only that sum and not the weight vector, so we cannot confirm from it that CAGrad's normalized weights are nonnegative, which the bound also requires. Read its row as an uncertified guarantee rather than as a measurement of angular deviation, since a weak lower bound on the cosine is consistent with a small true angle.

Stopping rule and fidelity of the ParetoTrace-Stream arm. This arm's acceptance test is $U \leq 1.05\ell_{\mathrm{c}}$ with $\ell_{\mathrm{c}} = \max\{0,\, 2m - U,\, m^2/U\}$ and $m \geq 0.95U$, where $m = \min_i(K\lambda)_i$. Those two candidate values for ℓ_{c} are the FW linearization and the one-action Cauchy–Schwarz endpoint. They are *not* the completion lower endpoint L_{sk} of Theorem 1, which the paired endpoint audit shows stops earlier than both. This arm therefore demonstrates that a streamed action interface can drive clean supervised training at high fidelity; it does not evaluate the sharp endpoint, and no conclusion about the endpoint's action count should be drawn from it.

It uses a median of eight actions (range four to eight) with no full-Gram fallback in 1,500 updates, reaching full rank on 1,135 of them, so 75.7% of its accepted sketches carry complete information. Both the small M and the weaker certificate contribute to that, and this run does not separate them. All 18 separately computed native/full-Gram audit checkpoints pass: the largest objective factor is 1.00005 and the smallest normalized gain is 0.9832. Those checkpoints are steps 1, 100, 200, 300, 400 and 500 of each run, and a failure raises and terminates the run, so they are an acceptance gate that these completed runs passed rather than a survey of all 500 updates. ParetoTrace-Direct is faster at the same measured peak allocation in this small-objective-count setting.

Cached ParetoTrace-Stream minus ParetoTrace-Direct test macro-CE differences are $+0.04396$, -0.01684, $+0.04378$ for seeds 2026091511, 2026091512 and 2026091513, respectively; these three paired observations do not establish trajectory equivalence or statistical significance. Uniform and

Table 6: Internal partition sizes and per-task test CE averaged over the three paired seeds. Each task CE first averages token-mean losses over its examples. Full seed ranges and paired per-task differences are retained in the numerical summary; these partitions are not official benchmark tests.

Task	Train	Val.	Test	ParetoTrace-Direct	Cached ParetoTrace-Stream	Uniform	CAGrad
glue/mnli	418	45	49	0.519	0.506	0.530	0.510
glue/qnli	414	52	46	0.201	0.262	0.218	0.203
glue/rte	424	46	42	0.224	0.211	0.188	0.186
super_glue/cb	205	26	19	0.660	0.730	0.262	0.262
anli/a1	412	53	47	0.636	0.651	0.547	0.536
sick/entailment_AB	419	47	46	0.169	0.200	0.176	0.163
snli	408	46	58	0.658	0.721	0.564	0.625
scitail/snli_format	417	47	48	0.279	0.253	0.314	0.255

Table 7: The acquisition rule run to its own certificate, on one A100-SXM4-40GB, medians of three warm repetitions over three seeds, timed from capture to parameter direction. Certified counts seeds meeting the 5% target with the gain condition. Materialization exhausts the card at 1024 and 2048. The fixed 18-action cap, reported per seed in the per-seed cost study, certifies 5/9 against the rule's 9/9.

M	Route	Seconds	Peak MB	Actions	Certified
512	Materialized	1.1898	24,692	512	3/3
	Contraction	0.7867	1,813	512	3/3
	Rule	**0.8036**	**2,046**	50	**3/3**
1024	Contraction	2.2582	3,497	1024	3/3
	Rule	**1.7031**	**3,617**	50	**3/3**
2048	Contraction	8.3639	6,993	2048	3/3
	Rule	**4.2516**	**6,998**	50	**3/3**

CAGrad have lower macro-CE in each paired seed, but improve on only four and seven of the eight seed-averaged task CEs, respectively. Defining a run's worst-task CE as the maximum of its eight task means, the means of these per-run maxima are 0.685, 0.788, 0.590, 0.625 for ParetoTrace-Direct, cached ParetoTrace-Stream, Uniform and CAGrad. Uniform has lower worst-task CE than ParetoTrace-Direct in all three paired seeds; cached ParetoTrace-Stream has higher worst-task CE in all three. These descriptive held-out outcomes do not establish a mean-loss or worst-task advantage for exact MGDA or an out-of-sample task-protection guarantee.

ParetoTrace-Direct and CAGrad share the factor-contraction implementation and have similar measured update costs. Their selected learning rates differ, so this is a comparison of separately calibrated methods, not an isolated weighting comparison at fixed step size. Native weighted-gradient checks and full-Gram numerical references audit implemented directions, not outward roundoff bounds or population outcomes.

C.2 Per-seed costs at 512 to 2048 objectives

We report the end-to-end cost of ParetoTrace-Stream against materialization and ParetoTrace-Direct, per seed, at 512 to 2048 objectives.

Section 5 reports medians. Because the acquisition arms do not all certify on every seed, the per-seed values matter and are given here in full, from the single A100-SXM4-40GB session recorded in the released records. Stopped marks whether the run met its own 5% relative target with the gain condition. Seed 0 has the lowest mean pairwise cosine at every size, at 0.625 to 0.637 against 0.792 to 0.808 for the other two, and is the seed that costs the rule four passes instead of two.

Three readings. The rule stopped on its own certificate on all nine runs, including every seed where the cap did not. The cap is a few percent faster where it works, and on the four runs where it does not it returns objective factors 1.010, 1.287, 1.410 and 1.522 at worst-task gain down to 0.384, which is a materially wrong direction that nothing in the run reports. And the cost of the rule's extra safety is small and shrinking: it is 8% slower than the cap at $M = 512$, 2% at 1024 and 1% at 2048, because the additional actions are absorbed into passes the cap was already paying for.

The seed-0 runs are the honest caveat. They need 244 actions over four passes rather than 50 over two, which at $M = 512$ makes the rule 1.36 s against materialization's 1.19 s, so on that seed alone

Table 8: All three seeds at each size, one A100-SXM4-40GB, medians of three warm repetitions. The rule stops on its own certificate; the cap stops at 18 actions regardless.

M	Route	Seed	Seconds	Peak MB	Actions	Passes	Stopped
512	Materialized	0	1.1867	24692.5	512	1	yes
	Materialized	1	1.1900	24691.9	512	1	yes
	Materialized	2	1.1898	24691.9	512	1	yes
	Contraction	0	0.7864	1813.2	512	1	yes
	Contraction	1	0.7871	1812.6	512	1	yes
	Contraction	2	0.7867	1812.6	512	1	yes
	Fixed 18-action cap	0	0.7402	2034.0	18	2	**no**
	Fixed 18-action cap	1	0.7713	2032.9	18	2	yes
	Fixed 18-action cap	2	0.7413	2032.9	18	2	yes
	Rule	0	1.3557	2097.1	244	4	yes
	Rule	1	0.8036	2045.5	50	2	yes
	Rule	2	0.7952	2045.5	50	2	yes
1024	Materialized	0	out of memory				
	Materialized	1	out of memory				
	Materialized	2	out of memory				
	Contraction	0	2.2642	3496.7	1024	1	yes
	Contraction	1	2.2582	3496.7	1024	1	yes
	Contraction	2	2.2569	3496.7	1024	1	yes
	Fixed 18-action cap	0	1.6508	3610.5	18	2	**no**
	Fixed 18-action cap	1	1.6664	3610.5	18	2	yes
	Fixed 18-action cap	2	1.6714	3610.5	18	2	yes
	Rule	0	2.9827	3642.2	244	4	yes
	Rule	1	1.6743	3616.6	50	2	yes
	Rule	2	1.7031	3616.6	50	2	yes
2048	Materialized	0	out of memory				
	Materialized	1	out of memory				
	Materialized	2	out of memory				
	Contraction	0	8.3596	6993.4	2048	1	yes
	Contraction	1	8.3639	6993.4	2048	1	yes
	Contraction	2	8.3763	6993.4	2048	1	yes
	Fixed 18-action cap	0	4.2012	6994.6	18	2	**no**
	Fixed 18-action cap	1	4.1806	6994.6	18	2	**no**
	Fixed 18-action cap	2	4.2269	6994.6	18	2	yes
	Rule	0	7.6882	7010.3	244	4	yes
	Rule	1	4.2285	6997.7	50	2	yes
	Rule	2	4.2516	6997.7	50	2	yes

certificate-driven acquisition loses to materialization on time. It still wins on memory by $11.8\times$. At 1024 and 2048 materialization is not available at all, so the comparison there is against contraction, which seed 0 beats at 2048 (7.69 s against 8.36 s) and loses to at 1024 (2.98 s against 2.26 s). The median figures quoted in Section 5 are not carried by this seed, and it is the seed the agreement predictor would have flagged as expensive.

C.3 Independent batches and stronger exact-MGDA controls

On the same pinned Pythia-160M/A100 setup, we compare complete direction construction across three separately sampled batches, with one warmup and two timed repetitions per batch. Times and peaks are first aggregated within batch; Table 9 reports the median across the three batches. The $M = 32$ case partitions examples into four strata within each of the eight source tasks; it is not a collection of 32 independent datasets.

All methods use the same clean answer-only losses, supervised-position head, FP32 model, and disabled TF32. Costs include derivative acquisition, Gram construction, QP solution and final direction. TorchJD is version 0.17.0. The example-Gram route aggregates its per-example Gram into task means; the grouped-output route differentiates the task means directly. Chunked vmap writes into a preallocated task table. Parameterwise vmap avoids a second flattened copy, accumulating

Table 9: Complete MGDA direction construction on three batch seeds. Peak is total allocated GPU memory. OOM means the configuration exceeded the 40-GB A100's available memory; it is not assigned a runtime.

M	Implementation	Batches	Seconds	Peak GiB
8	ParetoTrace-Direct	3	0.283	4.04
8	ParetoTrace-Direct, microbatch 32	3	0.284	4.04
8	Serial task gradients	3	0.895	8.71
8	Chunked vmap, 2 tasks	3	0.925	9.45
8	Chunked vmap, 8 tasks	3	0.891	14.02
8	Parameterwise vmap	3	0.888	6.71
8	TorchJD, example Gram	3	0.421	13.37
8	TorchJD, grouped output	3	1.043	9.65
32	ParetoTrace-Direct	3	0.857	14.21
32	ParetoTrace-Direct, microbatch 32	3	0.980	10.18
32	Serial task gradients	3	12.803	30.87
32	Chunked vmap, 2 tasks	3	13.527	31.96
32	Chunked vmap, 8 tasks	0	OOM	
32	Parameterwise vmap	0	OOM	
32	TorchJD, example Gram	0	OOM	
32	TorchJD, grouped output	0	OOM	

parameter-block Grams before reconstruction. Failures are retained. Full-Gram FP64 references and native backward checks are performed outside the timing region; every successful reported row meets relative objective error 10^{-3}, minimum response ratio 0.999, and native-direction relative error 10^{-3} in those numerical audits.

At $M = 8$, the fastest tested external control is TorchJD's example-Gram route: 0.421 seconds and 13.37 GiB, versus ParetoTrace-Direct's 0.283 seconds and 4.04 GiB. At $M = 32$, microbatching ParetoTrace-Direct lowers peak allocation while increasing runtime. These comparisons attribute the execution gain to the factor-contraction path; they do not establish an adaptive-query speedup.

C.4 The arms behind Figure 1(c)

Figure 1(c) is a separate study from Table 9: one held-out batch at $M = 8$ and batch size 32, five weighting rules rather than seven exact implementations, and its own frozen configuration selection. Its arms are not rows of Table 9 and the two should not be read against each other. Table 10 tabulates them so the panel is checkable.

Table 10: Complete-direction cost at $M = 8$ on one held-out batch, medians of three warm repetitions on one A100-SXM4-40GB. Total peak is `peak_allocated_bytes` and incremental peak subtracts the resident state, the two conventions of Appendix B. Frank–Wolfe here is the iterative simplex solver on the same contracted factors, so it shares ParetoTrace-Direct's memory and pays in iterations.

Method	Seconds	Total peak (GiB)	Incremental peak (GiB)
ParetoTrace-Direct	**0.287**	4.03	3.40
CAGrad	0.290	4.03	3.40
Materialized MGDA	0.892	8.70	8.08
Frank–Wolfe MGDA	3.843	4.03	3.40
Uniform	0.152	3.25	2.63

Uniform is the floor, since it builds no geometry at all. Against the exact arms, contraction is $3.11\times$ faster than materializing the Gram at $2.16\times$ less total peak, and $13.4\times$ faster than the Frank–Wolfe solver at equal memory. At $M = 8$ every route fits the card, so this panel prices arithmetic rather than feasibility.

C.5 What an action costs: the measured pass constant

Acquisition cost on this hardware is governed by two constants, both directly measurable, and they predict the execution ordering of Section 5 from first principles. Write f for the cost of one pass over

the factor cache, which forms the per-objective tile, and g for the marginal cost of one coordinate column once that tile is resident. Write P_{a} for the number of passes, distinct from the residual Gram R of Theorem 4. Forming the whole Gram itself pays one pass and M columns, so a P_{a}-pass recovery of k actions beats it exactly when $P_{\mathrm{a}} f + kg < f + Mg$, that is when

$$(P_{\mathrm{a}} - 1)\,(f/g) + k < M. \tag{24}$$

Both constants are directly measurable. On an A100 with an untrained ResNet-18 (11.23M parameters, BatchNorm affine frozen, FP32 tiles bounded by a 256 MB budget), timing a pass that serves no columns against one that serves 64 gives

M	Pass f (s)	Per column g (s)	f/g (columns)
128	0.0341	2.63×10^{-4}	129.7
256	0.0669	5.30×10^{-4}	126.1
512	0.1313	8.23×10^{-4}	159.6
1024	0.2666	1.60×10^{-3}	166.3

One pass therefore costs as much as 126–166 coordinate columns across the range, because forming the tile is bandwidth-bound while the column GEMM runs on tensor cores. This locates the cost in P_{a} rather than k, and it accounts for the ordering in Section 5: the coordinate policy of the deployed trajectory policy acquires about one coordinate per round, so P_{a} equals its action count and the left side of Equation (24) exceeds M at every size we run. The action saving that policy does deliver is real, a median 52 of 100, but it is a saving in information rather than in time.

The inequality is satisfiable, and a schedule keeping P_{a} small realises the saving it predicts. Acquiring a block of coordinates per pass and stopping on the certificate of Corollary 1, three seeds at $M = 1024$ on strongly agreeing objectives give the figures below. These come from an earlier schedule with a smaller first block than the block schedule's floor-16 rule, so at the same cosines the action counts and times differ from the end-to-end cost table's, 26 and 124 actions here against 50 and 244 there. Both stop on the same certificate and differ only in how many coordinates a pass requests. We report both rather than keeping the more favourable one.

Mean cosine	Passes	Actions	ParetoTrace-Stream (s)	Contraction (s)	Speedup	Certified
0.626	4	124	2.262	2.298	$1.02\times$	yes
0.793	2	26	1.453	2.296	$1.58\times$	yes
0.808	2	26	1.455	2.296	$1.58\times$	yes

All three satisfy both branches, at relative factors 1.002–1.003 and normalized gains 0.954–0.988 against an FP64 full-Gram reference. The saving tracks agreement rather than M: the seed at cosine 0.626 needs four passes and returns to parity, while the two above cosine 0.79 need two. The floor-16 rule shows the same ordering at a larger constant. So the condition for acquisition to pay is not size alone but size together with objectives whose gradients agree, which is what Equation (24) makes checkable in advance.

Equation (24) is device- and implementation-specific, and it is a statement about cost only. What decides whether a sketch is sufficient is the interval, not the pass count.

C.6 What no prior recovers, and what a weaker target buys

We test four explanations for the interval's width that do not require the information to be missing, and report what a weaker or recalibrated target buys instead.

The range is wide because the observations are incomplete. This section rules out the alternatives. Four explanations for the width do not require missing information: the model may be needlessly adversarial, a probabilistic prior may beat a worst case, the target may be too demanding, or worst-case certification may simply be conservative. We test all four. None survives. The last two subsections then report what a weaker or recalibrated target does buy.

C.6.1 TWO CONSTRAINED VERSION SPACES

We state the two constrained version spaces that are both cheap and sharp, and measure what each one buys on the recorded Grams.

Proposition 4 (Spectral budget). *Fix $\tau > 0$ and let $\mathcal{V}_\tau$ be the bounded version space of Equation* (16)*, written here with the equivalent spectral cap $B \preceq \tau I$: since $B \succeq 0$ and $B\Omega = 0$ force $\mathrm{range}(B) \subseteq \mathcal{S}^\perp$, the constraints $B \preceq \tau I$ and $B \preceq \tau P_{\mathcal{S}^\perp}$ cut out the same set. Writing $P_{\mathcal{S}^\perp}$ for the orthogonal projector onto $\mathcal{S}^\perp$, for every λ,*

$$\sup_{H \in \mathcal{V}_\tau} \lambda^\top H \lambda = \lambda^\top (K_{\mathrm{sk}} + \tau P_{\mathcal{S}^\perp})\lambda, \tag{25}$$

and, writing $U_\tau = \min_{\lambda \in \Delta_M} \lambda^\top (K_{\mathrm{sk}} + \tau P_{\mathcal{S}^\perp})\lambda$,

$$\sup_{H \in \mathcal{V}_\tau} q_H^\star = U_\tau. \tag{26}$$

So U_τ is the sharp *finite upper endpoint, its minimizer is the minimax-objective decision over $\mathcal{V}_\tau$, and as $\tau \to \infty$, $U_\tau \uparrow U_{\mathrm{sk}}$, so Theorem 1 is the limiting member of this family.*

Proof. $B\Omega = 0$ gives $\mathrm{range}(B) \subseteq \mathcal{S}^\perp$, hence $B = P_{\mathcal{S}^\perp} B P_{\mathcal{S}^\perp}$ and $\lambda^\top B \lambda \le \tau \|P_{\mathcal{S}^\perp}\lambda\|^2$ by $B \preceq \tau I$. The value is attained at $B = \tau P_{\mathcal{S}^\perp}$, which is feasible, giving Equation (25).

For Equation (26), write $f(B, \lambda) = \lambda^\top (K_{\mathrm{sk}} + B)\lambda$ on $\mathcal{B}_\tau \times \Delta_M$, where $\mathcal{B}_\tau = \{B \succeq 0 : B\Omega = 0,\ B \preceq \tau I\}$ is convex, closed and bounded, hence compact, and contains 0. For fixed λ, f is linear and therefore concave and continuous in B. For fixed B, f is a positive-semidefinite quadratic and therefore convex and continuous in λ. Both sets are convex and compact, so Sion's minimax theorem gives

$$\sup_{B \in \mathcal{B}_\tau} \min_{\lambda \in \Delta_M} f(B, \lambda) = \min_{\lambda \in \Delta_M} \sup_{B \in \mathcal{B}_\tau} f(B, \lambda) = U_\tau,$$

and the left side is exactly $\sup_{H \in \mathcal{V}_\tau} q_H^\star$. The limit follows because $\lambda^\top (K_{\mathrm{sk}} + \tau P_{\mathcal{S}^\perp})\lambda \to \infty$ for every $\lambda \notin \mathcal{S}$ and is constant in τ for $\lambda \in \mathcal{S}$. □

Compactness is what the unbounded space lacks, which is why its supremum is generally unattained and Appendix A.2 argues by hand. It also has a consequence the unbounded model hides: under $\mathcal{V}_\tau$ the minimax decision minimizes $\lambda^\top (K_{\mathrm{sk}} + \tau P_{\mathcal{S}^\perp})\lambda$ over the whole simplex, so it need not lie in $\mathcal{S}$, whereas under Theorem 1 any off-span decision has infinite worst-case regret and the span constraint is forced. We measured whether that matters. At the stopping prefix of every one of the 103 Grams, the $\mathcal{V}_\tau$ minimax decision puts a median ℓ_1 mass of 0.0003 off the span at the measurable budget $\tau = \mathrm{tr}(K) - \mathrm{tr}(K_{\mathrm{sk}})$, at most 0.0017, and U_τ / U_{sk} has median 1.0000 and minimum 0.9999. The span restriction is therefore a convenience on this data rather than an artifact that changes the answer.

The same interval under a trace budget. A measured trace defines the family $\mathcal{V}_\tau^{\mathrm{tr}} = \{K_{\mathrm{sk}} + B : B \succeq 0,\ B\Omega = 0,\ \mathrm{tr}\, B \le \tau\}$ directly, and it has the same sharp endpoints. For every λ, with $x = P_{\mathcal{S}^\perp}\lambda$, $\lambda^\top B \lambda \le \lambda_{\max}(B)\|x\|^2 \le \tau \|x\|^2$, with equality at $B = \tau x x^\top / \|x\|^2$, which is admissible because $x \perp \mathcal{S}$. The family is compact and convex, so Sion's theorem again gives $\sup q_H^\star = U_\tau$, and $B = 0$ attains L_{sk}. Since $H \mapsto q_H^\star$ is continuous and the family is convex, the compatible optima fill exactly $[L_{\mathrm{sk}}, U_\tau]$. The measurement above is taken at the stopping prefix, where the upper endpoint barely moves. Over every prefix of positive width on the uniform-first coordinate paths, the budget narrows the interval by a median 1.7%, at most 28.3%, and moves the 5% stop earlier on 2 of the 103 Grams while leaving the median stop at 13 actions.

Proposition 5 (Diagonal budget). *Suppose $d = \mathrm{diag}(K)$ is observed alongside $Y = K\Omega$, which costs M scalars and is the same information the agreement predictor already assumes. Put $\delta = d - \mathrm{diag}(K_{\mathrm{sk}})$ and $s = \delta^{1/2}$ componentwise, both well defined because $K_{\mathrm{sk}} \preceq K$ gives $\delta \ge 0$. Then for every $\lambda \in \Delta_M$ and every H in $\mathcal{V}_{\mathrm{diag}} = \{K_{\mathrm{sk}} + B : B \succeq 0,\ B\Omega = 0,\ \mathrm{diag}(B) = \delta\}$, a low-rank-plus-diagonal family of the kind studied directly by Yeon & Anitescu (2025),*

$$\lambda^\top H \lambda \le \lambda^\top (K_{\mathrm{sk}} + s s^\top)\lambda,$$

so $U_{\mathrm{diag}} = \min_{\lambda \in \Delta_M} \lambda^\top (K_{\mathrm{sk}} + s s^\top)\lambda$ is a finite valid upper endpoint for every sketch.

Proof. $B \succeq 0$ gives $|B_{ij}| \le \sqrt{B_{ii}B_{jj}} = s_i s_j$, so for $\lambda \ge 0$, $\lambda^\top B\lambda \le \sum_{ij} \lambda_i\lambda_j s_i s_j = (s^\top\lambda)^2$. Let $\hat\lambda$ attain the minimum on the right. Every $H \in \mathcal{V}_{\mathrm{diag}}$ has $q^\star_H \le \hat\lambda^\top H\hat\lambda \le U_{\mathrm{diag}}$. □

The diagonal bound is attained by $B = ss^\top$, so it is sharp for the version space defined by the diagonal alone and the only slack comes from dropping $B\Omega = 0$. U_τ is sharp by Proposition 4; U_{diag} is valid but need not be, because $ss^\top$ can violate $B\Omega = 0$. In both cases $[L_{\mathrm{sk}}, \min\{U_{\mathrm{sk}}, U_\tau, U_{\mathrm{diag}}\}]$ contains $q^\star$ and is never infinite.

C.6.2 What each one buys

Under each model the stop behaves as follows, summarized in Table 11. The measured trace and the spectral budget it implies move nothing: median 14 actions and 0 of 103 Grams improved. The oracle spectral radius improves 3, oracle residual eigenspaces at ranks 2, 4 and 8 improve 4, 6 and 6, and the free diagonal improves 8, the best of any. No model moves the median below 13. Every model here is a subset of the unconstrained space by construction, so none can be worse, and the informative quantity is how rarely any is better.

C.6.3 What the width survives

Table 11 collects every constrained version space and oracle we tested against the unconstrained model of Theorem 1. Each row grants the model information a deployment could not have, and none of them moves the median action count.

Table 11: Is the width missing information or a pessimistic model? Median actions to a 5% target on the frozen coordinate path over the 103 recorded Grams, the recorded set only, where each test grants the model something a deployment could not have. ‡Best setting $\nu = 1$; $\nu = 2, 4, 8$ give 14. †Median per-Gram ratio over these plus the six real Grams.

Test	Granted to the model	Median	Effect
Unconstrained (Thm. 1)	nothing	14	baseline
Trace and spectral budgets	measured scalars	14	improves 0 of 103
Diagonal budget	the M diagonal entries	14	improves 8 of 103
Oracle norm and eigenspaces	what no algorithm has	14	improves 3 to 6 of 103
Probabilistic residual, 95%	a matched-trace prior	$13^\ddagger$	covers 0 of 103
Oracle stopping time	hindsight on the true Gram	10	closes $1.46\times$ sooner†

C.6.4 Two explanations that are not the answer

We rule out two further explanations for the interval's width: a probabilistic residual prior, and the possibility that worst-case certification is merely conservative.

The negative result above concerns worst-case modelling, so the natural competitor is the probabilistic-numerics one (Hennig et al., 2015): put a prior on the residual and report a credible interval, which on a decaying spectrum can be far narrower than any worst case. Drawing $B = c\,PZZ^\top P$ with Z standard Gaussian of width ν, rescaled so $\operatorname{tr} B = \tau$, and stopping on the 95th percentile of $q^\star(K_{\mathrm{sk}} + B)$ over 64 draws, the credible rule buys at most one action in the median and buys it by being wrong. Its coverage is zero: on all 103 Grams the endpoint at the stopping prefix lies below the true optimum, and the most optimistic prior saves the most actions while certifying 5% and returning 74.9%. The reason is that the true residual is a tail event. At a quarter of each Gram, over 256 draws, the true optimum exceeds every draw on 100% of Grams at $\nu = 1$ by a median factor 1.273. A residual of the right size in a random direction is not what these Grams have: it concentrates on the directions acquisition has not yet spent a query on, which is close to where the worst case points. Softening the model narrows the interval only by excluding the answer.

The other explanation is that worst-case certification is merely conservative, so the plug-in decision is already accurate long before the certificate says so. Walking the same coordinate path and recording both the certified stop and the action count at which the weighting is in fact within 5% of the optimum, scored against the withheld full Gram, the oracle needs a median 10 actions against the certificate's 15. That population is 109 Grams, the 103 recorded ones together with the six real CIFAR-10 Grams,

and it is not the population of Table 11. Two statistics differ on it and we report both: the median of the per-Gram ratios is 1.46, which is the figure we quote, and the ratio of the two medians is 1.50. Restricted to the 103 recorded Grams alone the medians are 14 and 10, so the ratio of medians is 1.40 and the median per-Gram ratio 1.47. The four numbers are the same comparison read two ways on two populations, not four measurements. On the six real conflicting Grams the certificate spends a median 251 actions and the oracle 172.5, pooling the three at $M = 512$ with the three at $M = 1024$, and on no Gram does the oracle need a small fraction. At $M = 512$ alone the certificate's median is the 227 reported elsewhere. So the actions buy accuracy and not merely permission to stop. The $1.46\times$ is the honest price of the guarantee. The oracle chooses only when to stop, not what to ask, and Appendix C.8 shows query order is worth more than the stopping rule on this geometry.

Finally, the 103 Grams come from about five independent captures, so every headline count was also computed per capture. The sharp endpoint's advantage appears in all of them, at worst on three quarters of a capture's snapshots, and the constrained version spaces are flat per capture as well as pooled. The three dependent trajectories differ in absolute level, most visibly at the projected-gradient policy where the medians are 9, 28 and 14 actions before the sharp endpoint and 8, 11 and 7.5 after, so the pooled figure of 22 is a summary of five numbers rather than of 103.

C.6.5 A WEAKER TARGET DOES NOT MAKE THE CERTIFICATE CHEAP

We ask whether weakening what is certified, from objective accuracy to common descent, makes the certificate affordable.

An obvious response to an expensive certificate is to weaken what is certified: instead of the objective within 5%, ask only for a direction along which no task loss rises. The margin $m = \min_i(K\lambda)_i$ of Equation (22) is exact from the actions already taken, needs no completion and no lower endpoint, and $m > 0$ certifies common descent under the true Gram. On the same frozen coordinate paths over all 103 recorded Grams we record the first action count at which $m > 0$ and at which the certified gain reaches 0.5, 0.8 and 0.95, beside the 5% objective stop.

Table 12: Actions to certify each target, median per panel, frozen coordinate path. Uniform descends counts Grams on which the uniform weighting already has $\min_i(K\mathbf{1}/M)_i > 0$, so no recovery is needed for descent at all.

Panel	n	M	Obj. 5%	$m > 0$	$m/U \geq 0.8$	$m/U \geq 0.95$	Unif. descends
Historical EOS	60	32	13	1	10	14	60/60
Answer only	21	32	18	3	16	19	20/21
Paired EOS	21	32	16	2	14	17	21/21
ResNet	1	100	100	100	100	100	0/1

The descent target is met at a median of one action on the three $M = 32$ panels, a fourteenfold reduction, and the saving is worthless: on 101 of those 102 Grams uniform already descends, so a practitioner wanting only common descent should run no recovery. On the one panel where uniform does not descend, the ResNet Gram at 54% negatively correlated pairs, the descent certificate needs all 100 actions, exactly as the objective certificate does. The weaker target is cheap precisely where it is unnecessary and expensive precisely where it is needed, because certifying anything about a conflicting Gram requires seeing most of it.

Where uniform fails to descend, the retarget does pay. On the six real CIFAR-10 groups uniform ascends on some task in all six, at normalized gain -1.70 to -4.17.

Table 13: Real CIFAR-10 groups under a pretrained ResNet-18, frozen coordinate path, three seeds per size, medians. Ratio at the descent stop divides the achieved objective by the full-Gram reference. Uniform fails to descend on all six Grams, so the descent certificate is not vacuous here.

M	Objective 5%	Descent $m > 0$	Descent as % of M	Ratio at descent stop	Uniform gain
512	227	**55**	**10.7%**	1.99	-3.00
1024	330	**77**	**7.5%**	1.94	-3.87

Two things have to be said together. A direction that descends on every one of 1024 conflicting objectives costs 7.5% of the actions a full Gram needs, and the cheap alternative does not merely

do worse, it ascends. But at that stop the objective is about twice the MGDA optimum and the certified margin is thin, normalized gain 0.02 to 0.04. So what that buys is a feasible certified descent direction, not a good approximation of the optimum. The target should be chosen by which of the two a deployment needs, and the mean pairwise cosine tells a user which regime they are in before any action is spent.

C.6.6 THE TARGET IS FINER THAN THE QUANTITY'S OWN RESOLUTION

We measure how stable the minibatch MGDA optimum is, and what a noise-calibrated target would cost.

Every certificate here is conditional on one minibatch's Gram, which is only meaningful if the certified quantity is stable to better than the target. It is not. Fixing an ImageNet-pretrained ResNet-18 with a ten-way head and $M = 512$, resampling only the examples over six independent minibatches of two examples per objective, and solving MGDA exactly on each, the median cross-minibatch factor $q_{K_i}(\lambda^\star_{K_j})/q^\star_{K_i}$ over all 30 ordered pairs is 5.07, range 4.07 to 6.92. Against a sixteen-example reference the median is 2.86. Set beside a 5% target that is nearly two orders of magnitude, so the field, ourselves included, has been certifying the minibatch MGDA objective to a precision the population quantity does not possess. It also offers a mechanism for the null training result: chasing 7 mrad of fidelity to a direction itself a factor of three from the larger-batch answer should not be expected to change a trained model.

The constructive consequence is that ϵ should be set from the noise floor rather than by convention.

Table 14: Median actions to close the relative certificate at conventional and noise-calibrated targets, three seeds per size. The rightmost column is the fraction of a full Gram. Uniform's own objective ratio on these Grams is 2.94 to 6.46.

M	$\epsilon = 0.05$	ratio 2.00	ratio 2.86	ratio 5.07	ratio 2.86 as % of M
512	227	140	**114**	89	**22.3%**
1024	330	181	**150**	119	**14.6%**

Calibrated to the large-batch gap the objective certificate closes at 22.3% and 14.6% of a full Gram, where the conventional target needed 44% and 32%. But calibration does not rescue the schedule: replaying the block schedule at the calibrated target leaves it refusing on all six with the same action count, because descent and not the objective is what binds, and dropping descent closes the objective certificate on all six at a median 244 actions while three of those six weightings ascend on some task. The looser 5.07 is not usable either, since uniform already achieves 2.94 to 6.46 on these Grams and would pass on four of six. The useful calibration is the large-batch gap. This is one architecture, one dataset, six minibatches, and objective ratios rather than direction angles.

C.7 TESTING THE REGIME BOUNDARY OUT OF SAMPLE ON REAL DATA

We test whether a boundary fitted on synthetic label noise predicts behaviour on real data. It does, and it predicts failure.

Our label-noise sweep locates the edge of the operating regime on the synthetic panels, and finds certification fails past effective rank 3.15. That boundary was fitted on synthetic Grams, so it is fair to ask whether it means anything off that family. This appendix tests it on data the boundary was never calibrated on, and reports the answer whichever way it comes out.

Setup. Real CIFAR-10 images with their real labels, an ImageNet-pretrained ResNet-18 with a ten-way head, batch normalization affine frozen, and M objectives each scoring a disjoint pair of real examples. This is the per-example-group geometry that appears in data attribution and influence estimation, which is where large M actually arises. Everything else matches the large-M study: same harness, same tile budget, same numerical gates, three seeds differing in the sampled examples.

The Grams are nothing like the synthetic ones. At $M = 512$ the effective rank is 57.8, 63.0 and 60.3 across seeds, mean pairwise cosine 0.013 to 0.022, and 57 to 63% of gradient pairs are negatively correlated. At $M = 2048$ the effective rank is 66.1. For contrast the synthetic panels at

the same M sit at effective rank 1.5 to 2.1 with mean cosine 0.63 to 0.81 and under 4% negative pairs, and an untrained network on the same real images gives effective rank 23.2. So it is the real data and labels, not the architecture or the objective count, that produce conflict: these objectives genuinely disagree.

The boundary predicts refusal, and refusal is what happens. Effective rank near 60 is far past 3.15, so the sweep predicts the certificate cannot close. It does not, on any seed.

Table 15: Real CIFAR-10 groups under a pretrained ResNet-18 at $M = 512$, one A100-SXM4-40GB, medians of three warm repetitions over three seeds. The rule refuses on every seed, as the synthetic boundary predicts, and returns *not certified* rather than a direction it cannot vouch for.

Route	Seconds	Peak MB	Actions	Passes	Certified
Materialized	1.3841	24,692	512	1	3/3
Contraction	**0.9783**	**1,813**	512	1	**3/3**
Acquisition rule	2.9634	2,163	513	6	**0/3**

Why 513 actions at $M = 512$ does not contradict Corollary 1. The corollary forces $L = U = q^\star$ once the sketch reaches rank M, so a run charged 513 actions on a 512-objective problem looks as though it must certify. It does not, because the count is cumulative over a restarting schedule rather than a count of distinct directions. The predict-then-double rule re-queries a prefix from the start on each failure, spending $16 + 32 + 64 + 128 + 256 = 496$ actions over five passes and then 17 more before the budget cap, so the largest prefix it ever completes is 256 coordinates and the realized rank never exceeds half of M. The rank clause is therefore never triggered. The returned weighting is nonetheless accurate: its achieved objective is within 0.53%, 0.74% and 0.08% of $q^\star$ on the three seeds, so the rule fails to *certify* a direction that is in fact good, which is the conservativeness of the certificate on this geometry rather than a failure of the direction.

Three readings, and the first is against the acquisition route as deployed. On real data at this scale the block rule is the worst of the three: it spends every action over six passes, takes three times contraction's time, and certifies nothing. The speed result of Section 5 does not transfer here, and we say so rather than leaving the synthetic panels to imply that it would.

That failure is the schedule, not the certificate. Run instead on the frozen coordinate path, the same 5% objective certificate closes at a median 227 actions of 512 and 330 of 1024, so the information is there at 44% and 32% of a full Gram and the floor-16 doubling schedule overshoots straight past it. The block schedule's guarantee is on actions relative to the path's own $k^\star$ and it holds; what fails is the predictor that sets the first block. On conflicting geometry the coordinate path is the better acquisition order, and Appendix C.6.5 shows the descent target closes there at 7.5% and 10.7% of M.

The second reading is the one that matters. The boundary was fitted on a synthetic label-noise sweep and used, without refitting, to predict behaviour on a different architecture regime, different inputs, different labels and a pretrained rather than random network. It predicted this schedule would fail and it did, on three seeds of three, and the run reported that rather than returning its weighting. But the same certificate closes on the coordinate order at 227 actions, so the boundary is a statement about the schedule and not about the information available: read correctly it is a pre-flight test for which acquisition order to use. Told that the observations do not support early stopping, the user contracts, and contraction here is $1.41\times$ faster than materialization at $13.6\times$ less peak memory, both exact. So the two halves compose as the theory says, the interval deciding the route and the route it chooses winning. What it does not do on real data is save time by stopping early.

C.8 Certificate and policy are complements, not substitutes

We separate two effects that are usually confounded: how much the query order is worth, and how much the stopping rule is worth. We run the full factorial of both.

Two comparisons in this paper vary one thing at a time. The paired endpoint audit holds the queries frozen and varies the lower bound; the matched-policy control holds the target fixed and varies the query policy, where projected gradient beats our acquisition rule on two of four panels. Neither

says whether the sharp endpoint still earns its keep once a good policy chooses the queries, which is the question a practitioner has, so this appendix runs the missing factorial. We capture the query sequence a projected-gradient policy actually issues, then replay that sequence and evaluate both stopping tests on it, so policy and certificate vary independently. The coordinate arm is the frozen path used elsewhere. All four cells stop at the same 5% relative target on the same 103 recorded Grams, and the endpoint is valid on the dense queries because Theorem 1 is stated for an arbitrary Ω.

Table 16: Median actions to the 5% objective target. Rows are the query policy, columns the stopping test. The sharp endpoint helps more on projected-gradient queries than on coordinate queries, and most on the one panel whose objectives genuinely conflict. Projected-gradient queries are dense and its backtracking probes each count as an action, so unlike the coordinate path a cell may exceed M without the sketch having spanned $\mathbb{R}^M$; the 134 below is one such case.

Panel (n, M)	Coord. + FW	Coord. + sharp	PG + FW	PG + sharp
Historical EOS (60, 32)	15	13	13.5	**8**
Answer only (21, 32)	21	18	32	**20**
Paired EOS (21, 32)	17	16	24	**12**
ResNet (1, 100)	100	100	134	**28**
All 103	16	14	22	**11**

The endpoint is worth more once the policy is good. At frozen coordinate queries the sharp endpoint cuts the median from 16 to 14, a 12% saving that invites the reading that it is second order beside query selection. On projected-gradient queries the same endpoint cuts 22 to 11, a factor of two, and it improves 92 of the 103 Grams. The two ingredients are complements: the policy chooses informative directions but stops on a bound too weak to notice it has enough, and the coordinate order stops on a sharp bound but feeds it uninformative directions.

This overturns the paper's own worst case. The ResNet Gram is the only panel whose objectives genuinely conflict, at 54% negatively correlated pairs, and it is the panel on which the frozen-path audit reports the sharp endpoint saving nothing at all, 100 actions of 100. That is an artifact of the coordinate order. With projected-gradient queries the same certificate closes at **28 actions of 100**, and the recovered weighting is not merely certified but good: against the withheld full Gram its objective ratio is 1.0139, inside the 5% target with room, its worst-task gain is $+0.482$, and the certified interval at that stop is $U/L = 1.047$. So conflicting geometry does not defeat the interval, it defeats the coordinate acquisition order, and pairing the interval with a policy that chooses its own directions recovers a $3.6\times$ saving on exactly that case.

What this does not establish. Actions are not time. The pass constant of 126 to 166 columns still applies, and projected-gradient queries are dense and do not batch into passes the way a coordinate block does, so we make no wall-clock claim for that cell and have not implemented it on the GPU harness. The ResNet cell is one Gram. And this is a stopping comparison on a fixed query sequence, not a new acquisition rule: designing a policy that is jointly pass-economical and endpoint-aware remains open.

C.9 CelebA-40: the benchmark, and the substitution this paper is about

We give the CelebA protocol in full, the per-batch geometry behind Table 17, and the fixed-budget training comparison. The geometry study measures how faithful each cheap rule is at a checkpoint. The training study asks whether that fidelity changes a trained model.

Everything above measures either synthetic Grams or per-example groups, and neither is multi-task learning. This section runs the standard benchmark, at its own number of objectives and batch size, against the substitution the paper opens by naming.

Setup. CelebA with all forty binary attribute tasks, a shared ImageNet-pretrained ResNet-18 trunk and forty linear heads, twenty thousand images at 64×64, the trunk trained four hundred steps at batch 128. All forty tasks share one minibatch, so resampling perturbs every task's gradient jointly, unlike the disjoint groups used elsewhere. The MGDA-UB study reports the median over three

independently initialized and trained trunks, each with its own minibatch, and the certificate and baseline studies use three data seeds on one trunk.

C.9.1 MGDA-UB RETURNS AN ASCENT DIRECTION ONCE OBJECTIVES MULTIPLY

MGDA-UB (Sener & Koltun, 2018) replaces the shared-trunk Gram with gradients taken with respect to the shared representation. It is the cheap substitution this literature actually deploys, and the one our opening question is about. Exact MGDA's normalized worst-task gain is 1.000 by construction, so the question is what the surrogate returns in its place.

Table 17: MGDA-UB against exact MGDA on CelebA at batch 512, scored on the exact shared-trunk Gram. Median over three independently trained trunks, with the range across trunks for the gain and the count of trunks on which the returned direction ascends on some task.

M	Obj. factor	Worst-task gain	Ascends	Angle	$q^\star/\overline{K_{ii}}$
2	1.022	$+0.849$ $[+0.78, +0.94]$	0/3	8.4°	0.406
5	1.144	$+0.796$ $[+0.09, +0.81]$	0/3	20.8°	0.078
10	1.446	$+0.546$ $[+0.24, +0.60]$	0/3	33.7°	0.049
20	1.547	$+0.303$ $[-0.40, +0.40]$	1/3	34.2°	0.025
40	1.840	$\mathbf{-0.909}$ $[-1.36, -0.39]$	**3/3**	38.5°	0.016

The objective factor rises with M at every batch size, and the gain falls with M. The robust statement is at the ends: at two objectives the surrogate never returns an ascent direction on any trunk at any batch size, and at forty it does so on all three trunks at all three batch sizes, with median gain -4.81, -0.13 and -0.91 at batches 32, 128 and 512. Between those ends the crossing point moves with the trunk and the batch, so we give counts rather than a threshold. At two tasks the surrogate is nearly exact, which is presumably why it has survived. This is the paper's thesis where it can be checked: partial geometry suffices at small M and fails at large M, and the failure is invisible to the method.

The target is not vacuous here. On the same Grams uniform scalarization also ascends, at worst-task gain -0.318, and PCGrad and CAGrad descend but sit 38.3° and 32.0° off the exact direction at objective factors 1.881 and 1.586. No cheap method on this benchmark reproduces the exact direction, and two of the four share MGDA-UB's failure.

The last column is the mechanism. Relative to the mean diagonal, $q^\star$ falls from 0.406 at two objectives to 0.016 at forty: as objectives multiply the min-norm point collapses toward zero and the geometry approaches Pareto-stationarity, which is what makes the surrogate's error catastrophic in relative terms. It is also why the cross-minibatch objective ratio is not comparable across settings. At $M = 10$ and batch 512 that ratio is 25.15 on the MNIST one-vs-rest construction and 1.157 here, because the MNIST tasks are near-mutually-exclusive and their $q^\star/\overline{K_{ii}}$ is about 0.012 against CelebA's 0.049. The ratio tracks degeneracy of the min-norm point, and M is a proxy for it rather than the variable itself.

What the certificate costs, and the noise floor. At batch 128 and 512 the certificate does close, but at a median 33 actions of 40, at effective rank 21.8 and 17.0 with 40 to 47% of gradient pairs negatively correlated. That is far outside the regime where acquisition is cheap, so the certificate closes only by buying most of the Gram and the routing rule sends the user to contraction on cost, which is what Section 6.3 predicts for geometry like this. The hindsight oracle of Appendix C.6.4 needs 31 and 24 actions on the same paths, so certification overhead here is 1.06 to $1.37\times$ against the $1.46\times$ measured on the recorded set. At $M = 40$ and batch 512, resampling only the examples puts the cross-minibatch factor at 1.257, not the fivefold figure the disjoint per-example panels give, so the conventional 5% target is several times finer than the resolution here rather than a hundred times. The instability is specific to the min-norm point: a fixed uniform weighting scored across the same sibling minibatches moves far less.

C.9.2 DOES FIDELITY BUY ANYTHING? CELEBA-40 UNDER A FIXED BUDGET

Everything above is geometry. The question that decides whether any of it matters is whether a more faithful MGDA direction produces a better model, and the language-model training study points the other way. This runs the comparison properly.

Protocol. The same trunk and heads, 16k train and 4k held out, 1200 steps at batch 128, SGD at learning rate 0.02 with momentum 0.9. Identical learning rate, schedule, data order and initialization across all five arms, so only the weighting rule differs. Table 18 reports two seeds; a third paired exact-versus-MGDA-UB run from a separate additional-seed campaign is used only for the sign test below and is never pooled with the table. Worst-attribute accuracy is reported beside the mean because worst-task protection is what MGDA exists to provide. The representation Gram is computed in closed form, checked against forty backward passes to 2.1×10^{-7}.

Table 18: CelebA-40 held-out accuracy after 1200 steps, mean of two seeds. Arms differ only in the weighting rule.

Arm	Mean accuracy	Worst-attribute accuracy	Seconds
Uniform	0.8586	0.6541	53
PCGrad	0.8591	0.6547	66
CAGrad	**0.8609**	0.6702	60
MGDA-UB	0.8558	0.6734	90
Exact MGDA	0.8552	**0.6778**	726

Exact MGDA is first on worst-attribute accuracy among all five arms in both seeds, and against uniform the margin is $+2.36$ points, $+2.15$ and $+2.58$ by seed. So the certified quantity does correspond to something: on the metric MGDA exists to optimize, the exact rule wins and the cheapest rule is also the worst. We do not claim the ranking tracks fidelity, because it does not. CAGrad is the more faithful of the two cheap rules on both measures in Table 17, at objective factor 1.586 against 1.840 and angle $32.0°$ against $38.5°$, and it descends where MGDA-UB ascends, yet finishes below MGDA-UB on worst-attribute accuracy. Fidelity picks out the winner here and does not order the field. Exact MGDA is also last on mean accuracy, -0.34 points against uniform at 13.7 times its wall clock, so a practitioner who wants mean accuracy should not use MGDA at all.

The margin over MGDA-UB is unresolved, not refuted. Against MGDA-UB the margin is $+0.44$ points, $+0.02$ and $+0.85$ by seed, while MGDA-UB's own spread is 0.0098, more than twice the 0.0044 margin. This protocol cannot resolve that comparison, which is a statement about the protocol and not a finding against exact recovery. The sign is consistent: exact leads MGDA-UB in all three paired runs, by $+0.0002$, $+0.0085$ and $+0.0010$, a one-sided sign test at $p = 0.125$. Separating a 0.0044 margin at conventional power would take of order twenty seeds per arm, 4 GPU-hours for the exact arm alone, and we have not spent it. The margin over uniform, 0.0237 against uniform's own spread of 0.0058, needs none of that. So we claim the ordering against uniform and leave the ordering among the cheap rules open. It runs opposite to the language-model study, where uniform has lower worst-task CE than ParetoTrace-Direct on all three paired seeds, and we report them as one for and one against.

Limits. Two seeds in the table and a third paired run for the sign test only, 1200 steps, 16k of 162k images at 64×64, one architecture, one learning rate shared across arms. The geometry and the training run at different batch sizes: the -0.909 we headline is batch 512, while training runs at batch 128 where the median gain is -0.13, so the informative missing run is training at batch 32. The exact arm's 726 s prices forty separate backward passes rather than ParetoTrace-Direct's contraction, so the wall-clock penalty overstates what this paper makes available. And our MGDA-UB uses the same exact simplex solver rather than the Frank–Wolfe of Sener & Koltun (2018), which makes the geometric criticism conservative and the training comparison harder on us. No $\Delta m\%$ against single-task baselines, which would need forty more runs.

C.10 NYUv2: THE THREE-TASK BENCHMARK, MEASURED RATHER THAN ARGUED

We report the NYUv2 pilot that prices the certificate at three objectives.

Section 6.1 argues that dense-prediction suites sit outside the regime this paper studies. A measurement on the benchmark is stronger than an argument from our own $M = 2$ column, so we report one.

Protocol. SegNet on NYUv2-MTAN, the standard three-task setup of semantic segmentation, depth and surface normals, batch 2, one epoch, 398 updates, one seed. Only the recovery arm was run.

Per-update telemetry records the actions the certificate spent, the interval ratio, whether the returned direction descended on every task, and the wall clock split between selecting a weighting and taking the step.

Table 19: NYUv2 with SegNet, 398 updates. A full Gram is three actions because there are three tasks. The interval is already closed at the median, every update already descends on all three tasks, and selecting the weighting costs several times the update it protects.

Quantity	Value
Actions per update, median (of 3)	2
Updates needing the full Gram	195 of 398
Updates closing in one action	25 of 398
Interval ratio U/L, median	1.000000
Interval ratio U/L, maximum	1.049
Common descent observed	398 of 398
Selection time, median	0.2075 s
Rest of the step, median	0.0699 s
Complete update, median	0.2774 s
Selection as a share of the step	75%

What this settles: the cost, not the geometry. The certificate spends a median of two actions where a full Gram costs three and needs all three on 195 of 398 updates, so the largest possible saving is one backward pass in three and it is unrealized on roughly half the updates. Selection takes 0.2075 s of a 0.2774 s step, so the step without it would be 0.0699 s and running the certificate makes it about four times longer. At three tasks the ceiling on the saving is one third and the price is a factor of four, which is a sufficient reason not to deploy a stopping certificate there. This part needs no baselines and we claim it.

What this does not settle. The interval ratio of 1.000000 at the median is close to definitional, since 195 of 398 updates bought complete coverage where Theorem 1 forces the ratio to one, and the rest stopped under a rule that cannot report much above 1.05. The common descent on all 398 updates is a property of one epoch at batch 2, the first half percent of the standard MTAN protocol, on a network at mean IoU 0.065. Conflict in dense prediction is generally reported to emerge as tasks specialize, so an untrained SegNet showing none is weak evidence that a trained one shows none, and we do not claim three-task dense prediction is conflict free. The pilot has no baseline arms, so it supports no accuracy or $\Delta m\%$ claim, and it cannot separate the proposition we defend, that uniform is already adequate on NYUv2, from the weaker fact that uniform is adequate at initialization. We also cannot tell whether the selection time includes the per-task backward passes MGDA needs in any case, so the factor of four may price exact MGDA rather than the certificate on top of it.

D How tight is the direction-error certificate?

Theorem 4 bounds the squared recovered-direction error by the interval width. A guarantee that holds but is orders of magnitude loose would say little about what a user actually gets, so we measure both sides on all 103 recorded Grams.

Protocol. Each Gram is scaled by $\mathrm{tr}(K)/M$ and solved exactly to obtain $\lambda^\star$ and $q^\star$. Coordinate recovery then runs on actions alone and stops at the first action count where the observed feasible objective is within 5% of the completion lower bound, giving λ_{sk} and the width $U - L$. The realized error is $\|G\lambda_{\mathrm{sk}} - G\lambda^\star\|^2 = (\lambda_{\mathrm{sk}} - \lambda^\star)^\top K(\lambda_{\mathrm{sk}} - \lambda^\star)$, computed after selection and never visible to the recovery run. Records accompany the submission.

Result. The bound holds on all 103 Grams, which is the numerical check on Theorem 4 in this regime. It is conservative by a median factor of 33 and by at least 5.2 on every Gram, so a 5% objective stop delivers a direction much closer to the exact MGDA direction than the certificate promises. Relative to $q^\star$, the certified width has median 0.027 and maximum 0.048, while the realized squared error has median 7.8×10^{-4} and maximum 4.8×10^{-3}. The three trajectory panels agree closely, at median conservatism 29.5, 28.8 and 30.4.

Reading. Two things follow. The certificate is safe, because nothing exceeded it. It is also a bound rather than an estimate, so a user who needs the realized error should read the width as an upper limit roughly thirty times the typical outcome in these panels, not as a prediction of it. We do not claim that factor transfers to other architectures or loss constructions.

E MATCHED BUDGET, AND WHERE THE CERTIFICATE STOPS WORKING

Two questions decide how much Section 5's result is worth. Does the adaptive rule beat a fixed budget that spends the same actions, and how far does the regime extend? Both are measured on the same A100 harness.

E.1 A FIXED BUDGET MATCHED TO THE RULE'S OWN SPEND

The rule spends a median 50 actions, so the control is a fixed 50-action budget. Nesting in Corollary 1 makes the certificate monotone, so one might expect a fixed 50 to certify wherever the rule does. It does not.

Table 20: Matched-spend control at the same three sizes, three seeds each, medians of three warm repetitions. The fixed budgets stop at their action count whatever the interval says; the rule stops on the interval.

Rule	Actions	Passes	Certified	Worst factor	Min gain
Fixed 50 actions, one pass	51	1	0/9	2.752	0.502
Fixed 50 actions, two passes	52	2	6/9	1.401	0.447
Adaptive rule	50–244	2–4	**9/9**	**1.002**	**0.955**

A fixed 50 in one pass certifies nowhere, because splitting the same actions over two passes lets the second block be chosen from the first block's responses. A fixed 50 in two passes certifies on six of nine, and the three it misses are exactly the three seed-0 runs, the lowest-cosine seed at each size, where it returns objective factors 1.260, 1.368 and 1.401 at worst-task gain down to 0.447. The rule certifies on all nine because it does not stop at 50 when 50 is not enough: on those three seeds it escalates to 244 actions over four passes. Monotone nesting says a closed certificate stays closed, not that any particular budget closes one.

The cost of that adaptivity is small where it is not needed. On the six seeds that a fixed two-pass 50 does certify, the rule is within 1% of its time at 1024 and 2048 and 6% at 512. The adaptive part is what buys the other three, and there it costs roughly a doubling of time on that seed.

E.2 WHERE THE CERTIFICATE STOPS WORKING

The large-M Grams are near rank two, which invites the objection that the regime is constructed. Rather than argue, we sweep it. Holding $M = 1024$ and the harness fixed, we replace a fraction of the teacher labels with random ones, which raises conflict and effective rank continuously, and run both routes at each level.

The boundary is sharp and it is narrower than the certification boundary alone. Certification survives to effective rank 2.41 and fails from 3.15 on. The speed advantage is narrower still: only at effective rank 2.08 does the rule come near contraction on this seed, and by 2.41 it certifies but takes $2.6\times$ contraction's time because it has spent almost every action anyway. So the honest statement of the operating regime is effective rank near two, not merely agreeing gradients, and Section 5's numbers should be read that way.

What does not degrade is the reporting. At every level from 0.25 up the rule returns *not certified* rather than a direction it cannot vouch for, and those are exactly the levels where its returned weighting is bad: objective factor 1.005 to 2.584 with worst-task gain as low as -0.736, meaning the direction raises some task's loss. This is the impossibility reading of Theorem 1 doing its job: the user is told the observations do not support early stopping and should contract instead. That costs time rather than saving it. At noise 1.0 contraction takes 15.35 s against the uncertified acquisition run's 6.82 s, so what the certificate buys past the edge is correctness at $2.3\times$ the time, not speed. Records are in

Table 21: Label noise sweep at $M = 1024$ on seed 0, the lowest-cosine seed, so the noise-0 row matches that seed's entry in Table 8 rather than the end-to-end cost table's median over three seeds. Effective rank is the participation ratio $(\sum\lambda_i)^2/\sum\lambda_i^2$ of the Gram spectrum. Certified is whether the rule met its own 5% target with the gain condition before exhausting the objectives. Timings are from this sweep's own runs, so they differ from Table 8's capture by about a percent at the matching cell (3.01 against 2.9827 s). Contraction's time rises with noise because forming the Gram is noise independent while the simplex QP on a higher-rank spectrum is not, so the growth is in the solve rather than in the contraction.

Noise	Mean c	Eff. rank	Neg. pairs	Rule actions	Certified	Rule s	Contraction s
0.00	0.626	2.08	4.1%	244	yes	3.01	2.26
0.10	0.526	2.41	8.6%	1014	yes	5.92	2.27
0.25	0.394	3.15	17.9%	1014	**no**	6.38	2.33
0.50	0.212	5.46	36.3%	1014	**no**	6.73	2.99
0.75	0.075	11.41	54.6%	1014	**no**	7.39	7.17
1.00	0.013	18.36	63.2%	1014	**no**	6.82	15.35

F When the interval saves information: gradient agreement

Section 6.2 reports that the completion endpoint reduces actions on seven of eight panel–path cells and not on the eighth. This appendix identifies the variable that decides it, using the audit script and stopping protocol of the paired endpoint audit without modification.

The predictor. For a Gram K write $c = \mathrm{mean}_{i<j}\, K_{ij}/\sqrt{K_{ii}K_{jj}}$, the mean pairwise cosine of the task gradients, and let ν be the fraction of pairs with $K_{ij} < 0$. Both are read off K and neither needs a solve. Over the 103 recorded Grams, the actions the completion endpoint needs at a 5% objective target on the uniform-first path correlate with c at Pearson $r = -0.784$ and with ν at $r = +0.744$; normalized effective rank, by contrast, reaches only $r = +0.277$ and is not the operative variable.

Table 22: Recorded panels ordered by gradient agreement. Actions are medians of the completion endpoint at a 5% objective target on the uniform-first path, the same quantity Section 6.2 reports.

Panel	n	Mean cosine c	Negative pairs ν	Actions$/M$
`historical_eos`	60	0.188	1.0%	0.375
`paired_eos`	21	0.154	2.0%	0.469
`answer_only`	21	0.097	8.9%	0.500
`resnet100`	1	−0.005	54.0%	0.890

The ordering is monotone, and it reframes the one null cell. The `resnet100` Gram is not merely the largest problem we test; it is the only one whose tasks genuinely conflict, with 54% of pairs at negative inner product against 1–9% elsewhere. On that Gram the endpoint saves nothing (100 actions of 100 on the coordinate path), which the predictor anticipates.

The law across the full range. The recorded Grams span only $c \in [-0.005, 0.238]$, so we extend the range with a synthetic shared-component family, $G = c\,u\mathbf{1}^\top + \sqrt{1-c^2}\,Z$ at $M = 32$ with $P = 4000$, and a conflicting variant in which half the objectives oppose the shared direction. These 56 Grams are scored by the same script, with the same stopping rule, and released with the records.

The recorded Grams give Pearson $r = -0.784$, and they are not 103 independent problems: 60 are checkpoints of three trajectories, 42 are the same 21 checkpoints scored under two loss constructions, and one is the ResNet Gram, so the effective design is four panels. The synthetic family alone gives -0.906, but it is generated to vary c and is a range extension rather than evidence. Pooled, $n = 159$ with Pearson -0.862 and Spearman -0.812. The practical reading is a pre-flight test. Gradient agreement is computable from the Gram diagonal and one dense action, before any acquisition schedule is chosen, and it predicts whether the interval will close early: at $c \gtrsim 0.5$ a single action suffices at $M = 32$, at the agreement our recorded trajectories exhibit ($c \approx 0.10$–0.19) between a third and a half of the coordinates are needed, and at $c \approx 0$ with substantial conflict the sketch buys almost nothing. This is a statement about information, not about wall time, and it is measured at one objective count on one target; the mechanism is that agreeing objectives place $\lambda^\star$ near the interior of a low-dimensional face that few actions already resolve.

Table 23: Synthetic sweep at $M = 32$, four draws per setting, medians. Same audit script, path and target as Table 22.

Mean cosine c	Negative pairs ν	Median actions	Actions/M
-0.00	51.1%	20.0	0.625
0.02	7.9%	17.5	0.547
0.04	0.5%	14.0	0.438
0.09	0.0%	13.5	0.422
0.16	0.0%	8.0	0.250
0.25	0.0%	5.5	0.172
0.36	0.0%	3.0	0.094
0.49	0.0%	3.0	0.094
0.64	0.0%	1.0	0.031

F.1 Agreement decides cost, not whether MGDA is needed

Acquisition is cheapest where task gradients agree, which invites the objection that agreement is also where plain scalarization would do. The two are separate questions and we measure both on the same Grams. For every Gram in the release we score uniform weights $\lambda = \mathbf{1}/M$ against the full-Gram MGDA solution, giving the objective ratio and the normalized worst-task gain already used elsewhere.

On the 103 recorded neural Grams uniform costs a median 3.13 times the MGDA objective, at least 1.86 and as much as 5.16, with worst-task gain down to -4.78 and two Grams on which uniform raises some task loss. Those Grams span mean cosine -0.005 to 0.238. On the 56 synthetic Grams uniform costs between 1.004 and 1.014 at every cosine in the family, including $c \approx 0$ and $c \approx 0.64$. Uniform is therefore nearly optimal on one family at all agreement levels and far from optimal on the other at all agreement levels, so agreement is not what decides it. Pooled over all 159 Grams the rank correlation between c and uniform's objective ratio is $+0.21$.

What does track it is the spread of gradient norms $\max_i \|g_i\| / \min_i \|g_i\|$, read straight off the Gram diagonal. Its logarithm correlates with uniform's ratio at Pearson $+0.90$ and Spearman $+0.83$ pooled, and the synthetic family has spread 1.03 to 1.06 against 3.87 to 15.53 for the neural Grams. The two families differ in both properties, so most of that pooled correlation is between families rather than within. Within the 103 neural Grams alone the effect survives at Pearson $+0.56$ while agreement still predicts nothing, at Spearman -0.003.

The practical reading is that the diagonal answers two different questions. Mean pairwise cosine predicts how many actions certification will take. Norm spread predicts how much the MGDA weighting differs from uniform in the first place. A deployment with agreeing gradients and equal norms should use uniform weights and skip this paper. A deployment with agreeing gradients and unequal norms is where recovery is both cheap and worth doing, and the recorded trajectories here have unequal norms throughout. Records accompany the submission.

G The sharp endpoint in end-to-end training

Earlier versions of this work evaluated L_{sk} only in replay. This appendix reports a training run whose stopping decision is made by L_{sk} itself, the matched comparison against the cheaper rule that previous trained runs used, and the implementation barrier that had kept the two apart.

Protocol. Pythia-160M, $M = 16$ tasksource objectives, one example per task per update, 100 updates, one seed, FP32, on one A100-SXM4-40GB. The update is plain scalar SGD at $\eta = 3 \times 10^{-4}$ with no clipping, so the step is exactly $\theta + \eta d$ with $d = -G\lambda$ and Proposition 1 applies verbatim. Acquisition is matrix free: each action is one weighted reverse pass followed by one forward-mode pass, skipping tasks with zero coefficient. Recovery stops at the first action count where $U/L_{\mathrm{sk}} \leq 1.05$. Every step then completes the basis to rank M and audits the returned weighting against the completed Gram, which is separate instrumentation and is not charged to the method.

The run. The certificate is sound. Across 100 steps there is no false stop, the interval contains $q^\star$ at every step, and the achieved objective factor never exceeds 1.0002. Recovery costs a median

6.24 s per step and 609 s in total, covering acquisition, the endpoint solves and direction construction. Training reduces the worst task loss from 257.10 to 4.94 and reaches held-out worst-task loss 6.15 at token accuracy 0.831.

Two honest limits. The run stops before full rank on only 19 of 100 steps, so at $M = 16$ the sharp endpoint buys the guarantee rather than an action saving. And the recovered margin is positive on 63 of 100 steps. The other 37 are exactly the steps whose measured q is itself non-positive, which an exact PSD Gram cannot produce. At full rank simplex stationarity makes the margin equal $q^\star$, and the two agree throughout, so these are float32 round-off once the min-norm point collapses toward Pareto-stationarity, not a failure of the rank clause: after step 60 the median q is 9×10^{-4} while the negative excursions reach -1.8. All 19 early stops fall in this group, since the rule can only close when the upper endpoint is itself near zero. Where the sign is unresolvable at this precision, Proposition 1 has no hypothesis to act on.

Matched comparison at identical queries. Each step stored its query order and its audited full Gram, so the three stopping rules can be replayed on exactly the same queries with only the lower bound varying. This isolates the endpoint from the acquisition policy.

Table 24: Stopping rules replayed on the 100 training transcripts at identical queries, target $U/L \leq 1.05$. The cheap composite carries its gate $m/q \geq 0.95$.

Lower bound	Steps certified	Median actions	Range
L_{sk} **(sharp)**	$\mathbf{37/100}$	**11**	5–15
ℓ_{c} (cheap, gated)	$27/100$	15	10–16
Frank–Wolfe alone	$27/100$	15	10–16

On the 27 steps where both rules certify, L_{sk} is strictly earlier on all 27 and never later, at a median 11 actions against 15. It also certifies 10 steps the cheap rule cannot, while the cheap rule certifies none that it cannot. This is Corollary 1's pointwise dominance measured on a real trajectory rather than on recorded Grams. The Frank–Wolfe and composite columns coincide exactly, so the Cauchy–Schwarz term m_+^2/U never binds here and contributes nothing beyond the linearization in this regime.

Why this had not paid off live. The live rule does not use L_{sk} directly. It subtracts a numerical slack, proportional to the mean diagonal of the shorted completion K_{sk} and scaled by the measured interpolation and symmetry residuals, to keep the reported bound valid under an inexact completion. All three come from the acquired actions, so the slack uses no information beyond what the method is charged for. At these gradient magnitudes that slack is not a correction but the dominant term: over the 1,433 recorded prefixes with a positive upper endpoint its median is 7.1 times that endpoint, it exceeds it on 60% of them, and it drives the reported lower bound to zero on 1,112 prefixes whose completion primal is strictly positive. That is the whole gap between 19 live early stops and the 37 that exact arithmetic on the same queries supports.

The barrier is therefore a scaling choice in the guard rather than a limitation of the endpoint, and it is the concrete thing to fix: a slack tied to the interval width rather than to the Gram scale would leave the bound valid without swamping it. We report the measurement and do not claim the repair, which we have not run.

G.1 Choosing the step from the certificate

The run above holds the step size fixed, so Proposition 1 has no admissible step to act on and monotonicity is not tested. This study chooses the step from the certificate instead, and adds the control that isolates recovery from auditing.

The acceptance test. At each step the exact response $K\lambda$ is read off the certificate, so the directional derivative of task i is $-(K\lambda)_i$ at no extra cost. We accept the largest η on a halving grid from 3×10^{-4} satisfying, for every task,

$$L_i(\theta + \eta d) \;\leq\; L_i(\theta) - c\,\eta\,(K\lambda)_i, \qquad c = \tfrac{1}{2}. \tag{27}$$

A plain non-increase test would be vacuous, since below some η every measured change falls under any fixed tolerance and even an ascent direction is accepted. Equation (27) avoids that: for $0 < c < 1$

it is satisfiable at small η exactly when every $(K\lambda)_i$ is positive. We also floor η where the update stops being resolvable, and reaching the floor is recorded as *no admissible step*, at which the run holds its parameters rather than stepping. Both properties are verified on exact quadratics, where the guarantee holds on $296/296$ instances with a positive margin and the search correctly exhausts on $300/300$ ascent directions.

The matched control. Uniform weighting is audited with one Gram action, which yields its exact response, and is then put through the identical acceptance test. This is the comparison that separates recovering the MGDA direction from merely checking whichever direction one already has, since any weighting can be audited for one action.

Table 25: Sixty updates of Pythia-160M at $M = 16$, one seed, identical data order. The two certified arms share the identical acceptance test. The recovered certified arm found no admissible step once and held its parameters rather than moving them.

Arm	Actions/step	Step size	Every task non-incr.	Median η	Worst-task loss
Start	—	—	—	—	257.10
Recovered, fixed step	16	tuned	$48/60$	3×10^{-4}	$\mathbf{5.45}$
Uniform, audited	1	certified	$60/60$	2.34×10^{-6}	246.34
Recovered, certified	16	certified	$\mathbf{60/60}$	5.86×10^{-7}	65.91

Reading. Every step of the recovered arm has a positive margin, so Equation (27) demands a strict decrease of every task, and an admissible step is found on 59 of 60. No task loss rises on any of the 60 minibatches. The one step without an admissible step is a refusal rather than a failure: seven tasks still violated Equation (27) at $\eta = 2.3 \times 10^{-9}$, so the run held. When a response $(K\lambda)_i$ is negative, the right side of Equation (27) exceeds the current loss and the test no longer enforces non-increase, so a negative-response step of the uniform arm is observed rather than guaranteed. This is the behaviour Proposition 1 predicts, reached through a sufficient-decrease test with exact responses rather than through its step cap, since we do not estimate the β_i. Each non-increase holds on that update's minibatch and is not a telescoping guarantee for one fixed objective.

Against the matched control the recovered direction is worth its cost. Both certified arms accept a step at essentially every update and neither raises a task loss, so the guarantee is not what separates them. What separates them is progress under it: the recovered direction moves the worst task from 257.10 to 65.91 while uniform reaches only 246.34, a reduction of 191 against 11. It does so at a *smaller* median step, 5.86×10^{-7} against 2.34×10^{-6}, so the difference is the direction and not the distance travelled. Recovering it costs 16 actions and 7.89 s per step against one action and 0.99 s. This is the comparison that answers whether recovery earns its price once any weighting can be audited, and on this run it does.

The guarantee is not free in optimization terms, which the fixed-step row makes plain. The same recovered direction taken at the tuned 3×10^{-4} reaches worst-task loss 5.45 against the certified arm's 65.91 over the same sixty updates, and raises some task loss on 12 of them. So the certificate does not accelerate training. It converts an unmonitored trajectory into one where every step is checked, and the price is a smaller step and slower progress. A practitioner who wants speed should take the tuned step and give up the guarantee, and the certificate is what makes that an informed choice rather than an unexamined one.

What this does not show. One seed, $M = 16$, and 60 updates, so none of these differences carries a variance estimate. Equation (27) enforces sufficient decrease rather than monotonicity, and for a task whose response is negative it permits a bounded rise; on the one such step observed here no task in fact rose. Finally the margin is conditional on the minibatch Gram, so none of this bounds a held-out loss.

G.2 All six arms under one protocol

Table 26 runs every arm on the same seed, the same data order and the same sixty updates, so the columns are comparable. Held-out figures are the frozen evaluation set at the final checkpoint.

Three readings, and they do not all favour the method.

Table 26: Sixty updates of Pythia-160M at $M = 16$, one seed. Columns counts Gram columns touched: Direct contracts all 16 without an acquisition query, whereas Stream acquires them. Non-incr. counts updates on which no task loss rose, and is available only for arms that compute a margin. Held-out worst is the largest per-task loss on the frozen evaluation set. Starting worst-task loss is 257.10.

Arm	Columns	s/step	Non-incr.	Train worst	Held-out worst	Token acc.
Uniform	0	0.23	—	16.61	17.38	0.291
FAMO	0	0.23	—	6.14	7.60	0.426
Uniform, audited, cert. step	1	0.99	60/60	246.34	248.33	**0.663**
ParetoTrace-Direct, fixed step	16	3.31	20/60	14.48	15.22	0.563
ParetoTrace-Stream, fixed step	16	7.28	48/60	**5.45**	**6.40**	0.564
ParetoTrace-Stream, cert. step	16	7.89	**60/60**	65.91	77.48	0.399

At the tuned step this table supports no ordering. ParetoTrace-Stream reaches the lowest held-out worst-task loss of the six, 6.40, against FAMO's 7.60, Direct's 15.22 and uniform's 17.38, but on one seed and equal update counts, while Stream costs 7.28 seconds per update against FAMO's 0.23. Direct and Stream compute the same direction whenever Stream reaches full rank. At the first shared state their weightings differ by 1.22×10^{-4} in the sup norm and both have normalized margin 1.000000. Their worst-task training losses are identical after one update, differ in the fifth significant digit after two, and separate by more than 1% after four. By then the fixed step has raised both above the start, to 297.9 and 275.6 against 257.1 after five updates. Stream also stops before full rank on 22 of the 60 steps, so its later directions are certified rather than exact. The $2.4\times$ spread between the two runs therefore measures divergence under an unstable fixed step, not a methodological advantage, and it exceeds the gap between Stream and FAMO.

Token accuracy does not follow it. The arm that barely trains, uniform under a certified step, has the highest token accuracy at 0.663 simply because it stays near the pretrained model, while every arm that reduces loss substantially gives accuracy up at this learning rate. Worst-task loss and token accuracy are not measuring the same thing here and we do not claim the second.

Cost is real and ordered as the theory predicts. Uniform and FAMO build no geometry and run at 0.23 s per update. Direct contracts all sixteen blocks at 3.31 s. Stream adds acquisition and the full-rank audit at 7.28 s, and the certified step adds the line search at 7.89 s. Direct is the cheaper route to an exact direction, which is what Section 5 recommends at this objective count.

G.3 Live stopping and descent studies on real conflicting geometry

Both studies use an ImageNet-pretrained ResNet-18 with a fresh ten-way head, frozen BatchNorm, and two CIFAR-10 training examples per objective, on one A100-SXM4-40GB with FP32 factors and TF32 disabled. Each seed fixes a head and data panel. Three states per seed follow Direct's own SGD updates at step 0.01, and every arm is evaluated at the same states in rotated order. Timed cost covers capture, all actions and solves, and direction reconstruction. The images are the 50,000 CIFAR-10 training images, read from a public parquet copy whose SHA-256 is recorded with the results.

Completion against squared-margin stopping. At $M = 100$, over seeds 7, 11, 22, 33 and 44, the two Stream arms share one query basis, full-span upper solver and lower-completion acquisition rule, and differ only in the lower bound used for the $U \le 1.05\ell$ stop. Completion needs a median 54 actions and 2.420 seconds against 79 and 3.808 for the squared-margin bound, and is better on 14 of 15 states. The gain is concentrated at warm states, 46.5 against 77 median actions, while cold states are close, 73 against 79. Direct takes 0.142 seconds. Ten native-autograd audits give factor-Gram relative error at most 1.81×10^{-7} and factor-to-native direction error at most 1.41×10^{-6}. Completion's audited objective factor is at most 1.00151 with normalized gain at least 0.9504, and uniform has a negative task response in all ten.

Descent recovery against contraction. At $M = 512$ and 1024, over seeds 11, 22, 33, 44 and 55, Stream spends one uniform action and then coordinate blocks of the $M/8$ tasks with the smallest

current response, doubling each pass, and accepts when the margin is positive and the certified gain $m/\sqrt{q\,U_{\mathrm{sk}}}$ reaches the target.

Table 27: Live descent recovery, medians over 15 states. Speedup is Direct time over Stream time at the same state.

M	Arm	Time (s)	Speedup	Faster	Passes / columns
512	Direct	1.038	—	—	1 / 512
512	Stream, strict descent	0.839	1.24	15/15	2 / 65
512	Stream, gain ≥ 0.2	2.078	0.49	2/15	2–3 / 193
1024	Direct	3.044	—	—	1 / 1024
1024	Stream, strict descent	3.205	0.95	1/15	2 / 129
1024	Stream, gain ≥ 0.2	9.092	0.34	0/15	2–3 / 385

Every Stream decision is accepted, has a strictly positive margin against the reference Gram, and has certified gain no larger than its reference gain. Five native audits at $M = 512$ agree with the factor Gram to relative error at most 5.7×10^{-4}. At $M = 1024$ acquisition remains cheaper than contraction, 1.09 against 1.65 seconds, but the constrained programs take 1.35 seconds on the host CPU, and a third pass reverses the comparison at gain 0.2.