A SECOND-ORDER COMPLEXITY THEORY FOR SPECTRAL OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Neural-network training combines approximate gradients with approximate update computations. Recent advances in spectral optimization, including Muon for large-scale training, make the balance between computational cost and update accuracy especially important. How much gradient information and arithmetic accuracy suffice to preserve optimization progress? We address this question using second-order complexity theory, which accounts for the requested accuracy and the cost of reading and processing gradient approximations. We measure spectral update quality by the shortfall in linearized descent relative to the best direction under the same spectral-norm constraint. This criterion weights weak gradient components by their small contribution to descent, even when their normalized directions are sensitive to error. We establish the exact minimax tradeoff between update sensitivity, uncertainty, and descent, attained by spectral clipping calibrated to a gradient-error bound and a supplied rank bound. On a normalized gradient class, we determine the input precision required for a target descent tolerance to within one bit. A certified matrix-multiplication algorithm achieves the guarantee in second-order polynomial bit time, uniformly across rank changes. For smooth nonconvex objectives, these local guarantees yield stationarity bounds that separate gradient uncertainty from implementation error and have optimal uncertainty dependence. These results quantify, in precise mathematical terms, how accurately gradients and updates need to be computed to preserve optimization progress, providing a complexity-theoretic framework for understanding when spectral optimization methods remain computationally feasible at high precision.

1 INTRODUCTION

Neural-network training repeatedly makes two coupled decisions: how accurately to estimate a gradient, and how accurately to turn that estimate into a parameter update. Minibatches control the work spent obtaining gradient information (Bottou et al., 2018), while mixed-precision arithmetic changes the cost and accuracy of representing and processing it (Micikevicius et al., 2018). For optimizers that transform matrix-valued gradients, both choices affect the resulting direction. The central question of this paper is: *how much gradient information and arithmetic accuracy suffice to choose an update with nearly optimal linearized descent?*

We use *second-order complexity theory* (Kawamura & Cook, 2012) to make these accuracy choices computationally explicit. It represents a real-valued input by an oracle returning finite approximations at requested accuracies and accounts for the lengths of requests and complete answers, as well as the arithmetic used to produce an output. This separates the gradient information a decision needs from the computation used to act on it. A desired descent tolerance can then determine both an input-accuracy requirement and a sufficient arithmetic budget.

This question is especially pressing in spectral optimization, where recent advances in large-scale neural-network training have sharpened the need to balance the cost and accuracy of matrix-valued update computations (Liu et al., 2025; Amsel et al., 2026a). For a matrix parameter such as a linear layer's weight matrix, a spectral-norm bound on the update controls the change applied to every layer input. Under this constraint, the polar factor of the gradient maximizes the linearized objective decrease by assigning unit magnitude to each nonzero singular direction (Bernstein & Newhouse, 2024). Muon approximately applies this transformation to momentum. The quality of a spectral

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.

update therefore depends both on the information in the matrix supplied to the transformation and on the accuracy with which the transformation is evaluated.

The distinction matters most for weak gradient components. In one dimension, gradients a and $-a$, with $a > 0$, have normalized directions separated by two, but choosing the wrong direction loses only $2a$ in linearized descent. If the observation is zero and the error radius is a, either sign is possible; choosing the zero update limits the shortfall to a for both. In a matrix gradient, each singular component likewise contributes to descent in proportion to its singular value. This leaves room to attenuate uncertain components while retaining a direction with a strong descent guarantee. The optimization objective itself supplies the weighting needed to judge approximation error.

We formalize update quality through the *descent shortfall*: the difference between the best linearized decrease under the spectral-norm constraint and the decrease achieved by the chosen direction. This quantity enters additively into the usual smoothness bound on the objective after an update. The optimizer observes a matrix within a supplied Frobenius error radius of the true gradient, and its update rule has a prescribed Lipschitz constant. The radius describes the available gradient information; the Lipschitz constant controls how strongly changes in that information can change the direction. The resulting minimax problem asks for the smallest shortfall that any such stable rule can guarantee.

The error radius connects this model to training choices. A variance bound converts a minibatch size into a sampling-error certificate, and bounds on numerical error account for gradient computation and storage. When the observed matrix is momentum, its deviation from the current gradient also includes the lag caused by changing iterates. These contributions determine the information available to the spectral subroutine. Our bounds coordinate this information with the subroutine's arithmetic accuracy: calibrated clipping achieves the optimal uncertainty contribution, and more accurate computation reduces the implementation contribution. Improving gradient information then supports a correspondingly stronger descent guarantee.

Recent work develops efficient polar approximations (Amsel et al., 2026a), spectral clipping (Jiang et al., 2026), and the smoothing benefits of finite Newton-Schulz iterations (Li & Tsuchiya, 2026). Analyses of inexact Muon also connect approximation quality to step size and momentum (Shulgin et al., 2026). We establish the best descent guarantee available from uncertain gradient information under a stability constraint, and the input precision and finite arithmetic needed to attain it. Our contributions are as follows.

- We give an *optimal calibration of spectral updates* to gradient uncertainty and a prescribed stability requirement (Theorem 1). For a supplied uncertainty radius, rank bound, and sensitivity limit, calibrated clipping attains the smallest worst-case descent shortfall over all admissible stable maps, including nonspectral maps. When no rank restriction is imposed and the sensitivity limit is optimized, calibrated clipping halves the worst-case linearized descent shortfall of polar plug-in (Corollary 2).
- We determine the *gradient accuracy and computation sufficient for a target descent tolerance* (Theorem 4). Necessary and sufficient absolute Frobenius input-precision bounds differ by one bit on a normalized gradient class. A certified matrix-multiplication algorithm attains the guarantee in second-order polynomial bit time across rank changes, specifying both the information to request and sufficient arithmetic to produce a feasible direction.
- We derive *nonconvex stationarity guarantees with sharp uncertainty dependence* (Section 4). The bounds show how gradient quality, update accuracy, and the number of iterations determine progress, clarifying where additional computation helps. Matching lower bounds identify when the uncertainty dependence cannot be improved. At fixed dimensions, we also establish the *optimal deterministic cost of a complete finite-precision optimization trajectory*, with matching upper and lower bounds (Appendix G). Coordinate updates attain this cost. For certified clipping, we give an explicit guarantee for the full trajectory, accounting for finite-precision computation throughout.
- We quantify the *precision advantage of selecting an update by its descent* (Section 5). At fixed matrix size, direct selection needs input precision growing linearly in the requested number of descent-accuracy bits. This gives an exponential precision saving over a two-stage procedure that first chooses a fractional spectral transform with small descent bias and then reconstructs it accurately.

These results show how second-order complexity theory can calibrate gradient and arithmetic accuracy to guarantee a specified amount of optimization progress.

1.1 Comparison with related work

Spectral optimization and gradient information. Norm-based steepest descent provides the optimization interpretation of polar directions (Bernstein & Newhouse, 2024). Related formulations use norm-constrained LMOs (Pethick et al., 2025) and non-Euclidean trust regions (Kovalev, 2025); with weight decay, Muon also admits a stochastic Frank-Wolfe interpretation (Sfyraki & Wang, 2026). Convergence analyses study the spectral steepest-descent rule (Li & Hong, 2025) and benefits associated with low-rank Hessian structure (Shen et al., 2026). Under a Kronecker-factored curvature model, Nguyen et al. (2026) connect spectral flattening to learning-rate stability. We build on this optimization geometry to clarify how much gradient information suffices for a prescribed descent guarantee.

Polar approximation and spectral clipping. The Polar Express optimizes polynomial approximants for polar-factor computation and develops their finite-precision implementation (Amsel et al., 2026a). SPECTRA applies spectral clipping to optimizer updates and develops Newton-Schulz-based approximations (Jiang et al., 2026). MuCon caps singular values at a prescribed threshold (Yi, 2026); dividing its clipped matrix by that threshold gives exactly the clipping map used here. Musec feeds clipped momentum into subsequent iterations and establishes convergence in nonsmooth nonconvex stochastic optimization. Its soft variant approximates a smooth saturation map through coupled Newton-Schulz iterations (Liu et al., 2026). Our contribution is the optimal calibration of clipping to a supplied gradient-error bound, rank bound, and allowed update sensitivity (Section 2).

Smoothing and projection-based descent. The clipping map also arises from classical quadratic smoothing of a support function (Nesterov, 2005). Li & Tsuchiya (2026) establish the smoothing benefit of finite Newton-Schulz iterations and a tradeoff between sensitivity and integrated error in a smoothed spectral potential. Building on the noiseless projection guarantee of Woodstock (2026), we establish the exact optimal tradeoff between update sensitivity, gradient uncertainty, and descent, with matching lower bounds over all admissible stable update rules. Appendix B gives the detailed comparison.

Gradient noise and useful update accuracy. Under a signal-plus-perturbation model, Li et al. (2026) show that momentum suppresses perturbations and improves the spectral separation and alignment of the matrix passed to orthogonalization. Fan et al. (2026) identify amplification of noisy weak singular directions in vision-language-action training and reinforcement learning with verifiable rewards, and introduce Pion's high-pass spectral filter. Huang (2026) develops a cubic Newton-Schulz schedule with lower computational cost and finds that more accurate polar reconstruction does not necessarily improve training quality. We formalize useful update accuracy through descent shortfall and certify it for every perturbation within a supplied error bound (Section 2). Sampling, numerical error, and momentum lag can all contribute to this bound.

Inexact oracles and optimization guarantees. Classical Frank-Wolfe analysis accommodates approximate linear subproblems and inexact gradients (Jaggi, 2013). For Muon, Kim & Oh (2026) analyze finite Newton-Schulz updates, and Shulgin et al. (2026) quantify the coupling between LMO accuracy, step size, and momentum. Do et al. (2026) give stochastic non-convergence examples and error bounds depending on iteration count and minibatch size for generalized Newton-Schulz Muon, including Polar Express. More broadly, Devolder et al. (2014) relate oracle accuracy to complexity and error accumulation in smooth convex optimization using approximate function values and first-order models. Bellavia et al. (2019) extend adaptive accuracy control to nonconvex optimization with inexact function and derivative evaluations, providing evaluation-complexity bounds. Monnet & Orban (2025) analyze rounding errors throughout a multi-precision optimization method and adapt evaluation precision to preserve convergence. For smooth objectives, we derive stationarity guarantees that distinguish gradient observation error from update implementation error, combining the optimal local error budget with established inexact-descent analysis and algorithm-independent uncertainty lower bounds (Section 4).

Matrix computation and finite precision. Polar-factor computation and perturbation analysis include Newton iteration (Higham, 1986), dynamically weighted Halley iteration (Nakatsukasa et al.,

2010), and Fréchet-derivative methods for full-rank matrices (Gawlik & Leok, 2017). Regularity estimates in Frobenius norm control how singular-value transformations change under input perturbations (Andersson et al., 2016). Recent training-oriented advances analyze floating-point quantization for Adam and Muon (Tang et al., 2026) and improve orthogonalization and communication through Dion3's Gram Newton-Schulz, specialized kernels, and partial-row updates (Amsel et al., 2026b). Our input-accuracy bounds are necessary and sufficient to within one bit for a target descent tolerance on a normalized gradient class. Using second-order complexity (Kawamura & Cook, 2012), we account for reading finite gradient approximations and performing finite arithmetic, and certify polynomial bit time across rank changes (Section 3).

Spectral powers and direct descent. Spectral-power updates are developed through coupled Newton iterations (Qi et al., 2026), noise-aware exponent scheduling in DynMuon (Wu et al., 2026), and a related Schatten-power family in Freon (Shumaylov et al., 2026). We compare the gradient accuracy required by direct descent selection with a two-stage procedure that chooses a fractional transform with small descent bias and then reconstructs it accurately. Direct selection achieves the same descent tolerance with exponentially fewer input-accuracy bits at fixed matrix size (Section 5).

2 Exact minimax descent under gradient uncertainty

We first determine the best descent guarantee available from an uncertain gradient under a prescribed sensitivity constraint. This minimax characterization underlies our second-order complexity analysis: Section 3 translates a desired descent tolerance into an oracle precision requirement and accounts for the bit cost of computing a feasible update.

Notation and information model. Let $\mathbb{M} = \mathbb{R}^{d_1 \times d_2}$, $k = \min(d_1, d_2)$, and $d = d_1 + d_2$. In the training interpretation, $W \in \mathbb{M}$ is a matrix parameter block and $G = \nabla f(W)$ is the gradient of a differentiable objective, holding other blocks fixed. Write $\langle A, B\rangle = \operatorname{tr}(A^\top B)$, and use $\|\cdot\|_F$, $\|\cdot\|_{\mathrm{op}}$, and $\|\cdot\|_*$ for the Frobenius, spectral, and nuclear norms. Define

$$K = \{D \in \mathbb{M} : \|D\|_{\mathrm{op}} \le 1\}, \qquad \mathcal{E}(G, D) = \|G\|_* - \langle G, D\rangle. \tag{1}$$

For $G = \sum_i \sigma_i u_i v_i^\top$, its canonical polar factor is $P_0(G) = \sum_{\sigma_i > 0} u_i v_i^\top$, including $P_0(0) = 0$. Spectral/nuclear duality gives $\max_{D \in K} \langle G, D\rangle = \|G\|_*$, attained by $P_0(G)$; thus $\mathcal{E} \ge 0$ is the linearized descent shortfall per unit step size for $W^+ = W - \eta D$. The constraint ensures $\|(W^+ - W)x\|_2 \le \eta \|x\|_2$ for every layer input x, giving a common bound on the update's action.

Suppose the first-order Taylor remainder for every perturbation Z is at most $\beta\|Z\|_{\mathrm{op}}^2/2$, where β is the smoothness constant in (12). Then

$$f(W - \eta D) \le f(W) - \eta\|G\|_* + \eta\mathcal{E}(G, D) + \frac{\beta\eta^2}{2}.$$

The decision loss therefore enters directly as an additive cost in a bound on the training objective after the update.

The observation H satisfies $\|G - H\|_F \le \delta$, where δ is a certified error radius. Errors may vary arbitrarily within this bound across repeated queries. For stochastic gradients, a confidence event for this bound transfers directly to the descent guarantee; Appendix F.3 constructs such events from conditional variance bounds. A supplied rank bound $\operatorname{rank} G \le r \le k$ allows the guarantees to exploit low-rank structure; setting $r = k$ covers every matrix in $\mathbb{M}$. Lipschitz constants refer to the Frobenius norm on both input and output and measure the sensitivity of the update direction to its input. The minimax problem covers all matrix magnitudes, and Section 3 specifies the normalized domain for the input-precision result.

For $L > 0$, let $\mathcal{S}_L$ be the set of maps $S : \mathbb{M} \to K$ with $\mathrm{Lip}_F(S) \le L$. The optimal worst-case loss is

$$\mathfrak{M}_r(L, \delta) = \inf_{S \in \mathcal{S}_L} \sup_{\substack{G, H \in \mathbb{M},\ \operatorname{rank} G \le r \\ \|G - H\|_F \le \delta}} \mathcal{E}(G, S(H)). \tag{2}$$

The update rule uses L, δ, r and the observation H to choose its direction. The infimum covers the full class of Lipschitz maps into K, including both spectral and nonspectral rules.

A natural way to stabilize the optimal descent direction is to regularize its defining maximization using classical quadratic support-function smoothing (Nesterov, 2005):

$$C_\tau(H) = \arg\max_{D \in K} \left\{ \langle H, D \rangle - \frac{\tau}{2}\|D\|_F^2 \right\} = \Pi_K(H/\tau), \qquad \tau > 0, \tag{3}$$

where Π_K denotes Frobenius projection onto K. Projection is nonexpansive, so C_τ has Lipschitz constant at most $1/\tau$.

Theorem 1 (Exact minimax descent loss)**.** *For $1 \le r \le k$, $L > 0$, and $\delta \ge 0$,*

$$\mathfrak{M}_r(L, \delta) = \begin{cases} \dfrac{(\sqrt{r} + L\delta)^2}{4L}, & L\delta \le \sqrt{r}, \\ \sqrt{r}\,\delta, & L\delta \ge \sqrt{r}. \end{cases} \tag{4}$$

The infimum is attained by C_τ with $\tau = \max\{L^{-1}, \delta/\sqrt{r}\}$. Explicitly,

$$C_\tau(H) = U \operatorname{diag}(\min\{\sigma_i(H)/\tau, 1\})\, V^\top. \tag{5}$$

Proof idea. Projection optimality gives a quadratic loss bound whose optimization over $\tau \ge L^{-1}$ yields both regimes. Opposite rank-r gradients attain the matching lower bounds through nearby observations constrained by Lipschitz continuity or identical observations that hide their sign. Appendix B.2 gives the full proof. □

Interpretation of the two regimes. At zero uncertainty, the optimal loss is $r/(4L)$. With $\delta > 0$, increasing the allowed sensitivity improves the guarantee up to $L = \sqrt{r}/\delta$, where calibrated clipping achieves the optimal uncertainty contribution $\sqrt{r}\,\delta$. Reducing δ improves this guarantee further. This tradeoff identifies how to coordinate stability and gradient accuracy. Corollary 2 gives the unrestricted optimum and compares it with polar plug-in. The same minimax value holds for expected loss of randomized maps whose mean has Lipschitz constant at most L (Appendix B.3).

Corollary 2 (Optimal descent through uncertainty calibration)**.** *For $\delta > 0$, optimization over all update rules gives optimal loss $\sqrt{r}\,\delta$, attained by $C_{\delta/\sqrt{r}}$. For polar plug-in with $r = k$,*

$$\sup_{\|G - H\|_F \le \delta} \mathcal{E}(G, P_0(H)) = 2\sqrt{k}\,\delta. \tag{6}$$

Thus calibrated clipping halves the worst-case linearized loss of polar plug-in.

The proof of Corollary 2 is given in Appendix B.4.

Bounded gradients. The same minimax value also applies to bounded gradients under the following condition. For the class $\|G\|_{\mathrm{op}} \le B$, equality holds whenever

$$B \ge \begin{cases} (L^{-1} + \delta/\sqrt{r})/2, & L\delta \le \sqrt{r}, \\ \delta/\sqrt{r}, & L\delta \ge \sqrt{r}. \end{cases} \tag{7}$$

For every B, the global upper bound provides a valid descent certificate; condition (7) additionally certifies its sharpness. Appendix B.3 proves this extension.

Corollary 3 (Approximately low-rank gradients)**.** *Write $\mathrm{tail}_r(G) = \sum_{i>r} \sigma_i(G)$. For every matrix G, every $\tau > 0$, and every observation with $\|G - H\|_F \le \delta$,*

$$\mathcal{E}(G, C_\tau(H)) \le \min_{0 \le r \le k} \left\{ \mathrm{tail}_r(G) + \frac{(\tau\sqrt{r} + \delta)^2}{4\tau} \right\}. \tag{8}$$

The proof is given in Appendix B.5.

For any fixed threshold, the guarantee adapts to the gradient's spectral decay through the minimum over r. Full-rank neural-network gradients can therefore benefit from concentration of nuclear norm in a few singular directions: the tail term accounts exactly for the descent contribution of the remaining directions. Supplied rank or tail estimates make this dependence explicit.

A concrete uncertainty budget. At a fixed iterate, average N independent per-example gradients X_i with mean G and Frobenius variance $\mathbb{E}\|X_i - G\|_F^2 \le \sigma^2$. Suppose computing and storing their mean incurs total Frobenius error at most q, giving the observed matrix H. For confidence level $1-\nu$ with $\nu \in (0,1)$, a certificate based on the variance bound is

$$\delta_N = \frac{\sigma}{\sqrt{N\nu}} + q, \qquad \Pr(\|H - G\|_F \le \delta_N) \ge 1 - \nu. \tag{9}$$

For a known rank bound r and $\delta_N > 0$, clipping at $\tau = \delta_N/\sqrt{r}$ certifies $\mathcal{E}(G, C_\tau(H)) \le \sqrt{r}\,\delta_N$ on the same event, with sensitivity at most $\sqrt{r}/\delta_N$. Increasing the batch size reduces the sampling term; more accurate gradient computation or storage reduces q. The certificate therefore guides the balance between batch size and numerical accuracy: when the sampling term is small relative to q, reducing q improves the dominant component of the bound. For example, take $r = 16$, $\sigma = 0.02$, $N = 2000$, $\nu = 0.05$, and $q = 0.0005$. Then $\delta_N = 0.0025$ and $\tau = 0.000625$, certifying linearized descent loss at most 0.01 with probability at least 0.95. This calculation illustrates how a supplied variance bound turns sampling and numerical precision into an explicit descent certificate. Appendices F.3-F.4 give concrete variance bounds from per-example gradient norms, the proof and sample budget, and extensions to adaptive iterates and momentum.

3 Optimal input precision for descent

Theorem 1 determines the optimal update rule for a given gradient uncertainty and stability requirement. We now ask the converse design question: for a desired descent loss, how accurate must the gradient estimate be? We model access at a chosen accuracy by a request of precision p returning a dyadic matrix H_p (entries $m2^{-b}$, $m \in \mathbb{Z}$, $b \in \mathbb{Z}_{\ge 0}$) with $\|H_p - G\|_F \le 2^{-p}$. Such access is called a *matrix Cauchy oracle*. Increasing p requests a more accurate observation; every answer satisfying the error bound is permitted, including across repeated queries.

The parameter p measures *absolute Frobenius input accuracy* of the whole gradient matrix. Matrix dimension and scale translate hardware mantissa precision or entrywise rounding bounds into this accuracy (Appendix C). For minibatch gradients, a request at precision p corresponds to requiring a total error certificate at most 2^{-p}; the variance and numerical-error bounds above give one way to obtain it on a confidence event. The oracle expresses the information requirement, and a statistical sampling model translates it into a sample budget. Let $p^*_{\mathrm{desc}}(n, r)$ be the smallest uniform bound on the largest requested precision of a procedure that, for every $\|G\|_{\mathrm{op}} \le 1$, $\operatorname{rank} G \le r$, returns a finite rational $D \in K$ with $\mathcal{E}(G, D) \le 2^{-n}$.

Whereas Woodstock (2026) counts real arithmetic operations, our computational cost includes writing requests, reading complete answers, rational arithmetic with guard bits, and writing the finite output. Writing $\ell(p)$ for the maximum oracle-answer length up to precision p, our runtime is polynomial in $d, n, \ell(p)$; dependence on this length function gives the second-order bound. Appendix C specifies the encoding conventions and the formal second-order polynomial model (Kawamura & Cook, 2012).

Theorem 4 (Optimal descent precision)**.** *For $n \ge 1$ and $1 \le r \le k$,*

$$\lceil n + \tfrac{1}{2}\log_2 r \rceil \le p^*_{\mathrm{desc}}(n, r) \le \lceil n + \tfrac{1}{2}\log_2 r \rceil + 1. \tag{10}$$

The upper bound is attained deterministically with one matrix-oracle query and second-order polynomial bit time, uniformly across singular-value gaps and rank changes. The lower bound certifies precision optimality among procedures with adaptive repeated queries and randomized success probability at least $2/3$ on every input.

Proof idea. Opposite rank-r gradients with nuclear norm $\sqrt{r}\,2^{-p}$ admit identical zero answers, forcing the precision lower bound. One additional bit permits certified clipping and rescaling within the descent-error budget while ensuring spectral feasibility. Appendix C.2 gives the full argument, including randomized queries. □

For fixed r, each additional bit of target descent accuracy needs one additional bit of input accuracy, up to the fixed additive allowance. The next lemma supplies the iteration count and working precision needed to compute the update, including guard precision near small thresholds.

Lemma 5 (Certified multiplication-only clipping). *For rational H with $\|H\|_{\mathrm{op}} \le 2$ and rational $0 < \tau, \zeta \le 1$, one can compute a rational $\widetilde{C}$ satisfying $\|\widetilde{C} - C_\tau(H)\|_F \le \zeta$ using $O(j)$ dense matrix products, where*

$$j = O\big(\log d + \log(1/\tau) + \log(1/\zeta)\big). \tag{11}$$

Working precision for the rounded iterates is of the same order. Total bit time is polynomial in d, these parameters, and the complete input encoding length.

An absolute-value formulation of clipping also appears in Yi (2026), using a polar factor and a clipped positive-semidefinite factor. Our construction applies absolute values directly to a symmetric dilation, giving explicit approximation and rounding budgets valid across rank changes (Appendix D). For outputs each within Frobenius error ζ of ideal clipping, an additive 2ζ allowance extends the ideal Lipschitz bound to a finite-output stability certificate.

4 Optimization guarantees with sharp uncertainty dependence

Achievable stationarity. Under smoothness, the average descent shortfall controls stationarity along the optimization trajectory. Suppose $f : \mathbb{M} \to \mathbb{R}$ is differentiable, bounded below, and satisfies

$$f(W + Z) \le f(W) + \langle \nabla f(W), Z \rangle + \frac{\beta}{2}\|Z\|_{\mathrm{op}}^2. \tag{12}$$

Ordinary L_f-smoothness in Frobenius norm implies (12) with $\beta = kL_f$. Write $\Delta = f(W_0) - \inf f$. The stationarity measure $\|\nabla f(W)\|_*$ is the largest linearized decrease per unit spectral-norm step; it vanishes exactly at a stationary point.

Proposition 6 (Descent with a certified directional error). *For $\eta > 0$ and an integer $T \ge 1$, if $W_{t+1} = W_t - \eta D_t$, $D_t \in K$, and $\mathcal{E}(\nabla f(W_t), D_t) \le \rho_t$ for $0 \le t < T$, then*

$$\frac{1}{T}\sum_{t<T}\|\nabla f(W_t)\|_* \le \frac{\Delta}{\eta T} + \frac{\beta\eta}{2} + \frac{1}{T}\sum_{t<T}\rho_t. \tag{13}$$

In particular, if $\|H_t - \nabla f(W_t)\|_F \le \delta$ with $\delta > 0$, clipping at $\tau = \delta/\sqrt{k}$ and an additional implementation loss ω_t gives, with $\eta = \sqrt{2\Delta/(\beta T)}$ when $\Delta > 0$,

$$\frac{1}{T}\sum_{t<T}\|\nabla f(W_t)\|_* \le \sqrt{\frac{2\beta\Delta}{T}} + \sqrt{k}\,\delta + \frac{1}{T}\sum_{t<T}\omega_t. \tag{14}$$

For an independent uniform index $J \in \{0, \dots, T-1\}$, the left side equals $\mathbb{E}_J\|\nabla f(W_J)\|_*$ conditional on the trajectory; it also bounds the minimum over the iterates. Here $\omega_t \ge 0$ bounds $\mathcal{E}(\nabla f(W_t), D_t) - \mathcal{E}(\nabla f(W_t), C_\tau(H_t))$; Appendix F gives an explicit certificate.

The proof is given in Appendix F.1. For a target $\min_{t<T}\|\nabla f(W_t)\|_* \le \gamma$ with $\gamma > 0$ and $\Delta > 0$, it suffices to take $T \ge 18\beta\Delta/\gamma^2$, $\delta \le \gamma/(3\sqrt{k})$, and average implementation loss at most $\gamma/3$. For sampled gradients, Appendix F.3 constructs a simultaneous confidence event across adaptive iterates on which the bound applies.

The lower bounds below address algorithm-independent uncertainty limits for returned points and best iterates.

Optimality for a returned point. The following lower bound certifies that calibrated clipping achieves the optimal uncertainty dependence for the point an algorithm returns. It applies to arbitrary errors within a fixed certified radius, while (9) shows how independent sampling improves the guarantee by reducing that radius.

Theorem 7 (Optimal uncertainty dependence). *With access only to a gradient oracle of certified error $\delta > 0$, no algorithm can guarantee $\|\nabla f(\widehat{W})\|_* < \sqrt{k}\delta$ on every condition-number-one strongly convex quadratic. This remains true with any finite number of adaptive queries, an exact Hessian oracle, and success probability at least $2/3$ on every input. Moreover, for every randomized algorithm there is such a quadratic with $\mathbb{E}\|\nabla f(\widehat{W})\|_* \ge \sqrt{k}\delta$.*

Proof idea. Two quadratics with opposite minimizer shifts admit identical gradient responses and Hessians. The triangle inequality makes their success sets below $\sqrt{k}\,\delta$ disjoint and forces the expectation lower bound. Appendix F.5 gives the full proof. □

Sharpness for the best iterate. An upper bound on average stationarity also controls the best iterate. To compare this criterion directly, write $B_T(f) = \min_{0\le t<T}\|\nabla f(W_t)\|_*$. This evaluates a trajectory by the true stationarity of its best point. A family of 2^k indistinguishable quadratics certifies sharp uncertainty dependence for this criterion as well.

Theorem 8 (Uncertainty dependence for the best iterate)**.** *Fix an integer $T \ge 1$, $\delta > 0$, $\lambda > 0$, and an initial point W_0. Consider any randomized algorithm producing $W_0, \dots, W_{T-1}$ with access only to a gradient oracle of Frobenius error at most δ, optionally together with an exact Hessian oracle. Any finite number of adaptive queries is allowed. There is a condition-number-one strongly convex quadratic f with Hessian λI and a valid gradient oracle such that*

$$\mathbb{E}B_T(f) \ge \sqrt{k}\,\delta\left(1 - \sqrt{\frac{2\ln T}{k}}\right)_+, \qquad (x)_+ = \max\{x, 0\}. \tag{15}$$

Expectation and probability are over the algorithm's randomness, with the objective fixed independently of it. For every $0 < \epsilon < 1$, there is also such an instance, possibly different, satisfying

$$\Pr\Big(B_T(f) \le (1-\epsilon)\sqrt{k}\,\delta\Big) \le \min\left\{1, T\exp\left(-\frac{k\epsilon^2}{2}\right)\right\}. \tag{16}$$

Both instances have initial gap $\Delta = \delta^2/(2\lambda)$ and satisfy (12) with $\beta = k\lambda$.

Proof idea. Quadratics with hidden diagonal signs share gradient responses and Hessians, making the candidate list independent of those signs. Nuclear-norm duality reduces the argument to a maximum of T sign sums, with expectation at most $\sqrt{2k\ln T}$ and concentration yielding the probability bound. Appendix F.6 proves the theorem and sharper bounds. □

For $\ln T = o(k)$, Theorem 8 recovers the leading uncertainty term $\sqrt{k}\delta$ even for the best iterate. Moreover, guaranteeing $B_T(f) \le (1-\epsilon)\sqrt{k}\delta$ with probability at least $2/3$ on every instance requires

$$T \ge \frac{2}{3}\exp\left(\frac{k\epsilon^2}{2}\right). \tag{17}$$

This quantifies the strength of the uncertainty calibration: its leading term is optimal when $\ln T = o(k)$, and any uniform improvement by a fixed fraction requires exponentially many candidates in k. For methods that query gradients only at their iterates and remain at the initial point when the observation history is zero, including calibrated clipping, Appendix F.7 proves the exact best-iterate level $\sqrt{k}\delta$ for every horizon.

Total computational cost. To translate stationarity guarantees into a total computational budget, Appendix G analyzes complete finite-bit trajectories, including the identification of an approximately stationary iterate. For its normalized objective class at fixed matrix dimensions, the second-order oracle analysis establishes an optimal deterministic bit runtime of $\Theta(n4^n)$ for stationarity tolerance 2^{-n}, using gradient-only access, linear-length oracle names, and full-message accounting (Theorem 11). It also provides a certified clipping trajectory with explicit arithmetic and iterate-rounding budgets.

5 The Precision Advantage of Direct Descent

Section 3 establishes the input precision sufficient for direct descent. We now examine the precision needed to reconstruct an intermediate spectral transform. On the normalized class $\|G\|_{\mathrm{op}} \le 1$, we compare direct descent selection with a two-stage procedure that first chooses a fractional spectral transform with small descent bias and then reconstructs it accurately. Bringing the transform closer to polar reduces its descent bias and increases its reconstruction precision requirement. Direct descent weights each singular component by its contribution to optimization progress, achieving the same descent tolerance with less input precision.

Consider the fractional family $\Phi_\alpha(G) = U\,\mathrm{diag}(\sigma_i^\alpha)V^\top$, $0 < \alpha \le 1$, with zero singular values mapped to zero. At $\alpha = 1$, this map returns G itself. As $\alpha \downarrow 0$, each positive singular value σ_i^α tends to one, so $\Phi_\alpha(G)$ approaches the canonical polar factor $P_0(G)$ for every fixed G.

For $0 < \delta \le 1$, the optimal accuracy for reconstructing this transform is described by

$$e_{\alpha,k}(\delta) = \inf_{\mathcal{R}} \sup_{\substack{\|G\|_{\mathrm{op}} \le 1 \\ \|G-H\|_F \le \delta}} \|\Phi_\alpha(G) - \mathcal{R}(H)\|_F. \tag{18}$$

The infimum ranges over the full class of maps $\mathcal{R} : \mathbb{M} \to \mathbb{M}$, including arbitrary discontinuous and noncomputable maps.

Theorem 9 (Information precision for spectral reconstruction)**.** *For every $0 < \alpha \le 1$,*

$$k^{(1-\alpha)/2}\delta^\alpha \le e_{\alpha,k}(\delta) \le (8k)^{(1-\alpha)/2}\delta^\alpha. \tag{19}$$

Consequently, accuracy 2^{-n} requires $\alpha p \ge n + \frac{1-\alpha}{2}\log_2 k$ when $\delta = 2^{-p}$; replacing k by $8k$ in this inequality suffices for the ideal plug-in reconstruction.

Appendix E.1 gives a full proof.

The information-radius result covers every real exponent $0 < \alpha \le 1$; rational exponents additionally provide finite parameter descriptions for the computational analysis below.

Precision savings from targeting descent. The exact noiseless descent bias of this family is

$$\sup_{\|G\|_{\mathrm{op}} \le 1} \mathcal{E}(G, \Phi_\alpha(G)) = \frac{k\alpha}{(1+\alpha)^{1+1/\alpha}} \in [k\alpha/4, k\alpha/e]. \tag{20}$$

Consider a two-stage procedure that chooses α to make this bias at most $\varepsilon/2$, then returns a feasible reconstruction $D \in K$ with Frobenius error $\varepsilon/(2\sqrt{k})$. The bias lower bound $k\alpha/4$ forces $\alpha \le 2\varepsilon/k$. Meanwhile, Theorem 9 requires input precision at least $\log_2(2\sqrt{k}/\varepsilon)/\alpha$ for this reconstruction tolerance. Combining the small exponent required for descent with this inverse dependence on α gives

$$p = \Omega\left(\frac{k}{\varepsilon}\log\frac{k}{\varepsilon}\right). \tag{21}$$

Appendix E.3 proves both the exact bias and this precision bound, and supplies a finite rational realization of the two-stage procedure. Direct descent achieves the same descent tolerance with input precision $O(\log(\sqrt{k}/\varepsilon))$. With $\varepsilon = 2^{-n}$ and fixed k, its $O(n)$ input-precision budget gives an exponential saving over the $\Omega(n2^n)$ requirement of this two-stage fractional reconstruction procedure.

Efficient computation across rank changes. Theorem 4 uses one algorithm to select a feasible descent direction in second-order polynomial bit time across rank changes. For fractional reconstruction, the exponent also controls the information cost. If $\alpha_b = 2^{-b}$, even the fixed reconstruction tolerance $1/8$ requires $\Omega(2^b)$ requested precision. This cost is exponential in the $O(b)$-bit description of the exponent and persists when oracle answers have lengths polynomial in the requested precision. A precision-p request has unary length p, so these precision lower bounds also lower-bound bit runtime in the second-order oracle model.

Approximate descent retains its uniform second-order polynomial bit-time guarantee at the polar endpoint $\alpha = 0$. The scalar polar map is zero at zero and one at every positive input; its discontinuity yields a strict computability separation between global polar reconstruction and direct descent. Appendix E.2 proves this separation and the exponential precision bound above. By weighting components near zero according to their contribution to descent, the direct method produces a certified update across rank changes.

6 Conclusion

Our results show how to compute reliable updates from approximate gradients. In spectral optimization, calibrated clipping attains the optimal worst-case tradeoff between gradient uncertainty, update stability, and descent. Targeting descent directly supports certified computation in second-order polynomial bit time across rank changes. At fixed matrix size, it also requires exponentially less input precision than the two-stage fractional reconstruction procedure. Our complexity and stationarity bounds translate these guarantees into explicit budgets for gradient and arithmetic accuracy. Once implementation error is small relative to the gradient-uncertainty contribution, improving gradient information becomes the main way to achieve stronger stationarity guarantees.

AI USE STATEMENT

AI tools were used to aid and polish the writing of the paper.

ETHICS STATEMENT

This work consists of theoretical analysis and proofs; there is no particular ethical concern involved.

REPRODUCIBILITY STATEMENT

All mathematical claims are supported by proofs in the appendices; Appendix A maps the main results to their proof locations. Appendix D specifies a finite rational implementation, including iteration counts and rounding budgets.

REFERENCES

Noah Amsel, David Persson, Christopher Musco, and Robert M. Gower. The polar express: Optimal matrix sign methods and their application to the Muon algorithm. In *International Conference on Learning Representations*, 2026a.

Noah Amsel, Jack Zhang, Kwangjun Ahn, Ali Naeimi, Austin Feng, Berlin Chen, Tri Dao, and John Langford. Dion3: Full-stack orthogonal updates. arXiv:2608.11612, 2026b.

Fredrik Andersson, Marcus Carlsson, and Karl-Mikael Perfekt. Operator-Lipschitz estimates for the singular value functional calculus. *Proceedings of the American Mathematical Society*, 144(5): 1867-1875, 2016. doi: 10.1090/proc/12843.

Stefania Bellavia, Gianmarco Gurioli, Benedetta Morini, and Philippe L. Toint. Adaptive regularization algorithms with inexact evaluations for nonconvex optimization. *SIAM Journal on Optimization*, 29(4):2881-2915, 2019. doi: 10.1137/18M1226282.

Jeremy Bernstein and Laker Newhouse. Old optimizer, new norm: An anthology. Workshop on Optimization for Machine Learning (OPT), 2024.

Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223-311, 2018. doi: 10.1137/16M1080173.

Sinho Chewi, Sébastien Bubeck, and Adil Salim. On the complexity of finding stationary points of smooth functions in one dimension. In *International Conference on Algorithmic Learning Theory*, volume 201 of *PMLR*, pp. 358-374, 2023.

Olivier Devolder, François Glineur, and Yurii Nesterov. First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming*, 146(1-2):37-75, 2014. doi: 10.1007/s10107-013-0677-5.

Thang Do, Steffen Dereich, and Arnulf Jentzen. On MUON optimization: From non-convergence to an error analysis with Polar Express and the Newton-Schulz polynomial from implementations. arXiv:2608.04607, 2026.

Chongyu Fan, Gaowen Liu, Mingyi Hong, Ramana Rao Kompella, and Sijia Liu. Rethinking Muon beyond pretraining: Spectral failures and high-pass remedies for VLA and RLVR. arXiv:2605.19282, 2026.

Evan S. Gawlik and Melvin Leok. Iterative computation of the Fréchet derivative of the polar decomposition. *SIAM Journal on Matrix Analysis and Applications*, 38(4):1354-1379, 2017. doi: 10.1137/16M108971X.

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

Hua Huang. How much orthogonalization does Muon need? arXiv:2606.00371, 2026.

Martin Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *International Conference on Machine Learning*, volume 28(1) of *PMLR*, pp. 427-435, 2013.

Xiaowen Jiang, Andrei Semenov, and Sebastian U. Stich. Enhancing LLM training via spectral clipping. In *International Conference on Machine Learning*, 2026.

Akitoshi Kawamura and Stephen Cook. Complexity theory for operators in analysis. *ACM Transactions on Computation Theory*, 4(2):5, 2012. doi: 10.1145/2189778.2189780.

Gyu Yeol Kim and Min-hwan Oh. Convergence of Muon with Newton-Schulz. In *International Conference on Learning Representations*, 2026.

Dmitry Kovalev. Understanding gradient orthogonalization for deep learning via non-Euclidean trust-region optimization. arXiv:2503.12645, 2025.

Jiaxiang Li and Mingyi Hong. A note on the convergence of Muon. arXiv:2502.02900, 2025.

Mingyi Li and Taira Tsuchiya. Muon with finite Newton-Schulz: The smoothing benefit in nonsmooth nonconvex optimization. arXiv:2608.26288, 2026.

Xianliang Li, Zihan Zhang, Weiyang Liu, and Han Bao. Denoise first, orthogonalize later: Understanding momentum in Muon via spectral filtering. arXiv:2606.03899, 2026.

Jingyuan Liu, Jianlin Su, Xingcheng Yao, Zhejun Jiang, Guokun Lai, Yulun Du, Yidao Qin, Weixin Xu, Enzhe Lu, Junjie Yan, Yanru Chen, Huabin Zheng, Yibo Liu, Shaowei Liu, Bohong Yin, Weiran He, Han Zhu, Yuzhi Wang, Jianzhou Wang, Mengnan Dong, Zheng Zhang, Yongsheng Kang, Hao Zhang, Xinran Xu, Yutao Zhang, Yuxin Wu, Xinyu Zhou, and Zhilin Yang. Muon is scalable for LLM training. arXiv:2502.16982, 2025.

Zhuanghua Liu, Menglian Wang, and Luo Luo. Musec: MomentUm SpEctral Clipping for stable Muon-type training. arXiv:2609.11655, 2026.

Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. Mixed precision training. In *International Conference on Learning Representations*, 2018.

Dominique Monnet and Dominique Orban. A multi-precision quadratic regularization method for unconstrained optimization with rounding error analysis. *Computational Optimization and Applications*, 91:997-1031, 2025. doi: 10.1007/s10589-025-00676-x.

Yuji Nakatsukasa, Zhaojun Bai, and François Gygi. Optimizing Halley's iteration for computing the matrix polar decomposition. *SIAM Journal on Matrix Analysis and Applications*, 31(5):2700-2720, 2010. doi: 10.1137/090774999.

Yurii Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103: 127-152, 2005. doi: 10.1007/s10107-004-0552-5.

Tien-Phat Nguyen, Truong Nguyen, Minh-Phuc Truong, Tuc Nguyen, James Bailey, and Trung Le. Spectral flattening is all Muon needs: How orthogonalization controls learning rate and convergence. arXiv:2605.13079, 2026.

Thomas Pethick, Wanyun Xie, Kimon Antonakopoulos, Zhenyu Zhu, Antonio Silveti-Falls, and Volkan Cevher. Training deep learning models with norm-constrained LMOs. In *International Conference on Machine Learning*, volume 267 of *PMLR*, pp. 49069-49104, 2025.

Xianbiao Qi, Marco Chen, Jiaquan Ye, Yelin He, and Rong Xiao. Delving into Muon and beyond: Deep analysis and extensions. arXiv:2602.04669, 2026.

Maria-Eleni Sfyraki and Jun-Kun Wang. Lions and Muons: Optimization via stochastic Frank-Wolfe under heavy-tailed noise. In *International Conference on Machine Learning*, 2026.

Wei Shen, Ruichuan Huang, Minhui Huang, Cong Shen, and Jiawei Zhang. On the convergence analysis of Muon. arXiv:2505.23737, 2026.

Egor Shulgin, Sultan AlRashed, Peter Richtárik, and Francesco Orabona. Beyond the ideal: Analyzing the inexact Muon update. In *International Conference on Artificial Intelligence and Statistics*, volume 300 of *PMLR*, pp. 5167-5175, 2026.

Zakhar Shumaylov, Nathaël Da Costa, Peter Zaika, Bálint Mucsányi, Alex Massucco, Yoav Gelberg, Carola-Bibiane Schönlieb, Yarin Gal, and Philipp Hennig. Muon is not that special: Random or inverted spectra work just as well. arXiv:2605.11181, 2026.

Xuan Tang, Jichu Li, and Difan Zou. A convergence analysis of adaptive optimizers under floating-point quantization. In *International Conference on Learning Representations*, 2026.

Zev Woodstock. High-precision linear minimization is no slower than projection. *Optimization Letters*, 20:909-915, 2026. doi: 10.1007/s11590-025-02269-3.

Fangzhou Wu, Rikhav Shah, Sandeep Silwal, and Qiuyi Zhang. DynMuon: A dynamic spectral shaping view of Muon. arXiv:2605.17109, 2026.

Albert Yi. MuCon: Clipped Muon updates for LLM training. arXiv:2605.26459, 2026.

A Overview

The appendices provide the geometric and computational foundations of the main results, together with their full proofs, and extend the analysis to complete optimization trajectories.

Appendix B develops the projection geometry underlying spectral clipping and the minimax results of Section 2. Appendix B.2 proves the exact minimax formula (Theorem 1), and Appendix B.3 establishes its randomized and bounded-gradient extensions. Appendices B.4 and B.5 prove the optimal uncertainty calibration and polar comparison (Corollary 2) and the approximate-rank guarantee (Corollary 3), respectively.

Appendix C formalizes the matrix Cauchy oracle, finite encodings, and second-order bit-time model used in Section 3. Appendix C.2 proves the necessary and sufficient input-precision bounds of Theorem 4, including the lower bound for adaptive randomized queries. The appendix also distinguishes requested input accuracy from arithmetic precision.

Appendix D proves Lemma 5 by constructing certified multiplication-only clipping. It gives the symmetric-dilation formula, a Newton-Schulz error estimate valid at zero, and explicit iteration and rounding budgets. The complete precision-to-descent procedure then realizes the second-order polynomial-time upper bound in Theorem 4.

Appendix E supplies the reconstruction analysis and precision comparison of Section 5. Appendix E.1 proves Theorem 9; Appendix E.2 derives the input-bit interpretation, the exponential precision lower bound for small fractional exponents, and the computability separation at the polar endpoint. Appendix E.3 proves the exact fractional descent bias and the precision saving of direct descent over the two-stage reconstruction procedure, including a finite rational realization of that procedure.

Appendix F develops the optimization certificates and uncertainty lower bounds of Section 4. Appendix F.1 proves Proposition 6. The subsequent analysis certifies implementation error, while Appendices F.3 and F.4 derive minibatch and momentum uncertainty budgets, including the confidence guarantee in (9) and its extension to adaptive iterates. Appendices F.5 and F.6 prove Theorems 7 and 8, respectively, for a returned point and the best iterate. Appendix F.7 establishes exact best-iterate sharpness for zero-preserving updates.

Appendix G extends the local guarantees to complete deterministic trajectories in the second-order oracle model, accounting for finite iterate representations and total bit runtime. After specifying the model in Appendix G.1, Appendix G.2 proves Theorem 11, the optimal-runtime result summarized in Section 4. Appendix G.3 analyzes dependence on the oracle size function, and Appendix G.4 constructs a complete certified clipping trajectory with explicit arithmetic and iterate-rounding budgets.

B PROJECTION GEOMETRY AND PROOFS OF THE MINIMAX RESULTS

B.1 DUALITY AND THE CLIPPING FORMULA

For any $D \in K$, an SVD of G yields

$$\langle G, D\rangle = \sum_i \sigma_i u_i^\top D v_i \le \sum_i \sigma_i = \|G\|_*.$$

The canonical polar $P_0(G)$ belongs to K and attains equality, including at rank-deficient matrices. Its Frobenius norm is the square root of the number of positive singular values. The definition assigns zero to the null-space component, fixing a canonical choice at every rank.

The variational characterization (3) is the usual quadratically smoothed support-function construction. To verify the SVD formula directly, extend U, V to orthogonal matrices. Frobenius invariance reduces the projection to that of a rectangular diagonal matrix with entries σ_i/τ. A diagonal candidate with entries $d_i = \min\{\sigma_i/\tau, 1\}$ belongs to K. For every $Z \in K$, each corresponding diagonal entry satisfies $Z_{ii} \le 1$, and the residual $\sigma_i/\tau - d_i$ is zero unless $d_i = 1$. Consequently,

$$\langle H/\tau - C_\tau(H), Z - C_\tau(H)\rangle \le 0.$$

This variational inequality identifies the unique Frobenius projection and proves (5) directly from projection geometry.

For completeness, if $D = \Pi_K(X)$ and $E = \Pi_K(Y)$, the two projection inequalities imply

$$\|D - E\|_F^2 \le \langle X - Y, D - E\rangle \le \|X - Y\|_F \|D - E\|_F.$$

Thus projection is nonexpansive. The constant $1/\tau$ for C_τ is exact: on sufficiently small multiples of any fixed nonzero matrix, all singular values lie below τ and the map is $H \mapsto H/\tau$.

Comparison with smoothing bounds. For their admissible scalar spectral filters h, Li & Tsuchiya (2026) prove $\int_0^\infty (1 - h(x))\, dx \ge [2\,\mathrm{Lip}(h)]^{-1}$ (their Appendix F). This lower bound concerns the integrated error in a smoothed spectral potential. Theorem 1 instead concerns pointwise linearized descent loss under Frobenius observation uncertainty, with a matching lower bound over all admissible stable maps, including nonspectral maps.

B.2 PROOF OF THEOREM 1

Proof. Let $D = C_\tau(H)$ and $V = P_0(G)$. Projection optimality implies $\langle H - \tau D, V - D\rangle \le 0$. Also $\|V\|_F \le \sqrt{r}$. With $u = \|V - D\|_F$,

$$\begin{aligned}\mathcal{E}(G, D) = \langle G, V - D\rangle &\le \tau\langle D, V - D\rangle + \delta u \\ &\le (\tau\sqrt{r} + \delta)u - \tau u^2 \le \frac{(\tau\sqrt{r} + \delta)^2}{4\tau}.\end{aligned} \tag{22}$$

Here $\langle D, V - D\rangle = \langle V, V - D\rangle - u^2 \le \sqrt{r}\, u - u^2$. Projection is nonexpansive, so $\mathrm{Lip}_F(C_\tau) \le \tau^{-1}$. Choosing the stated τ gives both upper bounds.

For the converse, choose a rank-r partial isometry Q with all its nonzero singular values equal to one, and set $u_0 = Q/\sqrt{r}$. Thus $\|u_0\|_F = 1$ and $\|u_0\|_* = \sqrt{r}$. For any $h \ge 0$, take

$$H_\pm = \pm h u_0, \qquad G_\pm = \pm(h + \delta)u_0. \tag{23}$$

Both pairs are admissible. For every $S \in \mathcal{S}_L$, their average loss satisfies

$$\begin{aligned}\frac{\mathcal{E}(G_+, S(H_+)) + \mathcal{E}(G_-, S(H_-))}{2} &= (h + \delta)\sqrt{r} - \frac{h + \delta}{2}\langle u_0, S(H_+) - S(H_-)\rangle \\ &\ge (h + \delta)(\sqrt{r} - Lh).\end{aligned} \tag{24}$$

When $L\delta \le \sqrt{r}$, set $h = (\sqrt{r} - L\delta)/(2L)$. Otherwise set $h = 0$. The right side equals (4). At least one loss is at least its average, completing the matching lower bound. □

B.3 Randomization and bounded-gradient sharpness

The two witnesses (23) prove the minimax converse for every map in $\mathcal{S}_L$. In the regime $L\delta \geq \sqrt{r}$ they have the same observation $H = 0$; they identify the exact contribution of observation uncertainty. For any randomized rule returning values in K, the mean also lies in K, and loss is affine in the output. Hence randomized expected loss equals the loss of the mean rule. This proves both the unrestricted uncertainty lower bound and the lower bound in Theorem 1 whenever that mean is L-Lipschitz. The condition on the mean permits random outputs with arbitrary pathwise regularity.

When $L\delta \leq \sqrt{r}$, the witness norm is

$$\|G_{\pm}\|_{\mathrm{op}} = (h+\delta)/\sqrt{r} = \tfrac{1}{2}(L^{-1} + \delta/\sqrt{r}).$$

When $L\delta \geq \sqrt{r}$, it is $\delta/\sqrt{r}$. Thus under (7) these witnesses belong to the bounded-gradient class. The global upper bound applies on that subclass, and the same witnesses still attain its lower bound. This establishes the bounded-domain equality under the explicit witness-norm condition.

B.4 Proof of Corollary 2

Proof. Removing the Lipschitz restriction gives, for $\delta > 0$, the lower bound $\sqrt{r}\delta$ using $G_{\pm} = \pm\delta Q/\sqrt{r}$ at observation zero. The choice $\tau = \delta/\sqrt{r}$ in (22) attains it. At $\delta = 0$ the unrestricted minimum is zero and is attained by canonical polar; positive uncertainty is handled by the calibrated clipping threshold above.

For polar plug-in, write $V = P_0(G)$ and $P = P_0(H)$. Since $\langle H, V - P\rangle \leq 0$,

$$\mathcal{E}(G,P) = \langle G, V-P\rangle \leq \langle G-H, V-P\rangle \leq \delta(\|V\|_F + \|P\|_F) \leq 2\sqrt{k}\delta.$$

For $0 < t < \delta/\sqrt{k}$, choose $H = -tQ$ and $G = (\delta/\sqrt{k} - t)Q$ with a rank-k partial isometry Q. The loss is $2k(\delta/\sqrt{k} - t)$, converging to the upper bound. This family of observations approaches the upper bound and establishes the exact supremum in (6).

□

B.5 Proof of Corollary 3

Proof. For a fixed G, choose $V_r = \sum_{i \leq r, \sigma_i > 0} u_i v_i^\top$, with $V_0 = 0$. It satisfies $V_r \in K$, $\|V_r\|_F \leq \sqrt{r}$, and

$$\|G\|_* - \langle G, V_r\rangle = \mathrm{tail}_r(G).$$

With $D = C_\tau(H)$ and $u = \|V_r - D\|_F$, projection optimality gives

$$\begin{aligned}\mathcal{E}(G,D) &= \mathrm{tail}_r(G) + \langle G, V_r - D\rangle \\ &\leq \mathrm{tail}_r(G) + (\tau\sqrt{r} + \delta)u - \tau u^2 \\ &\leq \mathrm{tail}_r(G) + \frac{(\tau\sqrt{r}+\delta)^2}{4\tau}.\end{aligned}$$

This holds for every r, including zero, so taking the minimum proves Corollary 3. The guarantee retains the original observation radius $\|G - H\|_F \leq \delta$ and accounts for the remaining singular directions through their nuclear tail, the appropriate loss for linearized spectral descent.

□

C Cauchy names, information precision, and second-order bit complexity

C.1 Encoding convention

The cost of accessing an approximation depends on both its requested accuracy and the length of the returned encoding.

Finite encodings. The positive integers d_1, d_2, r are explicit finite inputs, and n is given in unary, as a string of n ones. Arithmetic work is polynomial in the explicit matrix dimension. Dyadic matrix entries are encoded by signed binary integers and explicit binary-fraction fields; general rational outputs are encoded by binary numerators and positive denominators. Parsing and all output writing are charged. Changing among standard polynomially equivalent encodings changes the time bound polynomially, without changing the accuracy exponent in Theorem 4.

Approximation guarantee. A valid matrix name φ answers the distinguished unary query 1^p (a string of p ones, costing p steps to write) with a finite encoding of H_p such that $\|G - H_p\|_F \le 2^{-p}$. The model permits any finite answer length and any sequence of individually valid approximations. Nondistinguished queries can return a fixed dummy payload.

Response lengths. For a second-order normal form, pad all responses of the same query length to the same length, with this length nondecreasing. The resulting length function is $\ell(p) = \max_{|u|\le p} |\varphi(u)|$, where $|\cdot|$ denotes string length. Padding does not alter the matrix payload. In upper bounds the time to read the selected response is charged; in lower bounds the same padding is chosen for the two indistinguishable inputs.

Second-order polynomial bit-time bound. A second-order polynomial is an expression generated from nonnegative integer constants, ordinary size variables, addition, multiplication, and applications $\ell(P)$. The realization in Theorem 4 makes one query of length

$$p = n + \lceil \tfrac{1}{2} \log_2 r \rceil + 1,$$

reads its answer of length at most $\ell(p)$, and runs rational arithmetic polynomial in $d, n, \ell(p)$. This is a second-order polynomial bound. The integer $\lceil \frac{1}{2} \log_2 r \rceil$ can be found without real logarithms: it is the smallest integer a with $r \le 4^a$.

C.2 Proof of Theorem 4

Proof. Lower bound. Fix an algorithm whose largest requested precision is at most p. Choose Q with ones in its first r diagonal positions and zeros elsewhere, so Q is a rank-r partial isometry, and let $G_\pm = \pm 2^{-p} Q/\sqrt{r}$. These are computable real matrices; their spectral norms are at most one. At every precision $q \le p$, the zero matrix is a valid response to both inputs. Supply this common response, along with identical padding and dummy responses. Any deterministic adaptive computation has the same entire transcript on the two inputs. After fixing its random tape, the same holds for a randomized computation.

These common finite answers can be extended to total valid names. At precision $q > p$, approximate the entries of $G_\pm$ by dyadics with component error at most $2^{-q}/\sqrt{d_1 d_2}$. Use the same response-length bound for both signs. The entries involve only the computable number $1/\sqrt{r}$, so their approximation lengths are polynomial in q, p, d, and efficient names can be chosen. The lower bound thus holds for computable matrices with efficiently generated names of polynomial length.

For every feasible output D,

$$\mathcal{E}(G_+, D) + \mathcal{E}(G_-, D) = 2\sqrt{r}\, 2^{-p}.$$

If $\sqrt{r}\, 2^{-p} > 2^{-n}$, the two sets of successful outputs are disjoint. Under the common output distribution their success probabilities sum to at most one. Hence at least one is at most $1/2$, ruling out a uniform $2/3$ success guarantee. The conclusion holds for every finite number of requests of precision at most p.

Upper bound. Write $\varepsilon = 2^{-n}$ and query at the stated upper precision. Then $\|H - G\|_F \le \varepsilon/(2\sqrt{r})$. Choose

$$\tau = \frac{\varepsilon}{8r}, \qquad \zeta = \frac{\varepsilon}{16k}, \qquad D = \frac{\widetilde{C}}{1+\zeta}, \quad \|\widetilde{C} - C_\tau(H)\|_F \le \zeta. \tag{25}$$

Equation (22) bounds the ideal clipping loss by $25\varepsilon/32$. The scaling ensures $\|D\|_{\mathrm{op}} \le 1$, and

$$\|D - C_\tau(H)\|_F \le \zeta(1 + \sqrt{k}), \quad |\langle G, D - C_\tau(H)\rangle| \le 2k\zeta = \varepsilon/8. \tag{26}$$

Thus the final loss is at most $25\varepsilon/32 + \varepsilon/8 = 29\varepsilon/32 < \varepsilon$. Lemma 5, proved in Appendix D, supplies the certified approximation in (25) in polynomial bit time, with input encoding length included. □

C.3 Source precision versus arithmetic precision

The optimality statement identifies the largest requested source precision. The arithmetic budget additionally accounts for $\log(1/\tau)$, $\log(1/\zeta)$, matrix dimensions, and the input encoding length. For the parameters in (25), the complete construction runs in second-order polynomial bit time, with a bound polynomial in $d, n, \ell(p)$, including the intermediate guard bits.

An entrywise absolute guarantee $|H_{ij}-G_{ij}| \le 2^{-b}$ gives $\|H-G\|_F \le \sqrt{d_1 d_2}\, 2^{-b}$. Thus converting an entrywise fixed-point precision to a Frobenius guarantee adds up to $\frac{1}{2}\log_2(d_1 d_2)$ bits. For a floating-point format, combining its relative-error guarantee with a magnitude bound similarly yields the absolute Frobenius radius used by the theorem.

D A fully specified multiplication-only realization

The construction has four steps: express clipping through two matrix absolute values, approximate those values by polynomial iterations, control the accumulated rounding error, and scale the output to ensure feasibility. The approximation and rounding bounds prove Lemma 5; the final scaling completes the upper bound of Theorem 4. Algorithm 1 collects the full procedure.

D.1 Clipping through a symmetric dilation

For real t,

$$\operatorname{clip}(t/\tau, -1, 1) = \frac{|t+\tau| - |t-\tau|}{2\tau}.$$

For a symmetric matrix $B = V \operatorname{diag}(\lambda_i) V^\top$ with V orthogonal, spectral calculus applies a scalar function g to its eigenvalues: $g(B) = V \operatorname{diag}(g(\lambda_i)) V^\top$. In particular, $|B| = (B^2)^{1/2}$ is the positive semidefinite square root of B^2. With $c_\tau(t) = \operatorname{clip}(t/\tau, -1, 1)$, applying the scalar identity gives

$$A = \begin{pmatrix} 0 & H \\ H^\top & 0 \end{pmatrix}, \qquad c_\tau(A) = \frac{|A+\tau I| - |A-\tau I|}{2\tau}. \tag{27}$$

The upper-right block is $C_\tau(H)$. For symmetric B of norm at most four, $Bq_j(B/4)$ approximates $|B|$, where

$$q_0(s) = s, \quad q_{j+1}(s) = \tfrac{3}{2} q_j(s) - \tfrac{1}{2} q_j(s)^3, \quad 0 \le s(1-q_j(s)) \le \tfrac{1}{2}(2/3)^j \tag{28}$$

for $0 \le s \le 1$. The estimate remains valid at zero and gives the geometric error bound behind (11) when applied to $A \pm \tau I$. We next prove it and specify the rounding budgets.

D.2 A scalar Newton-Schulz estimate valid at zero

On $[0,1]$, the approximation $sq_j(s)$ to $|s|$ has error $s(1-q_j(s))$. The factor s suppresses the error near zero, allowing a uniform bound without a positive spectral gap. We prove this weighted error bound in (28). Set

$$F(s) = \frac{3}{2}s - \frac{1}{2}s^3, \qquad q_0(s) = s, \quad q_{j+1} = F \circ q_j, \qquad a_j = (3/2)^j.$$

The map F is increasing on $[0,1]$ and maps it into itself. Consequently $0 \le q_j(s) \le 1$ for all $s \in [0,1]$. For $a \ge 1$ define

$$\phi_a(s) = \frac{as}{\sqrt{1+(a^2-1)s^2}}, \qquad 0 \le s \le 1.$$

Direct substitution gives $\phi_a \circ \phi_b = \phi_{ab}$. Also $F(s) \ge \phi_{3/2}(s)$. To check the latter, set $u = s^2 \in [0,1]$, cancel s when positive, and square the nonnegative sides. The difference after clearing the denominator is

$$\left(\frac{3}{2} - \frac{u}{2}\right)^2 \left(1 + \frac{5}{4}u\right) - \frac{9}{4} = \frac{u(1-u)(21-5u)}{16} \ge 0.$$

The equality at $s = 0$ is immediate. Induction, monotonicity of F, and the composition identity yield

$$q_j(s) \geq \phi_{a_j}(s) \geq \frac{a_j s}{\sqrt{1 + a_j^2 s^2}}.$$

Writing $t = a_j s$,

$$\begin{aligned} 0 \leq s(1 - q_j(s)) &\leq \frac{1}{a_j}\, t \left(1 - \frac{t}{\sqrt{1+t^2}}\right) \\ &\leq \frac{1}{a_j} \frac{t}{1+t^2} \leq \frac{1}{2a_j}. \end{aligned}$$

The middle inequality follows because $\sqrt{1+t^2}(\sqrt{1+t^2} + t) \geq 1 + t^2$ after rationalizing $1 - t/\sqrt{1+t^2}$. This proves the asserted estimate, including at zero. Oddness extends it to the absolute-value approximation on $[-1, 1]$.

If B is symmetric and $\|B\|_{\mathrm{op}} \leq 4$, spectral calculus therefore gives

$$\||B| - Bq_j(B/4)\|_{\mathrm{op}} \leq \frac{2}{a_j}, \qquad \||B| - Bq_j(B/4)\|_F \leq \frac{2\sqrt{d}}{a_j}. \tag{29}$$

This absolute-value estimate holds uniformly through zero, providing the spectral-gap independence needed by the clipping construction.

D.3 Analytic clipping error

For $\|H\|_{\mathrm{op}} \leq 2$ and $0 < \tau \leq 1$, form the dilation A from (27), and set $B_\pm = A \pm \tau I_d$. Their norms are at most three, so (29) applies. The upper-right block of $c_\tau(A)$ is $C_\tau(H)$: an SVD decomposes the dilation into blocks with eigenvalues $\pm\sigma_i$, and the odd function c_τ produces the corresponding singular-value transformation. Extra zero directions contribute zero.

Replacing each $|B_\pm|$ by $B_\pm q_j(B_\pm/4)$ incurs total Frobenius error at most

$$\frac{2\sqrt{d}}{\tau a_j}. \tag{30}$$

Block extraction cannot increase the Frobenius norm. Choose the smallest integer $j \geq 0$ with

$$a_j \geq \frac{32d}{\tau\zeta}. \tag{31}$$

Then (30) is at most $\zeta/(16\sqrt{d}) \leq \zeta/16$. The count j has the order in (11); it can be selected by exact rational comparisons with $(3/2)^j$.

D.4 Dyadic rounding and a certified bit budget

Let Round_b round each entry to the nearest multiple of 2^{-b}, with a fixed deterministic tie convention. Its Frobenius error on a $d \times d$ matrix is at most $d2^{-b}$, a bound that simplifies the bit budget. Choose

$$b = \left\lceil \log_2 \frac{2^{16} d}{\tau\zeta} \right\rceil + 3j, \qquad \eta = d2^{-b}. \tag{32}$$

For each sign, compute

$$\widetilde{X}_0^{\pm} = \mathrm{Round}_b(B_\pm/4), \qquad \widetilde{X}_{t+1}^{\pm} = \mathrm{Round}_b\left(\frac{3}{2}\widetilde{X}_t^{\pm} - \frac{1}{2}(\widetilde{X}_t^{\pm})^3\right). \tag{33}$$

All arithmetic between these explicit roundings is exact rational arithmetic. This ensures that the only recurrence errors are the specified rounding errors. Matrix multiplication itself can be performed on binary integers; a denominator 2^{3b+1} suffices for a cubic update before rounding, so intermediate lengths remain polynomial.

For matrices of spectral norm at most two,

$$\|F(X) - F(Y)\|_F \le \frac{3}{2}\|X - Y\|_F + \frac{1}{2}\|X^3 - Y^3\|_F \le \frac{15}{2}\|X - Y\|_F \le 8\|X - Y\|_F.$$

Here $X^3 - Y^3 = X^2(X-Y) + X(X-Y)Y + (X-Y)Y^2$, so the cubic contribution is bounded by $12\|X - Y\|_F$ for general matrix pairs. Exact iterates $X_t^{\pm} = q_t(B_{\pm}/4)$ have spectral norm at most one. Let $e_t = \|\widetilde{X}_t^{\pm} - X_t^{\pm}\|_F$. As long as the approximate iterates have norm at most two,

$$e_0 \le \eta, \qquad e_{t+1} \le 8e_t + \eta, \qquad e_t \le 2\,8^t\eta.$$

The budget (32) implies, for all $t \le j$,

$$e_t \le \frac{\tau\zeta}{2^{15}} < 1. \tag{34}$$

This proves the norm assumption inductively and closes the error argument.

Compute, using exact rational arithmetic from the rounded iterates,

$$\widetilde{T} = \frac{B_+\widetilde{X}_j^+ - B_-\widetilde{X}_j^-}{2\tau}, \qquad \widetilde{C} = (\widetilde{T})_{1:d_1,\, d_1+1:d}. \tag{35}$$

The additional Frobenius error from (34) is at most

$$\frac{4e_j + 4e_j}{2\tau} \le \frac{\zeta}{8192}.$$

Together with (30), this is strictly less than ζ. Thus (35) satisfies Lemma 5. The remaining error budget also permits an optional final rounding.

D.5 Complete precision-to-descent procedure

Algorithm 1 Certified spectral descent from one Cauchy-oracle query

Require: Dimensions d_1, d_2; known $1 \le r \le k$; $n \ge 1$; oracle for $\|G\|_{\mathrm{op}} \le 1$, $\operatorname{rank} G \le r$.

1: $\varepsilon \leftarrow 2^{-n}$; $a \leftarrow \min\{a \in \mathbb{Z}_{\ge 0} : r \le 4^a\}$; $p \leftarrow n + a + 1$.
2: Query $H \leftarrow \varphi(1^p)$; set $\tau \leftarrow \varepsilon/(8r)$, $\zeta \leftarrow \varepsilon/(16k)$.
3: Form A and $B_{\pm}$ as in (27); choose j, b by (31)-(32).
4: Compute both rounded recurrences (33) for j iterations.
5: Extract $\widetilde{C}$ using (35).
6: **return** $D = \widetilde{C}/(1+\zeta)$.

Ensure: D is rational, $\|D\|_{\mathrm{op}} \le 1$, and $\mathcal{E}(G, D) \le 2^{-n}$.

The norm promise $\|H\|_{\mathrm{op}} \le 2$ required by the lemma follows from $\|G\|_{\mathrm{op}} \le 1$ and $\|H - G\|_F \le 2^{-p} < 1$. The returned scaling is rational even if τ is not dyadic. Its feasibility follows from

$$\|\widetilde{C}\|_{\mathrm{op}} \le \|C_\tau(H)\|_{\mathrm{op}} + \|\widetilde{C} - C_\tau(H)\|_F \le 1 + \zeta.$$

Furthermore,

$$D - C_\tau(H) = \frac{\widetilde{C} - C_\tau(H) - \zeta C_\tau(H)}{1+\zeta},$$

which proves (26). This scaling certifies feasibility directly from the approximation error bound.

The dominant construction uses $O(j)$ dense $d \times d$ products. With schoolbook integer multiplication, a cubic recurrence uses $O(d^3(b + \log d)^2)$ bit operations, together with polynomial overhead for encoding and rounding. The final rational products include the original input encoding length. If that length is J, an overall bound is $\mathrm{poly}(d, J, j, b)$; substituting (25) gives the second-order polynomial bit-time bound $\mathrm{poly}(d, n, \ell(p))$. The explicit budgets give a reproducible finite-bit realization of the descent guarantee.

Stability with a finite output precision. Any two outputs approximating the ideal clipping map within ζ satisfy

$$\|\widetilde{C}(H) - \widetilde{C}(H')\|_F \le \tau^{-1}\|H - H'\|_F + 2\zeta.$$

The Lipschitz constant of C_τ and the additive 2ζ representation allowance together give an explicit stability certificate for finite outputs.

E Spectral reconstruction and direct-descent precision savings

E.1 Proof of Theorem 9

For $0 < \alpha \le 1$, $h_\alpha(t) = \operatorname{sign}(t)|t|^\alpha$ satisfies

$$|h_\alpha(x) - h_\alpha(y)| \le 2^{1-\alpha}|x - y|^\alpha. \tag{36}$$

For x, y of the same sign, the stronger constant one follows from subadditivity of t^α on nonnegative numbers. For opposite signs, concavity gives $|x|^\alpha + |y|^\alpha \le 2^{1-\alpha}(|x| + |y|)^\alpha$.

Let $A = \sum_i a_i u_i u_i^\top$ and $B = \sum_j b_j v_j v_j^\top$ be symmetric $m \times m$ matrices. Expanding the Frobenius square in their two orthonormal eigenbases gives the exact identity

$$\|h_\alpha(A) - h_\alpha(B)\|_F^2 = \sum_{i,j} |h_\alpha(a_i) - h_\alpha(b_j)|^2 |\langle u_i, v_j\rangle|^2. \tag{37}$$

Indeed, expansion of the left side produces $\operatorname{tr} h_\alpha(A)^2 + \operatorname{tr} h_\alpha(B)^2 - 2\operatorname{tr}(h_\alpha(A)h_\alpha(B))$, and each term has the asserted cross-basis expression. The same identity with $h(t) = t$ gives

$$\|A - B\|_F^2 = \sum_{i,j} |a_i - b_j|^2 |\langle u_i, v_j\rangle|^2.$$

The weights sum to m. Apply (36) and Jensen's inequality to the probability weights $|\langle u_i, v_j\rangle|^2/m$ and the concave map $z \mapsto z^\alpha$. It follows that

$$\|h_\alpha(A) - h_\alpha(B)\|_F \le 2^{1-\alpha} m^{(1-\alpha)/2} \|A - B\|_F^\alpha. \tag{38}$$

For a rectangular matrix G, let $\mathcal{D}(G) = \left(\begin{smallmatrix} 0 & G \\ G^\top & 0 \end{smallmatrix}\right)$. Then

$$h_\alpha(\mathcal{D}(G)) = \mathcal{D}(\Phi_\alpha(G)), \qquad \|\mathcal{D}(G)\|_F = \sqrt{2}\|G\|_F.$$

The two symmetric dilations of G, H vanish on the orthogonal complement of the sum of their ranges; because they are symmetric, this sum is an invariant subspace for both. Its dimension is at most $2\operatorname{rank} G + 2\operatorname{rank} H \le 4k$. Restricting (38) to that subspace and using the dilation norm identity yields

$$\|\Phi_\alpha(G) - \Phi_\alpha(H)\|_F \le (8k)^{(1-\alpha)/2} \|G - H\|_F^\alpha. \tag{39}$$

This proves the upper half of Theorem 9 with the plug-in rule $\mathcal{R}(H) = \Phi_\alpha(H)$. The estimate holds uniformly across spectral gaps and changes of eigenspaces.

For the lower half, let Q be a rank-k partial isometry with all its nonzero singular values equal to one, and take $H = 0$ and $G_\pm = \pm\delta Q/\sqrt{k}$. They satisfy $\|G_\pm\|_{\mathrm{op}} \le 1$ for $\delta \le 1$. Their transformed outputs are opposite and each has Frobenius norm $k^{(1-\alpha)/2}\delta^\alpha$. By the triangle inequality, any reconstruction at $H = 0$ has at least this error on one of the inputs. This proves the lower bound.

E.2 Input-bit interpretation and computability

Input precision for a fixed exponent. If $\delta = 2^{-p}$, the lower bound shows that 2^{-n} reconstruction accuracy necessitates

$$\alpha p \ge n + \frac{1-\alpha}{2}\log_2 k.$$

The plug-in map has information error at most 2^{-n-1} if

$$\alpha p \ge n + 1 + \frac{1-\alpha}{2}\log_2(8k).$$

The remaining half of the tolerance may be used to approximate the finite-dimensional spectral map numerically. For computable rational $\alpha > 0$, that map can be computed to arbitrary finite tolerance by ordinary spectral approximation or polynomial functional calculus. Thus the information precision, allowing a constant error-budget margin, is

$$\frac{n + \frac{1-\alpha}{2} \log_2 k + O(1)}{\alpha}.$$

For each rational $\alpha > 0$, this bound specifies the maximal input precision sufficient for a computable finite reconstruction. It complements the uniform descent construction of Appendix D, which runs in second-order polynomial bit time.

Computability for rational exponents. To establish computability uniformly across spectral gaps for rational α, enclose the symmetric dilation spectrum in a finite rational interval, approximate the continuous scalar h_α uniformly by a rational polynomial, and evaluate that polynomial on the matrix. A computable modulus (an algorithmic input-accuracy bound sufficient for any requested output accuracy) follows from (36).

Concretely, rescale the interval to $[0, 1]$ and denote the resulting scalar function by g. Equation (36) supplies a computable constant M_α with $|g(x) - g(y)| \le M_\alpha |x - y|^\alpha$. For an integer $m \ge 1$, its Bernstein polynomial is

$$\mathcal{B}_m g(x) = \sum_{i=0}^{m} g(i/m) \binom{m}{i} x^i (1-x)^{m-i}.$$

The weights are the probabilities of a binomial random variable Y with parameters m, x. Since $\mathbb{E}(Y/m) = x$ and $\mathrm{Var}(Y/m) = x(1-x)/m$, Jensen's inequality gives

$$\sup_{x \in [0,1]} |\mathcal{B}_m g(x) - g(x)| \le M_\alpha (4m)^{-\alpha/2}.$$

Replacing each sample $g(i/m)$ by a rational approximation within $\xi > 0$ adds at most ξ to this error, because the nonnegative weights sum to one. Thus m and the sample accuracy can be chosen effectively for any uniform tolerance, yielding a rational polynomial after rescaling back to the original interval. This supplies a computable realization of the reconstruction accuracy bound. The explicit descent construction additionally achieves second-order polynomial bit time, with specified iteration and rounding budgets.

Precision when the exponent varies. For $\alpha_b = 2^{-b}$, the pair $x = 0$ and $x_b = 2^{-2^b}$ has transformed outputs zero and $1/2$. Before requesting precision at least 2^b, a Cauchy machine can receive the same zero answers. Output tolerance $1/8$ then has disjoint success sets. In this oracle model, writing the precision request alone costs $\Omega(2^b)$ bit steps. Both inputs admit names whose response lengths grow polynomially in the requested precision; the rational exponent's description has only $O(b)$ bits. This proves an exponential uniform precision lower bound using names of polynomial length.

Noncomputability of canonical polar. In this oracle model, Type-2 computability requires an algorithm that, from any valid Cauchy name of the input and any requested output accuracy, terminates with a finite approximation within that accuracy. Canonical polar is discontinuous at the scalar zero input: $P_0(0) = 0$ and $P_0(x) = 1$ for every $x > 0$. A finite computation on a zero Cauchy name makes finitely many requests. A sufficiently small positive scalar admits the same answers to all those requests, although the required outputs at tolerance less than $1/2$ differ. Hence canonical polar across rank changes is not Type-2 computable. A supplied positive spectral gap supports polar reconstruction on the corresponding domain, while Theorem 4 provides a uniform procedure for approximate descent in second-order polynomial bit time across rank changes.

E.3 Fractional bias and the two-stage precision comparison

The transformed singular values lie in $[0, 1]$ whenever $\|G\|_{\mathrm{op}} \le 1$, so $\Phi_\alpha(G) \in K$. The loss separates as

$$\mathcal{E}(G, \Phi_\alpha(G)) = \sum_{i=1}^{k} \big(\sigma_i - \sigma_i^{1+\alpha}\big).$$

The scalar loss $s - s^{1+\alpha}$ has derivative $1 - (1+\alpha)s^{\alpha}$. The loss attains its maximum at $s_\alpha = (1+\alpha)^{-1/\alpha}$, with value $\alpha/(1+\alpha)^{1+1/\alpha}$. Taking all k singular values equal to s_α attains the sum of these maxima and proves (20).

To verify the bracketed constants, set $g(\alpha) = (1+\alpha)\log(1+\alpha)/\alpha$. Then

$$g'(\alpha) = \frac{\alpha - \log(1+\alpha)}{\alpha^2} \geq 0, \quad \lim_{\alpha \downarrow 0} g(\alpha) = 1, \quad g(1) = \log 4.$$

Thus $e \leq (1+\alpha)^{1+1/\alpha} \leq 4$ for $0 < \alpha \leq 1$.

We first make feasibility explicit. Set $\eta = \varepsilon/(2\sqrt{k})$ for $0 < \varepsilon \leq 1/2$. If $\|\widetilde{D} - \Phi_\alpha(G)\|_F \leq \eta$, nonexpansiveness and $\Phi_\alpha(G) \in K$ give

$$\|\Pi_K(\widetilde{D}) - \Phi_\alpha(G)\|_F \leq \eta.$$

When the exponent has descent bias at most $\varepsilon/2$, a feasible reconstruction D within this tolerance satisfies

$$\mathcal{E}(G, D) \leq \mathcal{E}(G, \Phi_\alpha(G)) + \|G\|_F \eta \leq \varepsilon/2 + \sqrt{k}\eta = \varepsilon.$$

For rational $\alpha > 0$, obtain a finite rational realization by first computing rational $\widetilde{D}$ within $\eta/2$ of $\Phi_\alpha(G)$, using the computability construction above. Then $\|\widetilde{D}\|_{\mathrm{op}} \leq 1 + \eta/2 < 2$. Apply Lemma 5 with threshold one and $\zeta = \varepsilon/(16k)$ to approximate $\Pi_K(\widetilde{D})$ by rational $\widehat{C}$ within ζ, and return $D = \widehat{C}/(1+\zeta)$. The feasibility and scaling argument in (26) gives $D \in K$ and

$$\|D - \Phi_\alpha(G)\|_F \leq \eta/2 + \zeta(1+\sqrt{k}) \leq 3\eta/4 < \eta.$$

The extra accuracy changes only the constant error-budget allowance in the reconstruction precision bound and requires no further access to G after computing $\widetilde{D}$.

Suppose that the first stage of the procedure described in Section 5 enforces bias at most $\varepsilon/2$. The lower estimate $k\alpha/4$ then gives $\alpha \leq 2\varepsilon/k$. Requiring reconstruction tolerance $\eta = \varepsilon/(2\sqrt{k})$ and applying the information lower bound gives

$$\begin{aligned} \alpha p &\geq \log_2(1/\eta) + \frac{1-\alpha}{2}\log_2 k \\ &= \log_2(2/\varepsilon) + (1 - \alpha/2)\log_2 k \\ &\geq \log_2(2/\varepsilon) + \tfrac{1}{2}\log_2 k. \end{aligned}$$

Combining these inequalities proves (21), uniformly for $k \geq 1$ and $0 < \varepsilon \leq 1/2$. The direct algorithm's logarithmic input-precision budget therefore gives a quantified saving over the two-stage fractional reconstruction procedure, while certifying a feasible direction with the same descent tolerance.

F Optimization certificates and sharp uncertainty dependence

F.1 Proof of Proposition 6

Proof. Substitute $Z = -\eta D_t$ in (12) to obtain $f(W_{t+1}) \leq f(W_t) - \eta\|\nabla f(W_t)\|_* + \eta\rho_t + \beta\eta^2/2$. Summation and $f(W_T) \geq \inf f$ yield (13). Corollary 2 gives $\rho_t \leq \sqrt{k}\delta + \omega_t$, and the stated step size yields (14). □

F.2 Certifying the implementation error

Here we make the implementation-error certificate explicit. If the exact clipped direction is $C_t \in K$, the feasible implementation is $D_t \in K$, and $\|D_t - C_t\|_F \leq e_t$, then

$$\mathcal{E}(\nabla f(W_t), D_t) \leq \mathcal{E}(\nabla f(W_t), C_t) + \|\nabla f(W_t)\|_F e_t.$$

Thus a supplied bound $M_t \geq \|\nabla f(W_t)\|_F$ certifies $\omega_t = M_t e_t$. Combining this gradient bound with the construction of Appendix D gives an explicit implementation budget for the convergence proposition. The certificate directly translates arithmetic accuracy into its contribution to optimization progress.

F.3 Obtaining the uncertainty radius from a minibatch

Let $\mathcal{F}_t$ contain the history before sampling at iterate W_t, and write $G_t = \nabla f(W_t)$. Conditional on $\mathcal{F}_t$, suppose the N_t sampled per-example gradients $X_{t,i}$ are independent and satisfy

$$\mathbb{E}[X_{t,i} \mid \mathcal{F}_t] = G_t, \qquad \mathbb{E}[\|X_{t,i} - G_t\|_F^2 \mid \mathcal{F}_t] \le \sigma_t^2.$$

The batch size $N_t \ge 1$ and bound $\sigma_t \ge 0$ are chosen before sampling. For example, independent sampling with replacement from a finite training set has mean equal to the gradient of its empirical objective. Define the exact minibatch mean $\overline{G}_t = N_t^{-1} \sum_i X_{t,i}$. Conditional independence and centering eliminate the cross terms, giving

$$\mathbb{E}[\|\overline{G}_t - G_t\|_F^2 \mid \mathcal{F}_t] = \frac{1}{N_t^2} \sum_{i=1}^{N_t} \mathbb{E}[\|X_{t,i} - G_t\|_F^2 \mid \mathcal{F}_t] \le \frac{\sigma_t^2}{N_t}. \tag{40}$$

One way to justify σ_t is a bound $\|X_{t,i}\|_F \le R_t$ on all possible sampled gradients. The conditional variance identity then gives $\mathbb{E}[\|X_{t,i} - G_t\|_F^2 \mid \mathcal{F}_t] = \mathbb{E}[\|X_{t,i}\|_F^2 \mid \mathcal{F}_t] - \|G_t\|_F^2 \le R_t^2$, so $\sigma_t = R_t$ is valid. For a layer with per-example gradient $X_{t,i} = u_{t,i} v_{t,i}^\top$, uniform bounds $\|u_{t,i}\|_2 \le A_t$ and $\|v_{t,i}\|_2 \le B_t$ give $R_t = A_t B_t$ by $\|uv^\top\|_F = \|u\|_2 \|v\|_2$.

For a confidence level $1 - \nu_t$ with $\nu_t \in (0, 1)$ fixed before sampling, Markov's inequality applied to this squared norm yields

$$\Pr\big(\|\overline{G}_t - G_t\|_F > s_t \mid \mathcal{F}_t\big) \le \nu_t, \qquad s_t = \frac{\sigma_t}{\sqrt{N_t \nu_t}}. \tag{41}$$

When $\sigma_t = 0$, the error is zero almost surely and the conclusion holds directly. If the computed observation satisfies $\|H_t - \overline{G}_t\|_F \le q_t$ almost surely, the triangle inequality gives the certified radius $\delta_t = s_t + q_t$, proving (9). Here q_t accounts for the total error in forming and representing the mean. For an exact mean followed by entrywise rounding with error at most u_t per entry, the choice $q_t = \sqrt{d_1 d_2}\, u_t$ gives a concrete certificate.

For fixed r, σ, q, ν, calibrating the clipping threshold to the uncertainty radius certifies a target ideal loss $\varepsilon > \sqrt{r}\, q$ with probability at least $1 - \nu$ whenever the integer batch size satisfies

$$N \ge \max\left\{1, \left\lceil \frac{r\sigma^2}{\nu(\varepsilon - \sqrt{r}\, q)^2} \right\rceil \right\}. \tag{42}$$

Indeed, this makes $\sqrt{r}(\sigma/\sqrt{N\nu} + q) \le \varepsilon$. The formula coordinates two resources: choose numerical accuracy with $\sqrt{r}\, q < \varepsilon$, then allocate enough samples to meet the remaining tolerance. A prescribed sensitivity budget L is incorporated through $\tau_t = \max\{L^{-1}, \delta_t/\sqrt{r}\}$ and the corresponding minimax guarantee in Theorem 1. Equation (42) supplies a sufficient sample budget from a second-moment bound. The per-example norm bounds above provide concrete choices of σ_t certified before sampling. Stronger distributional assumptions support tighter concentration bounds in place of (41).

For a fixed horizon T, choose deterministic allocations with $\sum_{t<T} \nu_t \le \nu$. Taking expectations of the conditional failure probabilities and using a union bound shows that all radii hold simultaneously with probability at least $1 - \nu$, including along adaptive iterates. For example, $\nu_t = \nu/T$ gives $s_t = \sigma_t \sqrt{T/(N_t \nu)}$. On this event, clipping at $\delta_t/\sqrt{k}$ for positive radii gives $\rho_t \le \sqrt{k}\, \delta_t + \omega_t$ in (13). A zero radius permits the exact polar direction. The resulting bound $T^{-1} \sum_{t<T} (\sqrt{k}\, \delta_t + \omega_t)$ on average directional error connects the confidence allocation across iterations directly to a guarantee for the optimization trajectory.

F.4 Adding momentum drift to the radius

Li et al. (2026) analyze the denoising and alignment benefits of momentum under a structured signal-plus-perturbation model. We bound momentum's error relative to the current gradient in terms of sampling error and temporal drift. Consider the ideal, unquantized exponential moving average

$$M_0^{\mathrm{mom}} = \overline{G}_0, \qquad M_t^{\mathrm{mom}} = \mu M_{t-1}^{\mathrm{mom}} + (1 - \mu)\overline{G}_t, \qquad 0 \le \mu < 1.$$

Suppose d_t bounds $\|G_t - G_{t-1}\|_F$. On the simultaneous sampling event above, define $b_0 = s_0$ and

$$b_t = \mu b_{t-1} + \mu d_t + (1-\mu)s_t. \tag{43}$$

Subtracting G_t from the momentum recurrence gives

$$M_t^{\mathrm{mom}} - G_t = \mu(M_{t-1}^{\mathrm{mom}} - G_{t-1}) + \mu(G_{t-1} - G_t) + (1-\mu)(\overline{G}_t - G_t).$$

The triangle inequality and induction prove $\|M_t^{\mathrm{mom}} - G_t\|_F \le b_t$. If the actual observation satisfies $\|H_t - M_t^{\mathrm{mom}}\|_F \le q_t$, the appropriate radius is therefore $\delta_t = b_t + q_t$. Here q_t bounds the accumulated numerical deviation from the ideal momentum, including earlier buffer roundings. For a different initialization, replace b_0 by a valid initial error bound.

One concrete drift bound follows from an L_f-Lipschitz gradient in Frobenius norm: for spectral-feasible steps $W_t = W_{t-1} - \eta_{t-1}D_{t-1}$ with $\eta_{t-1} > 0$, take $d_t = L_f\eta_{t-1}\sqrt{k}$. If $s_t \le s$, $d_t \le d$, and $b_0 \le s$, the recurrence gives $b_t \le s + \mu d(1-\mu^t)/(1-\mu)$. This certificate identifies the sampling and momentum-drift contributions explicitly. Together with the numerical term q_t, it guides the choice of batch size, step size, momentum, and arithmetic precision for a target descent guarantee.

F.5 Proof of Theorem 7 and finite-observation realization

Proof. Let $u = Q/\sqrt{k}$ for a rank-k partial isometry, and choose a positive rational β_0. Define

$$f_\pm(W) = \frac{\beta_0}{2}\|W \mp (\delta/\beta_0)u\|_F^2. \tag{44}$$

Their gradients are $\beta_0 W \mp \delta u$, so the common response $H(W) = \beta_0 W$ has error exactly δ for both functions. Their Hessians are both $\beta_0 I$. Hence their transcripts are identical for every adaptive strategy. At any common output, $\|\nabla f_+(\widehat{W})\|_* + \|\nabla f_-(\widehat{W})\|_* \ge 2\delta\|u\|_* = 2\sqrt{k}\delta$. The success sets below this threshold are disjoint, which also rules out randomized success probability $2/3$ for both inputs. Taking expectations under the common output distribution also gives $\max_{s\in\{+,-\}} \mathbb{E}\|\nabla f_s(\widehat{W})\|_* \ge \sqrt{k}\delta$. □

Taking $\beta_0 = 1$ and restricting query points to rational matrices makes the common gradient response $H(W) = W$ and the identity Hessian exactly representable. The same construction therefore proves Theorem 7 in the finite observation model.

Selecting an independent uniform trajectory index extends the returned-point lower bound to the expected average criterion in (14), certifying the sharpness of its uncertainty term. The larger family below addresses the true best iterate. Both constructions establish this optimality with gradient observations and exact Hessians.

F.6 Proof of Theorem 8 and sharper finite-horizon bounds

Proof. Step 1: Indistinguishable quadratics. Fix the initial point W_0 and curvature $\lambda > 0$. For $s \in \{-1,1\}^k$, let Q_s be the rectangular diagonal matrix with diagonal s, and set $a = \delta/\sqrt{k}$. Consider

$$f_s(W) = \frac{\lambda}{2}\|W - W_0 - (a/\lambda)Q_s\|_F^2, \qquad H(W) = \lambda(W - W_0). \tag{45}$$

Each Q_s has k singular values equal to one. Hence $\nabla f_s(W) = H(W) - aQ_s$, the common gradient response has Frobenius error exactly δ, and every Hessian equals λI. Conditional on the algorithm's random seed, its complete transcript and all its iterates are therefore independent of s. The same conclusion holds with any finite number of adaptive auxiliary queries. All objectives have initial gap

$$f_s(W_0) - \inf f_s = \frac{\delta^2}{2\lambda},$$

and satisfy the spectral-norm smoothness assumption with $\beta = k\lambda$, since

$$f_s(W+Z) - f_s(W) - \langle \nabla f_s(W), Z\rangle = \frac{\lambda}{2}\|Z\|_F^2 \le \frac{k\lambda}{2}\|Z\|_{\mathrm{op}}^2.$$

Step 2: Reduction to sign sums. For a fixed random seed, define

$$b_{ti} = \frac{\lambda(W_t - W_0)_{ii}}{a}, \qquad c_{ti} = \mathrm{clip}(b_{ti}, -1, 1), \qquad 1 \le i \le k.$$

Spectral/nuclear duality gives $\|B\|_* \ge \sum_{i=1}^{k} |B_{ii}|$ for every rectangular matrix B: pair B with the rectangular diagonal matrix of the signs of its diagonal entries, whose spectral norm is at most one. Also, $|b_{ti} - s_i| \ge |c_{ti} - s_i| = 1 - s_i c_{ti}$. Consequently,

$$B_T(f_s) \ge a\left(k - \max_{t<T} X_t\right), \qquad X_t = \sum_{i=1}^{k} s_i c_{ti}. \tag{46}$$

Step 3: Expectation bound. Temporarily choose the s_i independently and uniformly from $\{-1, 1\}$. For every $\theta > 0$,

$$\mathbb{E}_s e^{\theta X_t} = \prod_{i=1}^{k} \cosh(\theta c_{ti}) \le e^{\theta^2 \sum_i c_{ti}^2/2} \le e^{k\theta^2/2}. \tag{47}$$

Here $\cosh x \le e^{x^2/2}$ follows from evenness and $\frac{d}{dx}(x^2/2 - \ln\cosh x) = x - \tanh x \ge 0$ for $x \ge 0$. Jensen's inequality and the bound of a maximum by a sum of exponentials yield

$$\mathbb{E}_s \max_{t<T} X_t \le \frac{\ln T}{\theta} + \frac{k\theta}{2}.$$

For $T > 1$, take $\theta = \sqrt{2\ln T/k}$; for $T = 1$, use $\mathbb{E}_s X_0 = 0$ directly. Substitution into (46), followed by averaging over the algorithm's randomness, gives

$$\mathbb{E}_{s,\mathrm{alg}} B_T(f_s) \ge \sqrt{k}\,\delta\left(1 - \sqrt{\frac{2\ln T}{k}}\right).$$

Nonnegativity supplies the positive part. At least one fixed sign vector has algorithm-expected loss at least this average, proving (15). This sign vector is chosen independently of the realized algorithmic randomness.

Step 4: Probability bound. For $0 < \varepsilon < 1$, the event $B_T(f_s) \le (1-\varepsilon)\sqrt{k}\,\delta$ implies $\max_{t<T} X_t \ge k\varepsilon$ by (46). Applying Markov's inequality to (47) and optimizing at $\theta = \varepsilon$ gives $\Pr_s(X_t \ge k\varepsilon) \le e^{-k\varepsilon^2/2}$. A union bound over the T iterates, followed by averaging over algorithmic randomness, bounds the mean success probability by $Te^{-k\varepsilon^2/2}$. Thus some fixed sign vector has success probability no greater than this bound. Truncating at one proves (16). The sign vector may depend on ε and need not be the one witnessing the expectation bound. □

The argument applies equally to any list of at most T candidate points; a shorter nonempty list can be padded by repetitions. For rational W_0, rational $\lambda > 0$, and rational query points, the common responses in (45) and the exact Hessians have finite rational encodings. The hard minimizers may be real, as permitted by the oracle model. Thus the construction also realizes the lower bound with finite observations.

Sharper dependence on the horizon. Define

$$I(q) = \frac{1+q}{2}\ln(1+q) + \frac{1-q}{2}\ln(1-q), \qquad 0 \le q \le 1,$$

with $0\ln 0 = 0$. Since $|c_{ti}| \le 1$, the stronger moment bound $\mathbb{E}_s e^{\theta X_t} \le (\cosh\theta)^k$ holds. Optimizing the corresponding Chernoff bound at $\theta = \operatorname{arctanh}\varepsilon$ gives a fixed instance satisfying

$$\Pr\Big(B_T(f) \le (1-\varepsilon)\sqrt{k}\,\delta\Big) \le \min\{1, Te^{-kI(\varepsilon)}\}. \tag{48}$$

Here $I(0) = I'(0) = 0$ and $I''(q) = 1/(1-q^2) \ge 1$, so $I(q) \ge q^2/2$, recovering (16). In particular, a uniform success probability of at least $2/3$ below $(1-\varepsilon)\sqrt{k}\,\delta$ requires at least $(2/3)e^{kI(\varepsilon)}$ candidates.

The sharper moment bound also gives

$$\mathbb{E}_s \max_{t<T} X_t \le \inf_{\theta>0} \frac{\ln T + k \ln\cosh\theta}{\theta}.$$

For $1 < T < 2^k$, the minimizing θ satisfies $\theta\tanh\theta - \ln\cosh\theta = \ln T/k$. The left side is $I(\tanh\theta)$, so the infimum equals $kI^{-1}(\ln T/k)$, where $I^{-1} : [0, \ln 2] \to [0, 1]$ is the increasing inverse. The endpoints follow by limits. Averaging over algorithmic randomness and selecting a fixed sign vector as above proves, for $1 \le T \le 2^k$,

$$\mathbb{E}B_T(f) \ge \sqrt{k}\,\delta \left[1 - I^{-1}\left(\frac{\ln T}{k}\right)\right]. \tag{49}$$

Exploration and the best-iterate criterion. The horizon dependence accounts for the possibility of exploring several candidate minimizers. For example, in one dimension let $f_b(w) = \lambda(w-b)^2/2$ and observe h_0 at w_0 with $|h_0 - \lambda(w_0 - b)| \le \delta$. An exact Hessian reveals λ. Set $c = w_0 - h_0/\lambda$, so $b \in [c - \delta/\lambda, c + \delta/\lambda]$, and choose $w_1 = c - \delta/(2\lambda)$ and $w_2 = c + \delta/(2\lambda)$. These three iterates satisfy $\min_{0\le t<3} |f_b'(w_t)| \le \delta/2$ for every b and every valid initial observation. This illustrates the benefit of candidate exploration for the true best-iterate criterion. Theorem 7 complements this trajectory analysis by quantifying the information needed to identify a suitable returned point.

F.7 Exact best-iterate sharpness for zero-preserving updates

Consider methods that query gradients only at their current iterate and base their updates on this observation history, optional exact Hessians, and private randomness. Call such a method *zero preserving* if its update is zero almost surely whenever all gradient observations so far are zero and all available Hessian observations equal λI.

Proposition 10 (Sharp uncertainty dependence for zero-preserving updates)**.** *Fix $\delta, \lambda > 0$ and W_0. For every zero-preserving method initialized at W_0, there is a condition-number-one quadratic with Hessian λI and a globally δ-valid gradient oracle for which, almost surely,*

$$B_T(f) = \sqrt{k}\,\delta \qquad \textit{for every finite } T \ge 1. \tag{50}$$

Proof. Fix any sign vector in (45). Since $H(W_0) = 0$, zero preservation implies inductively that every iterate equals W_0 almost surely. At this point $\nabla f_s(W_0) = -aQ_s$ has nuclear norm $ka = \sqrt{k}\,\delta$, proving the claim. □

Calibrated clipping is zero preserving because $C_\tau(0) = 0$, as is its zero-initialized momentum version. For admissible rational inputs and parameters, the certified rounded clipping subroutine in Appendix D also preserves zero exactly. When $H = 0$, the dilation A in (27) is zero and both $B_\pm$ are diagonal. Entrywise dyadic rounding preserves zero entries, so every polynomial iterate in (33) remains diagonal. The extracted off-diagonal block in (35) is therefore exactly zero, and rational feasibility scaling preserves it. Using these outputs as update directions satisfies the zero-preservation requirement whenever the subroutine's input and parameter promises hold. The exact sharpness result therefore applies to these finite clipping updates, and Theorem 8 extends the analysis to unrestricted randomized exploration.

G Second-order complexity of complete optimization trajectories

This appendix extends the local precision and stationarity guarantees to complete deterministic trajectories, accounting for finite iterate representations. Related approaches control evaluation accuracy in nonconvex regularization (Bellavia et al., 2019) and track rounding errors in multi-precision optimization (Monnet & Orban, 2025). Building on the exact-gradient query rate $\Theta(\varepsilon^{-2})$ for deterministic gradient-only optimization (Chewi et al., 2023), we establish a joint precision and iteration bound: $\Omega(\varepsilon^{-2})$ requests require precision $\Omega(\log(1/\varepsilon))$, even in one dimension. Matching upper bounds include arithmetic and explicit trajectory output over the full spectrally feasible class. A further construction supplies a complete certified clipping trajectory.

G.1 Objective class, gradient oracle, and output convention

To isolate accuracy dependence, fix $W_0 = 0$ and consider differentiable, bounded-below objectives satisfying

$$\begin{aligned} f(W+Z) &\le f(W) + \langle \nabla f(W), Z\rangle + \tfrac{1}{2}\|Z\|_{\mathrm{op}}^2, \\ f(0) - \inf f &\le 1. \end{aligned} \tag{51}$$

Normalizing the smoothness and initial-gap bounds makes the optimal dependence on the target accuracy explicit. Let $m = d_1 d_2$, retain $d = d_1 + d_2$ and $k = \min(d_1, d_2)$, and put $\varepsilon = 2^{-n}$ for an integer $n \ge 2$.

A gradient-evaluation name extends the matrix name in Appendix C by accepting both a dyadic query point and a requested precision:

$$\Psi(\mathrm{enc}(W), 1^p) = H, \qquad \|H - \nabla f(W)\|_F \le 2^{-p}. \tag{52}$$

Correctness is required for every name satisfying (52). Requests may be adaptive or repeated; a fixed name returns the same answer to an identical request, and distinct requests may return any answers within their individual error bounds. The algorithm's objective information consists of these gradient observations. The name's size function is

$$\ell(u) = \max_{|v| \le u} |\Psi(v)|,$$

where v is a complete request string, including the query point. Nondistinguished requests return a fixed dummy payload; padding to a nondecreasing common response length at each query length is permitted.

We use the explicit signed-integer and binary-fraction fields of Appendix C, with self-delimiting fields of linear encoding overhead, and deterministic multitape bit computation. For this trajectory task, *each complete request and response is charged by its length*, including repeated fields and complete payloads. Arithmetic and all output writing are charged separately. In particular, each call at precision p costs at least p bit steps. This full-message convention gives an explicit accounting of communication, computation, and output; adding the oracle's internal evaluation cost gives the cost of a concrete gradient implementation. The sharp bounds below use this convention; polynomially equivalent encodings give polynomially related runtime bounds.

The required output consists of all dyadic iterates $W_0, \dots, W_T$ of a finite run with

$$W_{t+1} = W_t - \eta_t D_t, \qquad \eta_t > 0, \quad \|D_t\|_{\mathrm{op}} \le 1,$$

and an index $j \le T$ such that $\|\nabla f(W_j)\|_* \le \varepsilon$. Thus the algorithm identifies a valid output point, in addition to writing the trajectory. Intermediate oracle queries may be arbitrary dyadic matrices. The optimization class permits every spectrally feasible update rule, including basis-dependent and discontinuous rules.

Linear-length names are available throughout this objective class at fixed dimensions. Indeed, applying (51) in the polar direction with step length $\|\nabla f(W)\|_*$ gives

$$\|\nabla f(W)\|_*^2 \le 2\big(f(W) - \inf f\big). \tag{53}$$

In particular $\|\nabla f(0)\|_* \le \sqrt{2}$. Applying (51) at zero and then (53) yields

$$\|\nabla f(W)\|_* \le \sqrt{2} + \|W\|_{\mathrm{op}}.$$

Rounding each gradient entry to $p + \lceil \log_2 \sqrt{m} \rceil + 1$ fractional bits therefore gives valid answers of length $O_{d_1,d_2}(p + |\mathrm{enc}(W)| + 1)$. Fix a sufficiently large encoding constant C_0 and call a name *linear-length* when $\ell(u) \le C_0(u+1)$; correctness is required for every name in this class.

G.2 Optimal bit runtime in the second-order model at fixed dimensions

The upper bound below uses a coordinate update within the full class of spectrally feasible rules, with dimension dependence displayed in (57). Appendix G.4 gives a separate construction based on certified spectral clipping, with its own arithmetic bound.

Theorem 11 (Complete-trajectory accuracy complexity)**.** *In the second-order oracle model of Appendix G.1, for every fixed pair $d_1, d_2 \geq 1$, the optimal worst-case bit runtime over deterministic algorithms and linear-length names is*

$$R^*_{d_1,d_2}(n) = \Theta_{d_1,d_2,C_0}(n4^n) = \Theta_{d_1,d_2,C_0}\left(\varepsilon^{-2}\log\frac{1}{\varepsilon}\right). \tag{54}$$

The upper bound writes the entire trajectory and returns a certified final iterate. The lower bound applies even to algorithms required to return only one approximately stationary point, and already holds for objectives whose matrix gradients have rank at most one.

Proof idea. At fixed dimensions, coordinate updates on a dyadic grid decrease the objective by order ε^2 per nonterminal step and use $O(n)$ bit operations per step, including writing the iterates. For the lower bound, coarse queries hide a scalar gradient of magnitude 2ε, while triangular bumps interpolate the fine-query responses and keep the returned point nonstationary unless $\Omega(\varepsilon^{-2})$ fine queries are made. Each fine request has precision at least n and therefore costs at least n bit operations. □

Proof. Finite-bit upper bound. Let $s = 2^{\lceil \log_2 m \rceil}$ and $q = \varepsilon/s$, so $m \leq s < 2m$. At the current iterate request precision

$$p = n + \log_2 s + 4,$$

and truncate each returned entry toward zero to p fractional bits, obtaining $\bar{H}$. Each entry has total observation and truncation error at most $q/8$. If

$$\max_{ij} |\bar{H}_{ij}| \leq 3q/4,$$

terminate and designate the current iterate. Otherwise choose a maximum-magnitude entry, breaking ties in a fixed order, and update

$$W^+ = W - \frac{q}{2}\operatorname{sign}(\bar{H}_{ij})E_{ij}, \tag{55}$$

where E_{ij} is the corresponding coordinate matrix. Its spectral norm is one, so the update is feasible.

At termination each gradient entry has magnitude at most $7q/8$. Since $\|G\|_* \leq \sum_{ij}|G_{ij}|$, the returned point satisfies

$$\|\nabla f(W)\|_* \leq 7mq/8 \leq 7\varepsilon/8 < \varepsilon.$$

At a nonterminal step, the chosen observation and true gradient entry have the same sign and $|G_{ij}| > 5q/8$. Hence

$$f(W^+) \leq f(W) - \frac{q}{2}|G_{ij}| + \frac{q^2}{8} < f(W) - \frac{3q^2}{16}. \tag{56}$$

The initial-gap bound permits only $O(q^{-2}) = O(s^2 4^n)$ such steps. A terminating query must therefore occur within this bound plus one.

For the arithmetic cost, all iterates lie on the fixed dyadic grid of spacing $q/2$. A coordinate is this spacing times an integer of magnitude at most the number of updates. Consequently every coordinate has $O(n + \log m)$ bits, and there is no accumulation of rounding errors in the iterates. The descent in (56) and (53) also give $\|\nabla f(W_t)\|_* \leq \sqrt{2}$ along the run, so the truncated observations have bounded integer parts. Reading and truncating one complete answer, scanning its entries, updating a grid numerator, and writing a complete iterate take

$$O(\ell(Cm(n + \log m)) + m(n + \log m))$$

bit steps. Summing over the updates gives the second-order runtime bound

$$R(n, \ell) \leq Cs^2 4^n\left[\ell(Cm(n + \log m)) + m(n + \log m)\right] \tag{57}$$

for a sufficiently large absolute encoding constant C. Fixed dimensions and linear-length names yield $O(n4^n)$, including the final output index.

Joint precision and iteration lower bound. First take scalar inputs. Set $a = 2\varepsilon$ and answer a deterministic algorithm's requests by

$$H(x, p) = \begin{cases} 0, & p \leq n - 1, \\ -a, & p \geq n. \end{cases} \tag{58}$$

Call the latter requests fine. The responses can use a fixed linear padding rule, independent of the query point. Suppose the algorithm halts after Q fine requests at $x_1, \dots, x_Q$, returning a designated point $\widehat{x}$. Define

$$S = \{x_1, \dots, x_Q, \widehat{x}\}, \qquad h(x) = \max_{z \in S} (a - |x - z|)_+, \qquad f(x) = \int_x^\infty h(u)\, du. \tag{59}$$

Here $(v)_+ = \max\{v, 0\}$. The function h is nonnegative, compactly supported, and 1-Lipschitz, with $h \le a$. Thus $f' = -h$ is 1-Lipschitz and $\inf f = 0$. Each triangular bump has area a^2, so

$$f(0) - \inf f \le \int_{\mathbb{R}} h(u)\, du \le (Q+1)a^2. \tag{60}$$

If $Q + 1 \le a^{-2}$, this objective belongs to (51). Every coarse answer is valid because $2^{-p} \ge a \ge |f'(x)|$, and each fine answer is exact at its queried point. Nevertheless $|f'(\widehat{x})| = a > \varepsilon$, contradicting correctness.

To realize the resisting argument by a fixed valid input, put $N = a^{-2} = 4^{n-1}$. If the run reaches its Nth fine request, construct (59) from those N fine-query centers, without an output center; its gap is at most one, and its name reproduces that entire prefix. If the run halts earlier, the preceding construction with its output center contradicts correctness. If it neither halts nor reaches N fine requests, the finitely many fine-query centers alone define a valid objective and name reproducing the whole run; with no such centers use $h = 0$. This would violate termination on a valid input. Thus every correct algorithm has a valid input requiring at least N fine requests.

Each of these requests has $p \ge n$ and is charged at least n steps under the complete-message convention. Hence the bit runtime in this oracle model is $\Omega(n4^n)$. For dyadic query centers, the constructed objective is computable and piecewise quadratic. Its historical answers extend to a total valid name by using zero at all coarse precisions and dyadic approximations to $-h$ at other fine queries. These payloads, together with the fixed padding rule, have linear length. The lower bound therefore holds within the stated name class.

Finally embed the scalar objective as $F(W) = f(W_{11})$. Then $\nabla F(W) = f'(W_{11})E_{11}$ has rank at most one, its nuclear norm is $|f'(W_{11})|$, and scalar smoothness implies (51) because $|Z_{11}| \le \|Z\|_{\mathrm{op}}$. Answer coarse requests by zero and fine requests by $-aE_{11}$. The scalar argument now applies to the first coordinate of every queried matrix and of the designated output. This proves the lower bound for every fixed matrix dimension. □

The matching bounds determine the second-order bit complexity of the complete trajectory task as $\Theta(n4^n)$ at fixed dimensions for linear-length names.

G.3 Dependence on the oracle size function

We now extend Theorem 11 beyond linear-length names. Fix dimensions and a nondecreasing integer length budget Λ with $\Lambda(u) \ge C_1(u+1)$ for a sufficiently large encoding constant. Let $R^*(n; \Lambda)$ be the optimal worst-case bit runtime in the second-order oracle model over objectives (51) and valid names with $\ell(u) \le \Lambda(u)$. The budget Λ describes the full class of admissible names.

Corollary 12 (Oracle-size dependence)**.** *In the complete-message model, there are dimension-dependent constants $c, C > 0$ such that*

$$c4^n\big[n + \Lambda(n)\big] \le R^*(n; \Lambda) \le C4^n\big[n + \Lambda(Cn)\big]. \tag{61}$$

If $\Lambda(Cu) = O(\Lambda(u))$ for every fixed $C \ge 1$, then

$$R^*(n; \Lambda) = \Theta(4^n[n + \Lambda(n)])\,.$$

In particular, admissible fixed-degree polynomial length budgets satisfy this condition and yield the sharp runtime formula.

Proof. The upper bound follows from (57). For the lower bound, use the resisting responses (58) and pad every answer at query length u to length $\Lambda(u)$. The assumed lower bound on Λ accommodates every payload of the constructed name, and dummy queries can receive the same padding. The proof of Theorem 11 still forces $\Omega(4^n)$ fine requests. Each has total query length at least n, so writing its request and reading its complete padded answer cost at least $n + \Lambda(n)$. This proves (61); the last conclusion follows from the regularity of Λ. □

G.4 A COMPLETE FINITE-BIT CLIPPING TRAJECTORY

For the clipping construction of Appendix D, a final dyadic approximation of each direction controls iterate denominators and rounding errors while preserving the descent budget.

Proposition 13 (Certified clipping through stationarity)**.** *Under (51) and (52), there is a deterministic implementation using one gradient query per step, with*

$$\eta = \varepsilon/4, \qquad T = 16 \cdot 4^n, \qquad p = n + \left\lceil \tfrac{1}{2} \log_2 k \right\rceil + 6,$$

that writes $W_0, \dots, W_T$ and designates an iterate W_j, $j < T$, satisfying $\|\nabla f(W_j)\|_* \le \varepsilon$. Its directions are finite dyadic approximations to the certified clipping directions, and its total bit runtime in the second-order oracle model is

$$O\big(4^n \left[\ell(Cm(n + \log d)) + d^3(n + \log d)^3\right]\big). \tag{62}$$

For fixed dimensions this is $O(4^n[\ell(Cn) + n^3])$.

Proof. Write $G_t = \nabla f(W_t)$. We construct directions with

$$\|D_t\|_{\mathrm{op}} \le 1, \qquad \mathcal{E}(G_t, D_t) \le \varepsilon/8. \tag{63}$$

Given these inequalities at earlier steps, smoothness implies

$$f(W_t) - \inf f \le 1 + t\varepsilon^2/16 \le 2, \qquad \|G_t\|_* \le 2$$

for $t \le T$, using (53). We now establish (63) at the current step, closing this induction.

Round the raw oracle answer to $p + \lceil \log_2 d \rceil$ fractional bits, obtaining $\bar{H}_t$. With $a = \lceil \frac{1}{2} \log_2 k \rceil$, its error satisfies

$$\|\bar{H}_t - G_t\|_F \le 2^{-(p-1)} = \frac{\varepsilon}{32\, 2^a} \le \frac{\varepsilon}{32\sqrt{k}}.$$

Normalize the input by two and apply the single-query construction of Theorem 4, with $r = k$ and target accuracy $2^{-(n+5)} = \varepsilon/32$. The matrix $\bar{H}_t/2$ has the required approximation error, and $\|G_t/2\|_{\mathrm{op}} \le 1$. The resulting feasible rational direction D_t^0 therefore obeys

$$\mathcal{E}(G_t, D_t^0) \le \varepsilon/16.$$

This invocation uses the already obtained matrix approximation, so it entails no additional gradient query.

To obtain a common dyadic denominator, set $h = n + 7$ and $u = 2^{-h} = \varepsilon/128$. Round every entry of D_t^0 to $h + \lceil \log_2 d \rceil$ fractional bits, obtaining R_t with $\|R_t - D_t^0\|_F \le u$, and set

$$D_t = (1 - u)R_t.$$

Then

$$\|D_t\|_{\mathrm{op}} \le (1 - u)(1 + u) \le 1, \qquad \|D_t - D_t^0\|_{\mathrm{op}} \le 3u.$$

Consequently, spectral/nuclear duality and $\|G_t\|_* \le 2$ give

$$\mathcal{E}(G_t, D_t) \le \varepsilon/16 + 6u = 7\varepsilon/64 < \varepsilon/8.$$

This proves (63) and completes the induction. All D_t share one dyadic denominator with $O(n + \log d)$ bits. Since η is also fixed and dyadic, each update $W_{t+1} = W_t - \eta D_t$ is performed exactly on a common grid. The magnitude bound $\|W_t\|_{\mathrm{op}} \le T\eta = 4 \cdot 2^n$ shows that every iterate entry still has only $O(n + \log d)$ bits.

Proposition 6 yields

$$\frac{1}{T} \sum_{t<T} \|G_t\|_* \le \frac{1}{\eta T} + \frac{\eta}{2} + \frac{\varepsilon}{8} = \frac{\varepsilon}{2}.$$

An output index can be identified without an additional oracle call or a nuclear-norm computation. Evaluate the exact rational scores

$$U_t = \langle \bar{H}_t, D_t \rangle + \varepsilon/4.$$

Since $\|D_t\|_F \le \sqrt{k}$, the observation error contributes at most $\varepsilon/32$ to this inner product. Combining this fact with (63) gives

$$\|G_t\|_* \le U_t \le \|G_t\|_* + 9\varepsilon/32.$$

The index minimizing U_t thus satisfies $\|G_j\|_* \le U_j \le 25\varepsilon/32 < \varepsilon$.

Finally, put $B = n + \log d$. Complete query strings have length $O(mB)$. Each oracle answer is read and rounded once, at cost linear in its length plus $O(mB)$. Thereafter Appendix D uses $O(B)$ dense matrix products with $O(B)$ working bits. Its rounded recurrences have common dyadic denominators, and the final divisions by τ and $1+\zeta$ admit common rational denominators of $O(B)$ bits. Schoolbook integer multiplication therefore gives $O(d^3B^3)$ bit operations per step, including the final dyadic direction conversion. Updating and writing the iterate and comparing the rational scores fit within this bound. Summing over T establishes (62). □