SCALEDGD FOR UNDERPARAMETERIZED MATRIX FACTORIZATION: GLOBAL CONVERGENCE AND STEP-SIZE UNIVERSALITY

Anonymous authors
Paper under double-blind review

ABSTRACT

We study the dynamics and convergence of scaled gradient descent (ScaledGD), a preconditioned gradient method, for underparameterized nonconvex matrix factorization. Given a symmetric positive semidefinite matrix M of rank d, we seek to identify its best rank-k approximation, $k < d$, namely, the leading k-dimensional eigenspace and its associated eigenvalues. Assuming an eigengap at rank k, we prove that ScaledGD converges almost surely to a single global minimizer from a normalized randomized range initialization, for every fixed step size $\eta \in (0, 1)$. This interval is independent of the target spectrum; the explicit asymptotic contraction factor depends on the relative eigengap. A fixed linear change of coordinates reveals that ScaledGD is exactly ordinary gradient descent on a hidden objective. This representation exposes the coupled dynamics of subspace selection and eigenvalue recovery, enabling a global convergence analysis throughout the full step-size interval. We extend the analysis to rectangular matrices through symmetric dilation, under a suitable initialization and a spectrum-dependent step-size bound. For rank-one approximation in both the symmetric and rectangular settings, we obtain stronger, explicit geometric rates from initialization under a strict leading gap. We also establish gap-free $O(1/t)$ bounds on excess reconstruction error. Both guarantees hold for every fixed $\eta \in (0, 1)$ under the respective initialization conditions.

1 INTRODUCTION

Gradient descent for nonconvex matrix factorization has been extensively studied for matrix completion (Jin et al., 2016; Ma et al., 2020), sensing (Tu et al., 2016), and approximation (Chen et al., 2024). Beyond these applications, matrix factorization serves as a structured model for studying optimization on nonconvex landscapes. Motivated by the success of gradient methods in deep learning, researchers have used these models to investigate how landscape geometry, initialization, and implicit regularization shape optimization and the solutions it selects (Ge et al., 2017; Gunasekar et al., 2017; Du et al., 2018).

Preconditioning adds a further dimension to this study. Scaled gradient descent (ScaledGD), which uses a preconditioner determined by the current iterate, achieves condition-number-independent local convergence rates for matrix recovery under suitable assumptions (Tong et al., 2021). These local guarantees motivate studying its global behavior when the rank of the iterates is too small to represent the target.

We study symmetric rank-k matrix factorization. Given a positive semidefinite matrix $M \in \mathbb{R}^{n\times n}$ of rank d, we seek a best rank-k approximation $XX^\top$, with $X \in \mathbb{R}^{n\times k}$, by minimizing $f(X) = \frac{1}{4}\|M - XX^\top\|_F^2$. Our focus is the *underparameterized* regime $k < d$, in which $XX^\top$ cannot represent M exactly; our results also cover exact parameterization ($k = d$). ScaledGD preconditions the gradient with $(X^\top X)^{-1}$:

$$X_{t+1} = X_t - \eta\nabla f(X_t)(X_t^\top X_t)^{-1} = (1-\eta)X_t + \eta M X_t (X_t^\top X_t)^{-1}. \tag{1}$$

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.

Each iteration updates all k columns jointly using a matrix–block product and a $k \times k$ solve. A universal admissible step-size range is appealing in practice: it allows step-size selection without estimating unknown problem-dependent constants, such as spectral bounds or eigengaps. We ask:

Does ScaledGD converge globally to a best rank-k approximation?
If so, for which step sizes?

ScaledGD recovery analyses assume that the target is representable (Tong et al., 2021; Jia et al., 2023). For underparameterized approximation, the Nyström-initialized guarantees of Li et al. (2025) concern a weak optimality criterion that does not identify the leading subspace. Jedra & Shah (2025) establish rank-one convergence through a connection to Heron's method and obtain higher-rank approximations by deflation. These results do not establish convergence of the simultaneous rank-k update over a spectrum-independent step-size interval; §2 gives a detailed comparison.

The difficulty is a matrix-valued coupling between subspace selection and eigenvalue recovery. The current Gram matrix controls subspace progress, while the subspace determines how the scale of the iterates evolves. Consequently, the scalar rank-one analysis does not directly extend to the joint update. Moreover, the inverse Gram matrix must remain well-defined throughout the trajectory, and the limiting subspace must be leading rather than merely invariant.

We give an affirmative answer under a boundary eigengap and a normalized randomized range initialization. A fixed linear change of coordinates reveals that ScaledGD is ordinary gradient descent on a hidden objective separating radial and subspace errors. Although this objective is analytic, its gradient is not globally Lipschitz near rank deficiency; general convergence results for descent methods require additional decrease and trajectory conditions (Attouch et al., 2013). Determinant ascent and a weighted matrix Heron recursion establish finite total movement of our iterates without requiring global descent of the hidden objective. Once convergence to a critical point is established, we identify the optimal limit using a center-stable argument in the spirit of Lee et al. (2016), applying the criterion of Cheridito et al. (2024, Proposition 2.5). §4–§5 develop this argument.

Contributions. Below we present the main contributions of this work.

- *(i) Global rank-k convergence throughout a universal step-size range.* Theorem 3.1 establishes that, under a boundary eigengap and the randomized range initialization of §3, the symmetric ScaledGD iterates converge almost surely to a single global minimizer for every fixed $\eta \in (0, 1)$. We give an explicit asymptotic linear rate for the iterates and their reconstruction, depending on the relative eigengap. The general-rank entrance time remains unquantified.
- *(ii) Rectangular extension and step-size dependence.* For the shared-inverse update induced by symmetric dilation, we establish almost-sure optimal rank-k approximation under a suitable initialization and a spectrum-dependent step-size bound. The symmetric analysis does not extend to the full interval $0 < \eta < 1$ in general: local stability depends on the spectrum (§6), and Figure 1(b) provides empirical support through persistent reconstruction errors above the corresponding stability threshold.
- *(iii) Explicit geometric and gap-free rank-one rates.* For rank one, we strengthen the guarantees in both settings to explicit gap-dependent geometric rates from initialization and gap-free $O(1/t)$ bounds on excess reconstruction error. Both settings admit every fixed $\eta \in (0, 1)$ under the respective initialization assumptions (§7).

Experiments in §8 also show symmetric trajectories reaching numerical precision across tested step sizes and condition numbers. A selected example exhibits small principal-angle reversals despite determinant ascent, supporting our choice of subspace potential. These tests use unnormalized Gaussian initialization; the symmetric cases are covered by §B.8.

2 Related work

Our closest comparisons concern underparameterized GD and ScaledGD. Table 1 summarizes their guarantees; the objectives, initializations, and spectral assumptions differ across works.

Table 1: Selected approximation guarantees. Conditions and distinctions are discussed below and in §E.

Work	Method	Initialization	Guarantee	Step size
Jin et al. (2023)	GD	Small, aligned	Linear*	Problem-dep.
Chen et al. (2024)	GD	Scaled Gaussian	Incremental*	Problem-dep.
Li et al. (2025)	ScaledGD	Nyström	Weak optimality	Fixed / accuracy-dep.
Jedra & Shah (2025)	ScaledGD	Gaussian range	Rank-one linear	$\eta = 1/2$
This work	ScaledGD	Random range	Rank-k linear*	$0 < \eta < 1$

*Jin et al.: sensing under RIP and alignment; Chen et al.: quantitative global rate conditional on saddle escape; ours: almost-sure optimal limit, with unquantified entrance time.

Underparameterized GD and incremental learning. Jin et al. (2023) show that GD from sufficiently small initialization successively approaches best solutions of increasing rank for matrix sensing (Theorem 4.1). Their Theorem 4.2 establishes convergence to the best rank-k solution when $k \leq d$. This optimum is defined by the sensing loss; with full observation it coincides with the best Frobenius-norm approximation. The analysis assumes restricted isometry, initial alignment, distinct positive target eigenvalues, and a sufficiently small problem-dependent step size. Chen et al. (2024) study incremental learning from large initialization; their quantitative global rate is conditional on saddle escape. These works characterize ordinary GD. Our analysis addresses how preconditioning changes the admissible steps and the coupled subspace and scale dynamics.

Underparameterized ScaledGD. Jedra & Shah (2025, Theorem 1) prove global linear convergence for rank-one ScaledGD with $\eta = 1/2$, a leading eigengap, and Gaussian range initialization. Their Heron interpretation explains scale recovery, and their higher-rank construction applies rank-one updates sequentially by deflation. Our results analyze the simultaneous update of all k columns, requiring only a boundary eigengap and admitting every fixed $0 < \eta < 1$ under an absolutely continuous range initialization (Theorem 3.1 and Corollary B.10). Our rank-one results also cover this full interval and give explicit geometric and gap-free reconstruction bounds.

Li et al. (2025, Theorem 2) study ScaledGD from Nyström initialization using weak optimality, $X^\top M^\dagger X = I_k$. This condition holds for appropriately scaled eigenvectors from any k positive eigenpairs, so it does not identify the leading subspace. Their guarantees concern this condition or its neighborhood, depending on the step size. Our theorem identifies an optimal limit, although its general-rank entrance time remains unquantified.

Exact and overparameterized recovery. For representable targets, guarantees include local convergence from spectral initialization (Tu et al., 2016; Tong et al., 2021) and global recovery under suitable random initializations (Stöger & Soltanolkotabi, 2021; Jia et al., 2023; Xu et al., 2023). Underparameterization introduces a nonzero optimal residual and requires selection of the leading subspace in addition to scale recovery.

GD on structured nonconvex objectives. For objectives with the Kurdyka–Łojasiewicz (KL) property, sufficient decrease and relative-error bounds yield convergence of bounded sequences under an appropriate continuity condition (Attouch et al., 2013). Saddle-avoidance results exclude strict-saddle limits under suitable regularity and step-size assumptions (Lee et al., 2016); perturbed GD also admits quantitative escape guarantees (Jin et al., 2017). For our hidden objective, the gradient is not globally Lipschitz near rank deficiency. These general tools therefore require additional trajectory control to establish the full interval $0 < \eta < 1$ and explicit spectral rate dependence. We obtain this control through determinant ascent and radial stabilization, then apply the saddle-avoidance criterion of Cheridito et al. (2024). The rectangular proof also uses KL theory. §E expands these comparisons.

3 Setting and Main Result

→ Symmetric Problem and Initialization. Let $M \in \mathbb{R}^{n\times n}$ be symmetric positive semidefinite of rank $d \geq k$, with reduced eigendecomposition $M = U\Lambda U^\top$. Here $\Lambda = \operatorname{diag}(\lambda_1, \ldots, \lambda_d)$

and $\lambda_1 \geq \cdots \geq \lambda_d > 0$. We assume that the leading k-dimensional invariant subspace is unique, or equivalently that $\Delta_k := \lambda_k - \lambda_{k+1} > 0$, with $\lambda_{d+1} := 0$ when $k = d$. We write $\kappa := \lambda_1/\lambda_d$ for the condition number of M on its range. We denote the $\ell \times \ell$ identity matrix by I_ℓ.

We seek a best rank-k approximation to M in Frobenius norm by solving

$$\min_{X \in \mathbb{R}^{n\times k}} f(X), \qquad f(X) := \frac{1}{4}\|M - XX^\top\|_F^2. \tag{2}$$

Let $U_k = [u_1, \ldots, u_k]$ and $\Lambda_k = \operatorname{diag}(\lambda_1, \ldots, \lambda_k)$. Under the boundary eigengap assumption, the unique optimal matrix is $M_k = U_k \Lambda_k U_k^\top$, and the minimizing factors form the manifold

$$\mathcal{X}_\star := \left\{ U_k \Lambda_k^{1/2} R : R \in O(k) \right\}. \tag{3}$$

The gradient is $\nabla f(X) = (XX^\top - M)X$. We study the preconditioned gradient iteration

$$X_{t+1} = X_t - \eta \nabla f(X_t)(X_t^\top X_t)^{-1} = (1-\eta)X_t + \eta M X_t (X_t^\top X_t)^{-1}, \tag{4}$$

where $\eta > 0$ is a fixed step size.

For simplicity of exposition, we use a normalized randomized range initialization. Let $G \in \mathbb{R}^{n\times k}$ have an absolutely continuous distribution, such as the standard Gaussian. Fix $\beta > 1$, and independently draw $S_{\text{init}} \in \mathbb{S}^k_{++}$ from a distribution that is absolutely continuous on $\{S \in \mathbb{S}^k_{++} : I_k \prec S^2 \prec \beta I_k\}$. Set

$$X_0 := MG(G^\top MG)^{-1/2} S_{\text{init}}. \tag{5}$$

Because $d \geq k$, the matrix $G^\top MG$ is positive definite almost surely. Thus the initialization is well-defined, satisfies $\operatorname{range}(X_0) \subseteq \operatorname{range}(M)$ and $I_k \prec X_0^\top M^\dagger X_0 = S_{\text{init}}^2 \prec \beta I_k$, and has the same column space as MG. The additional randomization by S_{init} gives a density in the transformed coordinates used for saddle avoidance. Initialization requires only MG and a $k \times k$ inverse square root; the eigendecomposition is used only in the analysis.

→ GLOBAL CONVERGENCE OF THE ITERATES.

Theorem 3.1 (Global convergence and linear rate for the iterates)**.** *Under the assumptions and initialization above, fix $0 < \eta < 1$. Almost surely, every iterate of (4) has full column rank, and there is a random $X_\star \in \mathcal{X}_\star$ such that*

$$X_t \xrightarrow[t\to\infty]{} X_\star, \qquad X_t X_t^\top \xrightarrow[t\to\infty]{} M_k. \tag{6}$$

Define

$$\rho_\eta := \left(1 - \frac{\eta(1-\eta)}{2}\frac{\lambda_k - \lambda_{k+1}}{\lambda_k}\right)^{1/2} \in (0,1). \tag{7}$$

There are almost-surely finite random constants $C_X, C_M > 0$ such that, for every $t \geq 0$,

$$\|X_t - X_\star\|_F \leq C_X \rho_\eta^t, \qquad \|X_t X_t^\top - M_k\|_F \leq C_M \rho_\eta^t. \tag{8}$$

The admissible step-size interval is independent of the spectrum. The rate factor depends on the relative boundary eigengap, and the random constants absorb the initial phase. These constants are not quantified here. §5 develops the proof; Appendix B.7 gives the passage from the intermediate estimate to (8).

4 THE HIDDEN GRADIENT STRUCTURE OF SCALEDGD

→ HIDDEN GRADIENT DYNAMICS. Define the transformed coordinates by

$$Z_t := \Lambda^{-1/2} U^\top X_t \in \mathbb{R}^{d\times k}. \tag{9}$$

The iteration preserves $\operatorname{range}(M)$, so $X_t = U\Lambda^{1/2} Z_t$. In these coordinates, $I_k \prec Z_0^\top Z_0 \prec \beta I_k$ and the iteration becomes

$$Z_{t+1} = (1-\eta)Z_t + \eta \Lambda Z_t (Z_t^\top \Lambda Z_t)^{-1}. \tag{10}$$

The transformed iteration is exactly ordinary gradient descent, with the same step size η, on the hidden objective

$$h(Z) := \frac{1}{2}\|Z\|_F^2 - \frac{1}{2}\log(\det(Z^\top \Lambda Z)). \tag{11}$$

Indeed, $\nabla h(Z) = Z - \Lambda Z(Z^\top \Lambda Z)^{-1}$, so $Z_{t+1} = Z_t - \eta \nabla h(Z_t)$. The gradient of h is not globally Lipschitz on the full-column-rank domain, so this equivalence alone does not justify convergence throughout $0 < \eta < 1$. The analysis below controls the subspace and scale separately, without requiring global descent of h.

The minimizers of h form the set $\mathcal{Z}_\star := \{E_k R : R \in O(k)\}$, where $E_k = [e_1, \ldots, e_k] \in \mathbb{R}^{d\times k}$. The transformation maps the minimizers of h exactly onto those of f: $\mathcal{X}_\star = U\Lambda^{1/2}\mathcal{Z}_\star$, since $U\Lambda^{1/2}E_k = U_k\Lambda_k^{1/2}$. For any $Z_\star \in \mathcal{Z}_\star$,

$$h_\star := h(Z_\star) = \frac{k}{2} - \frac{1}{2}\log(\det \Lambda_k). \tag{12}$$

→ RADIAL DECOMPOSITION. For a full-column-rank Z, consider its polar decomposition $Z = QW^{1/2}$, where $W = Z^\top Z \succ 0$ and $Q^\top Q = I_k$. The hidden objective separates into radial and subspace errors:

$$h(Z) - h_\star = \underbrace{\frac{1}{2}\big(\operatorname{tr} W - \log(\det W) - k\big)}_{\text{radial error}} + \underbrace{\frac{1}{2}\log\left(\frac{\det \Lambda_k}{\det(Q^\top \Lambda Q)}\right)}_{\text{subspace error}}. \tag{13}$$

The first term is nonnegative and vanishes exactly at $W = I_k$; the second is nonnegative and vanishes exactly when Q spans the leading subspace (Lemma A.6). This separates the scale of the factors from the subspace they represent.

→ WELL-POSEDNESS. Recall that $W_t = Z_t^\top Z_t$ and $Q_t = Z_t W_t^{-1/2}$. Substituting $Z_t = Q_t W_t^{1/2}$ into (10) and using $Z_{t+1} = Z_t - \eta\nabla h(Z_t)$ gives

$$\begin{aligned}\nabla h(Z_t) W_t^{1/2} &= Q_t(W_t - I_k) - (I_d - Q_t Q_t^\top)\Lambda Q_t (Q_t^\top \Lambda Q_t)^{-1} \\ &= Q_t(W_t - I_k) - R_t H_t^{-1},\end{aligned} \tag{14}$$

where we define $H_t := Q_t^\top \Lambda Q_t$ and $R_t := (I_d - Q_t Q_t^\top)\Lambda Q_t = \Lambda Q_t - Q_t H_t$. Since $Q_t Q_t^\top$ is the orthogonal projector onto $\operatorname{range}(Q_t)$, R_t is the component of ΛQ_t orthogonal to the current subspace. It vanishes exactly when this subspace is invariant under Λ. Combining (14) with the gradient-descent update and using $Q_t^\top R_t = 0$ yields the exact recursion

$$W_{t+1} = \Big(\underbrace{(1-\eta)W_t^{1/2} + \eta W_t^{-1/2}}_{\text{weighted Heron update}}\Big)^2 + \underbrace{\eta^2 W_t^{-1/2}(R_t H_t^{-1})^\top (R_t H_t^{-1}) W_t^{-1/2}}_{\text{subspace residual correction}}. \tag{15}$$

The first term is the square of a weighted Heron update on $W_t^{1/2}$; at $\eta = 1/2$, this is the classical map $S \mapsto \frac{1}{2}(S + S^{-1})$ for the square root of I_k. The second term is a positive-semidefinite correction that vanishes when the column space of Q_t is invariant under Λ. This structure yields the following uniform bounds.

Lemma 4.1 (Uniform matrix bounds and well-posedness)**.** *Under the assumptions and initialization of §3, fix $0 < \eta < 1$ and define*

$$\underline{m}_\eta := \begin{cases} 1, & \eta \le 1/2, \\ 4\eta(1-\eta), & \eta > 1/2, \end{cases} \qquad \overline{m}_{\eta,\beta} := \max\left\{\beta, \frac{2(1-\eta) + \eta\kappa/\underline{m}_\eta}{2-\eta}\right\}. \tag{16}$$

Then almost surely every iterate of (10) is full column rank and, for all $t \ge 0$,

$$\underline{m}_\eta I_k \preceq Z_t^\top Z_t \preceq \overline{m}_{\eta,\beta} I_k. \tag{17}$$

In particular, $X_t^\top X_t \succeq \underline{m}_\eta \lambda_d I_k$ and every update is well-defined.

The proof is given in Lemma B.1.

5 GLOBAL CONVERGENCE ANALYSIS

The global convergence proof has two main steps. First, determinant ascent and radial stabilization yield finite total movement and convergence to a critical point. This step does not rely on descent of h. Second, a center-stable argument excludes convergence to nonminimizing critical points almost surely. Local geometry at the limiting minimizer then gives the local linear rate. The complete proof is in §B. Throughout, assume the setting and initialization of §3 and fix $0 < \eta < 1$.

5.1 GLOBAL CONVERGENCE TO A CRITICAL POINT

We retain the notation H_t and R_t from §4.

Lemma 5.1 (Determinant ascent and angular summability)**.** *There is a deterministic constant $c_{\mathrm{ang}} > 0$, depending only on η, β, and the spectrum, such that almost surely, for every $t \geq 0$,*

$$\log\left(\frac{\det(Q_{t+1}^\top \Lambda Q_{t+1})}{\det(Q_t^\top \Lambda Q_t)}\right) \geq c_{\mathrm{ang}} \|R_t H_t^{-1}\|_F^2. \tag{18}$$

Consequently,

$$\sum_{t=0}^{\infty} \|R_t H_t^{-1}\|_F^2 < \infty. \tag{19}$$

The determinant increases strictly whenever $R_t \neq 0$, and its upper bound $\det(Q_t^\top \Lambda Q_t) \leq \det \Lambda_k$ gives summability. The proof and an explicit value of c_{ang} are in Lemma B.2 and Corollary B.3.

Angular summability also yields radial stabilization.

Lemma 5.2 (Radial contraction and summability)**.** *There is a deterministic constant $\rho_{\eta,\beta} \in (0,1)$, depending only on η, β, and the spectrum, such that almost surely, for every $t \geq 0$,*

$$\|\log(W_{t+1})\|_F \leq \rho_{\eta,\beta} \|\log(W_t)\|_F + \frac{\eta^2}{\underline{m}_\eta^2} \|R_t H_t^{-1}\|_F^2. \tag{20}$$

Consequently,

$$\sum_{t=0}^{\infty} \|W_t - I_k\|_F < \infty. \tag{21}$$

The weighted Heron update contracts the radial log-error up to the summable residual correction from Lemma 5.1; Lemma B.4 gives the proof and an explicit $\rho_{\eta,\beta}$. In particular, $W_t \xrightarrow[t\to\infty]{} I_k$; angular residual decay and (14) then give $\nabla h(Z_t) \xrightarrow[t\to\infty]{} 0$.

To prove convergence to one point, we establish summable update lengths. The gradient identity (14) and the lower bound (17) give

$$\|Z_{t+1} - Z_t\|_F = \eta \|\nabla h(Z_t)\|_F \leq \frac{\eta}{\sqrt{\underline{m}_\eta}} \left(\|W_t - I_k\|_F + \|R_t H_t^{-1}\|_F\right). \tag{22}$$

The radial errors are summable by Lemma 5.2. It remains to strengthen (19) from summability of squared angular residual norms to summability of the norms.

Proposition 5.3 (Convergence to one critical point)**.** *Almost surely,*

$$\sum_{t=0}^{\infty} \|R_t H_t^{-1}\|_F < \infty, \qquad \sum_{t=0}^{\infty} \|Z_{t+1} - Z_t\|_F < \infty. \tag{23}$$

Consequently, $Z_t \xrightarrow[t\to\infty]{} Z_\star$ for some $Z_\star$ satisfying $\nabla h(Z_\star) = 0$.

The additional ingredient is the multi-level determinant inequality in Lemma B.5. Combined with determinant ascent, it gives eventual geometric decay of the gap between $\log(\det(Q_t^\top \Lambda Q_t))$ and its limiting critical value. The ascent bound then makes the angular residual norms summable; Lemma B.6 gives the details.

5.2 Saddle avoidance and local linear convergence

Proposition 5.4 (Almost-sure saddle avoidance)**.** *The probability that Z_t converges to a critical point outside $\mathcal{Z}_\star$ is zero.*

Every nonleading critical manifold has an unstable angular direction. Absolute continuity of the initialization and a center-stable argument exclude convergence to these manifolds almost surely; see Lemma B.8.

Together, Propositions 5.3 and 5.4 give $Z_t \xrightarrow[t\to\infty]{} Z_\star \in \mathcal{Z}_\star$ almost surely. The local PL inequality and sufficient decrease near this minimizer give, for an almost-surely finite T and ρ_η from (7),

$$h(Z_{t+1}) - h_\star \le \rho_\eta^2\big(h(Z_t) - h_\star\big), \qquad t \ge T. \tag{24}$$

Summing the resulting geometric bound on update lengths yields the iterate and reconstruction rates in Theorem 3.1. Theorem B.9 and §B.7 give the details. The entrance time T is not quantified.

6 Extension to asymmetric matrices

The rectangular extension retains almost-sure convergence to an optimal approximation, but its negative dilation eigenvalues restrict the step size. Let $M \in \mathbb{R}^{m\times n}$ have compact SVD $U\Sigma V^\top$, rank $r \ge k$, and $\sigma_k > \sigma_{k+1}$, with $\sigma_{r+1} := 0$. Write $M_k = U_k\Sigma_k V_k^\top$ and $\kappa = \sigma_1/\sigma_r$.

→ Problem and Update. For $\mathcal{H}(M) = \left[\begin{smallmatrix} 0 & M \\ M^\top & 0 \end{smallmatrix}\right]$ and $Z = (Y^\top, X^\top)^\top$, consider

$$g(Z) = \frac{1}{4}\|\mathcal{H}(M) - ZZ^\top\|_F^2. \tag{25}$$

This objective equals $\frac{1}{4}\|M - 2YX^\top\|_F^2 + \frac{1}{4}\|X^\top X - Y^\top Y\|_F^2 + \frac{1}{4}\|M\|_F^2$; its minimizers satisfy $2Y_\star X_\star^\top = M_k$. ScaledGD on this dilation gives the shared-inverse update

$$\begin{aligned} X_{t+1} &= (1-\eta)X_t + \eta M^\top Y_t G_t^{-1}, \\ Y_{t+1} &= (1-\eta)Y_t + \eta M X_t G_t^{-1}, \end{aligned} \qquad G_t = X_t^\top X_t + Y_t^\top Y_t. \tag{26}$$

Each step uses products with $M, M^\top$ and a $k \times k$ solve; the dilation need not be formed.

→ Initialization and Convergence. Draw a standard Gaussian $\Omega \in \mathbb{R}^{m\times k}$ and, independently, $A_X \in \mathbb{R}^{m\times k}$ and $A_Y \in \mathbb{R}^{n\times k}$ with densities supported in their unit Frobenius balls. Set

$$Q = M^\top\Omega(\Omega^\top MM^\top\Omega)^{-1/2}, \qquad X_0 = Q + \frac{M^\top A_X}{4\|M\|_F}, \qquad Y_0 = M\left(Q + \frac{1}{4}A_Y\right). \tag{27}$$

This construction requires no singular values or singular vectors. Lemma C.4 verifies its density and boundedness.

Theorem 6.1 (Asymmetric convergence and linear rate)**.** *With the initialization above, fix $0 < \eta < \eta_{\max}$, where $\eta_{\max}$ is given in (169) with $\mu = 1/2$ and the deterministic R in (189). Almost surely, every update is well-defined and X_t, Y_t converge to limits $X_\star, Y_\star$ satisfying $2Y_\star X_\star^\top = M_k$. Define*

$$q_\eta := 1 - \frac{\eta}{2}\left(1 - \frac{\sigma_{k+1}}{\sigma_k}\right) \in (0,1). \tag{28}$$

There are almost-surely finite random constants $C_{XY}, C_M > 0$ such that, for every $t \ge 0$,

$$\|X_t - X_\star\|_F + \|Y_t - Y_\star\|_F \le C_{XY}q_\eta^t, \qquad \|2Y_tX_t^\top - M_k\|_F \le C_M q_\eta^t. \tag{29}$$

Theorem C.1 gives the general initialization conditions and explicit sufficient step-size bound. In signed coordinates, the negative component decays geometrically and the positive component approaches the symmetric update. Appendix C establishes convergence to a critical point, saddle avoidance, and the rate.

→ Why the Step-Size Range Shrinks. At a minimizer, negative-eigenspace perturbations have multipliers $1 - \eta(1 + \sigma_j/\sigma_i)$ for $i \le k$, $j \le r$. All normal directions are strictly contracting precisely when

$$0 < \eta < \frac{2}{1 + \sigma_1/\sigma_k}. \tag{30}$$

Thus $\sigma_1 > \sigma_k$ creates local instability for some $\eta \in (0,1)$ (Lemma C.8). The sufficient global bound is stronger. For $k = 1$, the threshold equals 1, and the next section obtains global convergence throughout this interval under a nonzero projection condition.

7 Rank-one approximation: explicit and gap-free rates

For rank-one approximation, we obtain explicit geometric rates from initialization when the leading eigenvalue or singular value is separated, and gap-free bounds on excess reconstruction error otherwise. Both guarantees hold for every $0 < \eta < 1$ in the symmetric and asymmetric settings.

→ Symmetric Matrices. Consider rank-one ScaledGD for $M \succ 0$ in dimension $d \geq 2$. Write $x_t = U\Lambda^{1/2} z_t$, and let $P_\star$ project orthogonally onto the leading eigenspace of M. Use the ScaledGD initialization of §3 with $k = 1$ and $\beta > 1$.

Theorem 7.1 (Symmetric rank-one convergence and rates)**.** *With this initialization, fix $0 < \eta < 1$ and assume $P_\star x_0 \neq 0$. Then the iterates converge to a limit $x_\star$ such that $x_\star x_\star^\top$ is a best rank-one approximation of M.*

(i) Gap-dependent rate. If $\lambda_1 > \lambda_2$, let $\rho \in (0,1)$ be the explicit factor in (219), whose improvement $1-\rho$ is proportional to $(\lambda_1 - \lambda_2)/\lambda_1$. There are finite constants $K_X, K_M > 0$ such that, for every $t \geq 0$,

$$\|x_t - x_\star\| \leq K_X \rho^t, \qquad \|x_t x_t^\top - x_\star x_\star^\top\|_F \leq K_M \rho^t. \tag{31}$$

(ii) Gap-free rate. For $\lambda_1 \geq \lambda_2$, there are finite $C_1, C_2 \geq 0$ such that, for every $t \geq 0$,

$$0 \leq \|M - x_t x_t^\top\|_F^2 - \|M - x_\star x_\star^\top\|_F^2 \leq \frac{C_1}{t+1} + \frac{C_2}{(t+1)^2}. \tag{32}$$

The constants K_X, K_M, C_1, C_2 are given explicitly in the proofs in §D.1 and §D.2 and depend only on η, β, M and x_0. In particular, C_1, C_2 depend only on $\eta, \lambda_1, \|x_0\|$ and $\|P_\star x_0\|$.

Unlike the general-rank bound, these constants are explicit from initialization. The proof combines contraction of $\sum_{i=2}^d z_{t,i}^2 / z_{t,1}^2$ with the recursion for $\log(\|z_t\|^2)$. For the gap-free bound, eigen-coordinate errors are weighted by their spectral differences from the leading eigenvalue.

→ Asymmetric Matrices. Consider rank-one ScaledGD for $M \neq 0$, using the rectangular update and singular-vector notation of §6; set $\sigma_2 := 0$ if $r = 1$. Write $z_t = (y_t^\top, x_t^\top)^\top$ and let $P_\star$ project orthogonally onto the leading positive eigenspace of $\mathcal{H}(M)$. The initialization may be arbitrary, subject to the nonzero projection condition below, which holds almost surely for any law with a density, including independent Gaussian initial iterates.

Theorem 7.2 (Asymmetric rank-one convergence and rates)**.** *Fix $0 < \eta < 1$ and assume $P_\star z_0 \neq 0$. Then every iterate is well-defined, and x_t, y_t converge to limits $x_\star, y_\star$ such that $2y_\star x_\star^\top$ is a best rank-one approximation of M.*

(i) Gap-dependent rate. If $\sigma_1 > \sigma_2$, let $\rho \in (0,1)$ be the explicit factor in (272), with $1-\rho$ proportional to $(\sigma_1 - \sigma_2)/\sigma_1$. There are finite constants $K, K_M \geq 0$ such that, for every $t \geq 0$,

$$\left(\|x_t - x_\star\|^2 + \|y_t - y_\star\|^2\right)^{1/2} \leq K\rho^t, \qquad \|2y_t x_t^\top - 2y_\star x_\star^\top\|_F \leq K_M \rho^t. \tag{33}$$

(ii) Gap-free rate. For $\sigma_1 \geq \sigma_2$, there are finite $C_1, C_2 \geq 0$ such that, for every $t \geq 0$,

$$0 \leq \|M - 2y_t x_t^\top\|_F^2 - \|M - 2y_\star x_\star^\top\|_F^2 \leq \frac{C_1}{t+1} + \frac{C_2}{(t+1)^2}. \tag{34}$$

The constants K, K_M, C_1, C_2 are given explicitly in the proofs in §D.3 and §D.4 and depend only on η, M, x_0 and y_0. In particular, C_1, C_2 depend only on $\eta, \sigma_1, \|z_0\|$ and $\|P_\star z_0\|$.

For rank one, all other dilation eigen-coordinates contract relative to the leading positive coordinate for every $0 < \eta < 1$, including the negative ones. Combining this with the Heron recursion gives explicit rates without the initialization restrictions of Theorem 6.1. Thus the rank-one result attains the full local stability interval in (30).

8 NUMERICAL ILLUSTRATIONS

Figure 1 illustrates three mechanisms using independent standard Gaussian entries for the initial iterates, without the theorem-specific normalization. Appendix F gives the protocols.

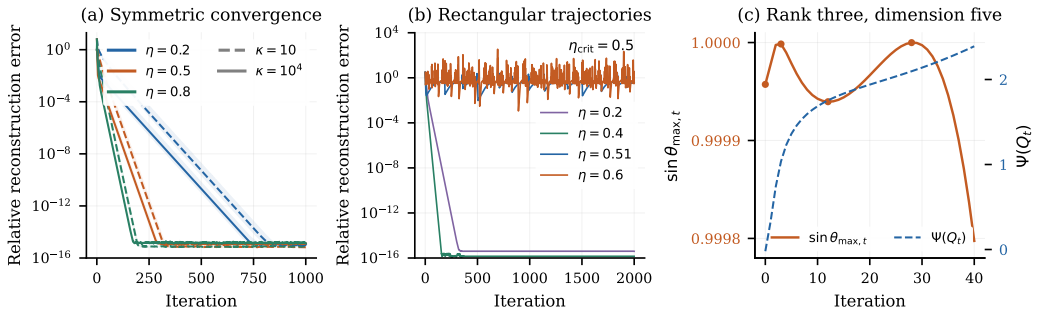

Figure 1: **Dynamics of ScaledGD.** (a) Symmetric rank-4 approximation in dimension 50, with relative gap 0.2: median reconstruction errors and min–max bands over six seeds, shown down to 10^{-16}. (b) Rectangular rank-2 approximation with the same Gaussian initialization across step sizes. The threshold in (30) is $1/2$; the larger step sizes show persistent reconstruction errors. (c) A selected rank-3 example in dimension 5: small reversals of $\sin\theta_{\max,t}$ near one (zoomed left axis), while the log-determinant potential $\Psi(Q_t) := \log\det(Q_t^\top \Lambda Q_t)$ increases (right axis).

→ CONVERGENCE AND STEP SIZES. For $\kappa \in \{10, 10^4\}$ and $\eta \in \{0.2, 0.5, 0.8\}$, all tested symmetric trajectories reach relative reconstruction error below 2×10^{-15} within 1000 iterations, then fluctuate near the numerical precision floor. For a rectangular matrix with singular values $(6, 2, 1)$ and $k = 2$, step sizes $\eta \in \{0.2, 0.4\}$ reach numerical precision, whereas $\eta \in \{0.51, 0.6\}$ produce persistent errors and excursions over 2000 iterations. The same contrast holds across six tested initializations. This is numerical evidence of failure above the local stability threshold, not a proof of nonconvergence for every initialization.

→ ANGLE REVERSALS AND DETERMINANT ASCENT. For $k = 3$ and $n = d = 5$, we select a Gaussian-initialized trajectory whose sine of the largest principal angle between $\mathrm{range}(Q_t)$ and the leading eigenspace increases, decreases, then increases again, although $\Psi(Q_t)$ ascends. The reversals are of order 10^{-5} in the sine, near one; the axis is zoomed to show them. An independent 80-digit replay confirms these small changes. This illustrates failure of monotonicity, not large or typical oscillations.

9 CONCLUSION

We establish almost-sure convergence of the symmetric ScaledGD iterates to a single global minimizer for every fixed $0 < \eta < 1$, under a boundary eigengap and absolutely continuous random initialization within $\mathrm{range}(M)$. The hidden gradient representation exposes the coupled evolution of the subspace and the scale of the iterates. Determinant ascent and a matrix Heron recursion control this coupling and yield finite trajectory length, even though the argument does not require global descent of the hidden objective. The admissible step-size interval is independent of the spectrum, while the asymptotic linear rate depends on the relative eigengap.

The analysis of the dilation-based rectangular extension yields a spectrum-dependent step-size bound; rank one admits explicit geometric and gap-free reconstruction bounds in both settings. The main quantitative challenge is to bound the general-rank entrance time. Further directions include sharpening the rectangular sufficient step-size bound, extending the initialization guarantees beyond the target range, and establishing accelerated rates.

AI USE STATEMENT

We used generative AI tools to assist with developing and writing proofs, reviewing related literature, implementing and analyzing numerical experiments, generating figures, and editing the manuscript for clarity. The authors reviewed the AI-assisted proofs, checked their relation to existing literature, and verified and tested the experimental code. The authors take responsibility for the final text, mathematical claims, proofs, code, and numerical results.

ETHICS STATEMENT

We do not anticipate ethical concerns arising from the theoretical results and synthetic numerical experiments presented in this paper.

REPRODUCIBILITY STATEMENT

The assumptions are stated alongside each result, with detailed proofs in Appendices A–D. For the numerical illustrations in Figure 1, Appendix F specifies the matrix constructions, initializations, step sizes, random seeds, iteration counts, and error metrics, together with the high-precision verification of the angle-reversal example.

REFERENCES

Hédy Attouch, Jérôme Bolte, and Benar Fux Svaiter. Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward-backward splitting, and regularized Gauss–Seidel methods. *Mathematical Programming*, 137(1–2):91–129, 2013. doi: 10.1007/s10107-011-0484-9.

Hengchao Chen, Xin Chen, Mohamad Elmasri, and Qiang Sun. Gradient descent in matrix factorization: Understanding large initialization. In *Proceedings of the Fortieth Conference on Uncertainty in Artificial Intelligence*, volume 244 of *Proceedings of Machine Learning Research*, pp. 619–647. PMLR, 2024. URL https://proceedings.mlr.press/v244/chen24a.html.

Patrick Cheridito, Arnulf Jentzen, and Florian Rossmannek. Gradient descent provably escapes saddle points in the training of shallow ReLU networks. *Journal of Optimization Theory and Applications*, 203(3):2617–2648, 2024. doi: 10.1007/s10957-024-02513-3.

Simon S Du, Wei Hu, and Jason D Lee. Algorithmic regularization in learning deep homogeneous models: Layers are automatically balanced. *Advances in neural information processing systems*, 31, 2018.

Rong Ge, Chi Jin, and Yi Zheng. No spurious local minima in nonconvex low rank problems: A unified geometric analysis. In *International Conference on Machine Learning*, pp. 1233–1242. PMLR, 2017.

Suriya Gunasekar, Blake E. Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://papers.nips.cc/paper_files/paper/2017/hash/58191d2a914c6dae66371c9dcdc91b41-Abstract.html.

Yassir Jedra and Devavrat Shah. k-SVD with gradient descent. *arXiv preprint arXiv:2502.00320*, 2025.

Xixi Jia, Hailin Wang, Jiangjun Peng, Xiangchu Feng, and Deyu Meng. Preconditioning matters: Fast global convergence of non-convex matrix factorization via scaled gradient descent. In *Advances in Neural Information Processing Systems*, volume 36, pp. 76202–76213, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/f02f1185b97518ab5bd7ebde466992d3-Abstract-Conference.html.

Chi Jin, Sham M. Kakade, and Praneeth Netrapalli. Provable efficient online matrix completion via non-convex stochastic gradient descent. In *Advances in Neural Information Processing Systems*, volume 29, 2016. URL https://proceedings.neurips.cc/paper/2016/hash/38651c4450f87348fcbe1f992746a954-Abstract.html.

Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M Kakade, and Michael I Jordan. How to escape saddle points efficiently. In *International conference on machine learning*, pp. 1724–1732. PMLR, 2017.

Jikai Jin, Zhiyuan Li, Kaifeng Lyu, Simon Shaolei Du, and Jason D. Lee. Understanding incremental learning of gradient descent: A fine-grained analysis of matrix sensing. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 15200–15238. PMLR, 2023. URL https://proceedings.mlr.press/v202/jin23a.html.

Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. Gradient descent only converges to minimizers. In *Proceedings of the 29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pp. 1246–1257, 2016. URL https://proceedings.mlr.press/v49/lee16.html.

Bingcong Li, Liang Zhang, Aryan Mokhtari, and Niao He. On the crucial role of initialization for matrix factorization. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=YTEwJaBdh0.

Cong Ma, Kaizheng Wang, Yuejie Chi, and Yuxin Chen. Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion, and blind deconvolution. *Foundations of Computational Mathematics*, 20(3):451–632, 2020. doi: 10.1007/s10208-019-09429-9.

Boris S. Mityagin. The zero set of a real analytic function. *Mathematical Notes*, 107:529–530, 2020. doi: 10.1134/S0001434620030189.

Dominik Stöger and Mahdi Soltanolkotabi. Small random initialization is akin to spectral learning: Optimization and generalization guarantees for overparameterized low-rank matrix reconstruction. *Advances in Neural Information Processing Systems*, 34:23831–23843, 2021.

Tian Tong, Cong Ma, and Yuejie Chi. Accelerating ill-conditioned low-rank matrix estimation via scaled gradient descent. *Journal of Machine Learning Research*, 22(150):1–63, 2021. URL https://jmlr.org/papers/v22/20-1067.html.

Stephen Tu, Ross Boczar, Max Simchowitz, Mahdi Soltanolkotabi, and Ben Recht. Low-rank solutions of linear matrix equations via procrustes flow. In *International Conference on Machine Learning*, pp. 964–973. PMLR, 2016.

Xingyu Xu, Yandi Shen, Yuejie Chi, and Cong Ma. The power of preconditioning in overparameterized low-rank matrix sensing. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 38611–38654. PMLR, 2023. URL https://proceedings.mlr.press/v202/xu23o.html.

CONTENTS

1 Introduction **1**

2 Related work **2**

3 Setting and Main Result **3**

4 The hidden gradient structure of ScaledGD **4**

5 Global convergence analysis **6**

5.1 Global convergence to a critical point 6

5.2 Saddle avoidance and local linear convergence 7

6 Extension to asymmetric matrices **7**

7 Rank-one approximation: explicit and gap-free rates **8**

8 Numerical illustrations **9**

9 Conclusion **9**

A Shared symmetric identities and geometry **14**

A.1 Transformed iteration and initialization 14

A.2 Polar decomposition and matrix inequalities 15

A.3 Objective decomposition and curvature 15

A.4 Critical points and local convergence geometry 16

B Proofs :: Rank-k approximation for symmetric matrices **19**

B.1 Proof of Lemma 4.1 19

B.2 Proof of Lemma 5.1 20

B.3 Proof of Lemma 5.2 21

B.4 Proof of Proposition 5.3 22

B.5 Proof of Proposition 5.4 24

B.6 Objective contraction near the limiting minimizer 25

B.7 Proof of Theorem 3.1 26

B.8 Removing the normalization of the initialization 27

C Proofs :: Rank-k approximation for asymmetric matrices **28**

C.1 Decay of the negative component 28

C.2 Proof of Theorem 6.1: well-posedness and initialization 29

C.3 Proof of Theorem 6.1: convergence to a critical point 30

C.4 Proof of Theorem 6.1: saddle avoidance 32

C.5 Proof of Theorem 6.1: local rate and conclusion 32

D Proofs :: Rank-1 approximation **34**

D.1 Proof of Theorem 7.1(i): gap-dependent rate 34

D.2 Proof of Theorem 7.1(ii): gap-free rate 36

D.3 Proof of Theorem 7.2(i): gap-dependent rate 39

D.4 Proof of Theorem 7.2(ii): gap-free rate 41

E Further comparisons with existing convergence guarantees **42**

E.1 General convergence theory and the ScaledGD trajectory 42

E.2 Interpreting the closest approximation guarantees 42

F Numerical protocols and verification **43**

F.1 Symmetric convergence 43

F.2 Rectangular trajectories across the stability threshold 43

F.3 Multiple angle reversals for rank three in dimension five 43

A SHARED SYMMETRIC IDENTITIES AND GEOMETRY

This appendix collects the identities and geometry used in the convergence proofs. The transformed iteration and polar decomposition underpin the analysis in §B; the curvature estimates and critical-point geometry support saddle avoidance and local convergence rates.

A.1 TRANSFORMED ITERATION AND INITIALIZATION

We justify the change of variables in §4, establish the initialization density needed for saddle avoidance, and identify the transformed update as gradient descent on h.

Lemma A.1 (Whitened iteration)**.** *Almost surely, X_0 has full column rank,* $\operatorname{range}(X_0) \subseteq \operatorname{range}(M)$*, and*

$$X_0^\top M^\dagger X_0 = S_{\text{init}}^2. \tag{35}$$

As long as the iterates are well-defined, $\operatorname{range}(X_t) \subseteq \operatorname{range}(M)$*. Writing*

$$X_t = U\Lambda^{1/2} Z_t, \tag{36}$$

iteration (4) *is equivalent to*

$$Z_{t+1} = (1-\eta)Z_t + \eta \Lambda Z_t (Z_t^\top \Lambda Z_t)^{-1}, \tag{37}$$

and $Z_0^\top Z_0 = S_{\text{init}}^2$.

Proof. Since $d \geq k$ and G is absolutely continuous, MG has full column rank almost surely. The stated normalization gives

$$X_0^\top M^\dagger X_0 = S_{\text{init}}^2 \tag{38}$$

and preserves $\operatorname{range}(X_0) \subseteq \operatorname{range}(M)$. If $\operatorname{range}(X_t) \subseteq \operatorname{range}(M)$, the columns of both terms in (4) lie in $\operatorname{range}(M)$, proving invariance. Substitution of $X_t = U\Lambda^{1/2} Z_t$ into (4), followed by left multiplication by $\Lambda^{-1/2} U^\top$, gives (37). Finally,

$$X_t^\top M^\dagger X_t = Z_t^\top Z_t, \tag{39}$$

so the initialization yields $Z_0^\top Z_0 = S_{\text{init}}^2$. □

Lemma A.2 (Absolute continuity of the whitened initialization)**.** *The random matrix Z_0 is absolutely continuous with respect to Lebesgue measure on $\mathbb{R}^{d\times k}$.*

Proof. Set $Y = \Lambda^{1/2} U^\top G$. Since G is absolutely continuous and the map $G \mapsto Y$ is surjective, Y is absolutely continuous on $\mathbb{R}^{d\times k}$. The polar-coordinate map

$$\operatorname{St}(d,k) \times \mathbb{S}_{++}^k \longrightarrow \{Z \in \mathbb{R}^{d\times k} : \operatorname{rank} Z = k\}, \qquad (Q,S) \longmapsto QS, \tag{40}$$

is a smooth diffeomorphism, with inverse given by the polar decomposition. The change-of-variables formula in local coordinates therefore gives a joint density for the polar factors of Y with respect to the product of Stiefel volume and Lebesgue measure on symmetric matrices. Taking the marginal shows that $Q_0 = Y(Y^\top Y)^{-1/2}$ has a density on $\operatorname{St}(d,k)$. In whitened coordinates, $Z_0 = Q_0 S_{\text{init}}$. Independence and the assumed density of S_{init} give a joint density for (Q_0, S_{init}); applying the same diffeomorphism proves that Z_0 has a Lebesgue density. □

Lemma A.3 (Gradient representation)**.** *On the full-rank domain,*

$$\nabla h(Z) = Z - \Lambda Z (Z^\top \Lambda Z)^{-1}, \tag{41}$$

and (37) *is precisely*

$$Z_{t+1} = Z_t - \eta \nabla h(Z_t). \tag{42}$$

Proof. The differential identity

$$d\big[\log(\det(Z^\top \Lambda Z))\big][E] = 2\operatorname{tr}\big((Z^\top \Lambda Z)^{-1} Z^\top \Lambda E\big) \tag{43}$$

gives the gradient formula. The update identity follows immediately. □

For the remainder of this appendix, write

$$W_t := Z_t^\top Z_t, \qquad S_t := Z_t^\top \Lambda Z_t, \qquad K_t := Z_t^\top \Lambda^2 Z_t. \tag{44}$$

A.2 Polar decomposition and matrix inequalities

The polar decomposition gives a recursion for $W_t = Z_t^\top Z_t$ with a positive semidefinite correction determined by $R_t H_t^{-1}$. These identities yield the uniform matrix bounds used in §B. Write

$$Z_t = Q_t W_t^{1/2}, \qquad Q_t^\top Q_t = I_k, \tag{45}$$

and define

$$H_t := Q_t^\top \Lambda Q_t, \qquad R_t := (I_d - Q_t Q_t^\top)\Lambda Q_t. \tag{46}$$

Lemma A.4 (Recursion for $Z_t^\top Z_t$)**.** *Whenever Z_t has full column rank, its update satisfies*

$$\begin{aligned} W_{t+1} = (1-\eta)^2 W_t + 2\eta(1-\eta) I_k + \eta^2 W_t^{-1} \\ + \eta^2 W_t^{-1/2} H_t^{-1} R_t^\top R_t H_t^{-1} W_t^{-1/2}. \end{aligned} \tag{47}$$

In particular,

$$S_t^{-1} K_t S_t^{-1} = W_t^{-1} + W_t^{-1/2} H_t^{-1} R_t^\top R_t H_t^{-1} W_t^{-1/2} \succeq W_t^{-1}. \tag{48}$$

Proof. Since $S_t = W_t^{1/2} H_t W_t^{1/2}$ and $\Lambda Q_t = Q_t H_t + R_t$, the whitened update becomes

$$Z_{t+1} = Q_t\big((1-\eta) W_t^{1/2} + \eta W_t^{-1/2}\big) + \eta R_t H_t^{-1} W_t^{-1/2}. \tag{49}$$

Since $Q_t^\top R_t = 0$, the cross terms vanish and

$$W_{t+1} = Z_{t+1}^\top Z_{t+1} = \big((1-\eta) W_t^{1/2} + \eta W_t^{-1/2}\big)^2 + \eta^2 W_t^{-1/2} H_t^{-1} R_t^\top R_t H_t^{-1} W_t^{-1/2}. \tag{50}$$

Expanding the square gives (47). Comparing it with the direct expansion

$$W_{t+1} = (1-\eta)^2 W_t + 2\eta(1-\eta) I_k + \eta^2 S_t^{-1} K_t S_t^{-1} \tag{51}$$

gives the second identity. □

Lemma A.5 (Uniform bound on the scaled projection residual)**.** *Let $\Lambda \in \mathbb{R}^{d\times d}$ be symmetric positive definite, with largest eigenvalue λ_1 and smallest eigenvalue $\lambda_d > 0$. Let $Q \in \mathbb{R}^{d\times k}$ satisfy $Q^\top Q = I_k$. Write $\kappa = \lambda_1/\lambda_d$ and define*

$$H := Q^\top \Lambda Q, \qquad R := (I_d - QQ^\top)\Lambda Q. \tag{52}$$

Then

$$0 \preceq (RH^{-1})^\top (RH^{-1}) \preceq (\kappa - 1) I_k, \qquad \|RH^{-1}\|_2 \le \sqrt{\kappa - 1}. \tag{53}$$

Proof. Set $J = Q^\top \Lambda^2 Q$. The projector identity gives $R^\top R = J - H^2$. Since $\Lambda^2 \preceq \lambda_1 \Lambda$ and $H \succeq \lambda_d I_k$,

$$I_k + (RH^{-1})^\top (RH^{-1}) = H^{-1} J H^{-1} \preceq \lambda_1 H^{-1} \preceq \kappa I_k. \tag{54}$$

The product $(RH^{-1})^\top (RH^{-1})$ is positive semidefinite, and its largest eigenvalue is $\|RH^{-1}\|_2^2$, proving both bounds. □

A.3 Objective decomposition and curvature

The objective and gradient identities separate the errors in W and the column space of Q. The upper Hessian bound is used in the saddle-avoidance argument in §B.

Lemma A.6 (Objective and gradient decompositions)**.** *For every full-rank $Z = QW^{1/2}$,*

$$h(Z) - h_\star = \frac{1}{2}\big(\operatorname{tr} W - \log(\det W) - k\big) + \frac{1}{2}\log\left(\frac{\det \Lambda_k}{\det H}\right), \tag{55}$$

$$\|\nabla h(Z) W^{1/2}\|_F^2 = \|W - I_k\|_F^2 + \|RH^{-1}\|_F^2. \tag{56}$$

Proof. Since

$$Z^\top \Lambda Z = W^{1/2} H W^{1/2}, \tag{57}$$

the objective identity follows by taking determinants. For the gradient,

$$\nabla h(Z) W^{1/2} = QW - \Lambda Q H^{-1}. \tag{58}$$

Using $\Lambda Q = QH + R$ gives

$$\nabla h(Z) W^{1/2} = Q(W - I_k) - R H^{-1}. \tag{59}$$

The two terms are orthogonal because $Q^\top R = 0$. □

Lemma A.7 (One-sided smoothness)**.** *For every full-column-rank Z and every $E \in \mathbb{R}^{d\times k}$,*

$$\nabla^2 h(Z)[E,E] \le L(Z)\|E\|_F^2, \qquad L(Z) := 1 + \|S^{-1}\|_2 \big\|S^{-1/2} K S^{-1/2}\big\|_2, \tag{60}$$

where $S = Z^\top \Lambda Z$ and $K = Z^\top \Lambda^2 Z$. Moreover,

$$L(Z) \le 1 + \frac{\lambda_1}{\lambda_{\min}(S)}. \tag{61}$$

In particular, if $W = Z^\top Z \succeq I_k$, then $L(Z) \le 1 + \kappa$. Consequently, along any line segment $[Z, Y]$ in the full-rank domain,

$$h(Y) \le h(Z) + \langle \nabla h(Z), Y - Z\rangle + \frac{L_{[Z,Y]}}{2}\|Y - Z\|_F^2, \qquad L_{[Z,Y]} := \sup_{V\in[Z,Y]} L(V). \tag{62}$$

Proof. **Step 1: Hessian quadratic form.** Let $S = Z^\top \Lambda Z$ and define

$$A = \Lambda^{1/2} E S^{-1/2}, \qquad B = \Lambda^{1/2} Z S^{-1/2}. \tag{63}$$

Then $B^\top B = I_k$, and direct differentiation gives

$$\nabla^2 h(Z)[E,E] = \|E\|_F^2 - \|A\|_F^2 + \frac{1}{2}\|A^\top B + B^\top A\|_F^2. \tag{64}$$

Write $A = BC + A_\perp$, where $C = B^\top A$ and $B^\top A_\perp = 0$. Then

$$\frac{1}{2}\|C + C^\top\|_F^2 \le 2\|C\|_F^2, \tag{65}$$

so

$$\nabla^2 h(Z)[E,E] \le \|E\|_F^2 + \|C\|_F^2. \tag{66}$$

Step 2: pointwise upper bound. Since

$$C = S^{-1/2} Z^\top \Lambda E S^{-1/2}, \tag{67}$$

the operator norm of the linear map $E \mapsto C$, from Frobenius norm to Frobenius norm, satisfies

$$\|E \mapsto C\|^2 = \|S^{-1}\|_2 \big\|S^{-1/2} K S^{-1/2}\big\|_2. \tag{68}$$

This proves the pointwise estimate. The matrix inequality $\Lambda^2 \preceq \lambda_1 \Lambda$ gives $K \preceq \lambda_1 S$, and hence $L(Z) \le 1 + \lambda_1/\lambda_{\min}(S)$. If $W \succeq I_k$, then $S \succeq \lambda_d W \succeq \lambda_d I_k$, and hence $L(Z) \le 1 + \lambda_1/\lambda_d = 1 + \kappa$.

Step 3: bound along a line segment. Integrating the pointwise Hessian bound along $[Z, Y]$ proves the final statement. □

A.4 Critical points and local convergence geometry

We characterize the critical points and their Hessians. Negative curvature identifies the strict saddles; positive curvature normal to $\mathcal{Z}_\star$ gives the local PL inequality used for convergence rates. Let $\mathrm{Crit}(h)$ denote the set of full-column-rank matrices Z satisfying $\nabla h(Z) = 0$.

Lemma A.8 (Critical manifolds). *A full-column-rank matrix Z is critical for h if and only if*

$$Z^\top Z = I_k \qquad \text{and} \qquad \Lambda Z = Z(Z^\top \Lambda Z). \tag{69}$$

Thus its column space is a k-dimensional invariant subspace of Λ. For an index set $I \subseteq \{1, \dots, d\}$ of cardinality k, the corresponding critical manifold is

$$\mathcal{Z}_I = \{E_I R : R \in O(k)\}, \tag{70}$$

when the eigenvalues are simple. More generally, let $\theta_1 > \cdots > \theta_m$ be the distinct eigenvalues of Λ, with eigenspaces $\mathcal{E}_1, \dots, \mathcal{E}_m$ and multiplicities $d_1, \dots, d_m$. Each critical manifold is indexed by a profile $\mathbf{r} = (r_1, \dots, r_m)$ satisfying $0 \le r_\ell \le d_\ell$ and $\sum_\ell r_\ell = k$; its column space is

$$\mathcal{V}_1 \oplus \cdots \oplus \mathcal{V}_m, \qquad \mathcal{V}_\ell \in \mathrm{Gr}(r_\ell, \mathcal{E}_\ell), \tag{71}$$

and its matrices are all orthonormal bases of that subspace. There are only finitely many such profiles. The unique minimizing subspace is the leading one,

$$\mathcal{Z}_\star = \{E_k R : R \in O(k)\}. \tag{72}$$

Every other critical manifold has a strictly negative normal-curvature direction.

Proof. **Step 1: characterization of the critical points.** The critical equation is

$$Z = \Lambda Z S^{-1}, \qquad S = Z^\top \Lambda Z, \tag{73}$$

or $\Lambda Z = ZS$. Multiplication by $Z^\top$ gives $S = (Z^\top Z)S$, and $S \succ 0$ implies $Z^\top Z = I_k$. Therefore the column space is invariant. Conversely, any orthonormal basis of an invariant subspace satisfies this critical equation. Such a subspace is a direct sum of subspaces of the eigenspaces of Λ, giving the stated profiles.

Step 2: minimizing subspace. At $Z = E_I R$,

$$h(Z) = \frac{k}{2} - \frac{1}{2}\log(\det \Lambda_I). \tag{74}$$

The same formula, with the corresponding eigenvalue multiplicities, holds for every profile. The product of selected eigenvalues is uniquely maximized by the leading k-dimensional eigenspace because $\lambda_k > \lambda_{k+1}$; when $k = d$, the whole space is the only k-dimensional subspace.

Step 3: negative curvature at every other critical manifold. If a nonleading invariant subspace contains an eigenvector of eigenvalue λ_i while omitting one of eigenvalue $\lambda_j > \lambda_i$, rotate the former toward the latter. Along this rotation,

$$\frac{d^2}{d\theta^2} h\big|_{\theta=0} = -\frac{\lambda_j - \lambda_i}{\lambda_i} < 0. \tag{75}$$

Hence every nonleading critical manifold is a strict saddle in a normal direction. □

Lemma A.9 (Hessian at a critical point). *At a critical point Z of h, the Hessian has eigenvalue 2 on perturbations ZC with $C = C^\top$, eigenvalue 0 along right-orthogonal rotations, and eigenvalues $1 - \lambda_j/\lambda_i$ on perturbations from a selected eigenvector i toward an omitted eigenvector j. Its kernel is exactly the tangent space of the critical manifold. In particular, its largest eigenvalue is 2.*

Proof. **Step 1: block form of the Hessian.** Choose orthonormal eigenbases in the invariant subspace of the critical point and in its orthogonal complement. After an orthogonal change of ambient coordinates and right multiplication by an orthogonal matrix, write

$$Z = \begin{bmatrix} I_k \\ 0 \end{bmatrix}, \qquad \Lambda = \mathrm{diag}(\Lambda_I, \Lambda_{I^c}), \qquad E = \begin{bmatrix} C \\ D \end{bmatrix}, \tag{76}$$

where Λ_I contains the selected eigenvalues and Λ_{I^c} the omitted ones. These orthogonal changes preserve the Hessian spectrum. Differentiating the gradient formula gives

$$\nabla^2 h(Z)[E] = \begin{bmatrix} C + C^\top \\ D - \Lambda_{I^c} D \Lambda_I^{-1} \end{bmatrix}. \tag{77}$$

Step 2: eigenvalues and tangent directions. Thus the selected-row block has eigenvalue 2 on symmetric matrices C and eigenvalue 0 on skew-symmetric matrices C. Each entry of D rotates a selected eigenvector of eigenvalue λ_i toward an omitted eigenvector of eigenvalue λ_j, and has Hessian eigenvalue $1 - \lambda_j/\lambda_i$. When $\lambda_i = \lambda_j$, these zero modes are exactly tangent to the Grassmannian factors inside the repeated eigenspaces. Together with the right-orthogonal directions, they exhaust the tangent space of the critical manifold. All other normal Hessian eigenvalues are nonzero. The positive eigenvalues associated with D lie in $(0, 1)$, whereas Lemma A.8 supplies at least one negative eigenvalue for every nonleading profile.

The space of symmetric C has positive dimension, and all eigenvalues associated with D are smaller than 1, so the largest eigenvalue is 2. □

For $k < d$, a global PL inequality is impossible because nonleading invariant subspaces are stationary with positive objective gap. Near the leading manifold, the normal Hessian gives the following sharp local constant.

Lemma A.10 (Sharp local PL constant)**.** *Let*

$$K_\star := \begin{cases} \dfrac{\lambda_k}{2(\lambda_k - \lambda_{k+1})}, & k < d, \\ \dfrac{1}{4}, & k = d. \end{cases} \tag{78}$$

For every $\overline{K} > K_\star$, there is a neighborhood $\mathcal{U}_{\overline{K}}$ of $\mathcal{Z}_\star$ such that every $Z \in \mathcal{U}_{\overline{K}}$ has full column rank and

$$h(Z) - h_\star \le \overline{K} \|\nabla h(Z)\|_F^2. \tag{79}$$

Moreover, $K_\star$ is the limiting worst-case ratio of $h(Z) - h_\star$ to $\|\nabla h(Z)\|_F^2$ as Z approaches $\mathcal{Z}_\star$ from outside $\mathcal{Z}_\star$.

Proof. **Step 1: smallest Hessian eigenvalue normal to $\mathcal{Z}_\star$.** The objective and gradient norm are invariant under the right action of $O(k)$, so it suffices to work on a normal slice through E_k. By Lemma A.9, the Hessian kernel at E_k is exactly the tangent space of $\mathcal{Z}_\star$. Its restriction to the normal space is positive definite: perturbations $E_k C$ with $C = C^\top$ have eigenvalue 2, and perturbations toward omitted eigenvectors have eigenvalues $1 - \lambda_j/\lambda_i$ for $i \le k < j$, all positive by the spectral gap. The smallest normal eigenvalue is

$$\mu_\star = \begin{cases} 1 - \dfrac{\lambda_{k+1}}{\lambda_k} = \dfrac{\lambda_k - \lambda_{k+1}}{\lambda_k}, & k < d, \\ 2, & k = d. \end{cases} \tag{80}$$

Step 2: Taylor expansion relative to a closest minimizer. Choose a closest point $Z_\star = E_k R$ to Z on $\mathcal{Z}_\star$; it exists because $\mathcal{Z}_\star$ is compact. First-order optimality implies that $Z - Z_\star$ is orthogonal to the tangent space at $Z_\star$. After right multiplication by $R^\top$, write $ZR^\top = E_k + e$, where $e = (Z - Z_\star)R^\top$ lies in the normal space at E_k. Taylor expansion gives, uniformly along the compact solution manifold,

$$h(Z) - h_\star = \frac{1}{2}\langle e, \nabla^2 h(E_k) e\rangle + O(\|e\|_F^3) \tag{81}$$

and

$$\nabla h(Z) R^\top = \nabla^2 h(E_k) e + O(\|e\|_F^2). \tag{82}$$

If H_N denotes the Hessian restricted to the normal space, then $H_N \succeq \mu_\star I_{dk - k(k-1)/2}$ and hence

$$\frac{1}{2}\langle e, H_N e\rangle \le \frac{1}{2\mu_\star}\|H_N e\|_F^2. \tag{83}$$

The gradient expansion also shows that $\nabla h(Z) \neq 0$ for Z sufficiently close to, but outside, $\mathcal{Z}_\star$. The two uniform Taylor expansions therefore imply

$$\limsup_{\substack{Z \to \mathcal{Z}_\star \\ Z \notin \mathcal{Z}_\star}} \frac{h(Z) - h_\star}{\|\nabla h(Z)\|_F^2} \le \frac{1}{2\mu_\star}. \tag{84}$$

Step 3: sharpness and a uniform neighborhood. Equality is attained in the limit along a normal eigenmode with eigenvalue $\mu_\star$. Thus the limiting worst-case ratio is

$$\frac{1}{2\mu_\star} = K_\star. \tag{85}$$

Uniformity along the compact solution manifold shows that every $\overline{K} > K_\star$ is valid on a sufficiently small neighborhood of the whole manifold. Shrinking that neighborhood if necessary also ensures full column rank. □

B PROOFS :: RANK-k APPROXIMATION FOR SYMMETRIC MATRICES

This appendix proves Theorem B.9 for every $0 < \eta < 1$ and then derives the iterate and reconstruction rates in Theorem 3.1. Determinant ascent controls $R_t H_t^{-1}$, while the recursion for $W_t = Z_t^\top Z_t$ controls $W_t - I_k$. Summability gives convergence to one critical point; saddle avoidance then identifies the limit as a minimizer. Descent of h is used only to establish the eventual rate. The shared identities and local geometry are in §A.

B.1 PROOF OF LEMMA 4.1

We first bound $W_t = Z_t^\top Z_t$ without using determinant ascent. For a full-rank $Z = QW^{1/2}$, retain the notation

$$H = Q^\top \Lambda Q, \qquad R = (I_d - QQ^\top)\Lambda Q. \tag{86}$$

Here $QQ^\top$ is the orthogonal projector onto $\mathrm{range}(Q)$, so R is the component of ΛQ orthogonal to the current subspace. Equivalently, $R = \Lambda Q - QH$; it vanishes exactly when $\mathrm{range}(Q)$ is invariant under Λ. For $W \succ 0$ and $0 < \eta < 1$, define the weighted Heron map

$$\Phi_\eta(W) := \Big((1-\eta)W^{1/2} + \eta W^{-1/2}\Big)^2. \tag{87}$$

Lemma B.1 (Full-range spectral annulus). *Let $0 < \eta < 1$, assume $I_k \preceq W_0 \preceq \beta I_k$, and define*

$$\underline{m}_\eta := \begin{cases} 1, & 0 < \eta \le 1/2, \\ 4\eta(1-\eta), & 1/2 < \eta < 1, \end{cases} \qquad \overline{m}_{\eta,\beta} := \max\left\{\beta, \frac{2(1-\eta) + \eta\kappa/\underline{m}_\eta}{2-\eta}\right\}. \tag{88}$$

Then every iterate is full rank and

$$\underline{m}_\eta I_k \preceq W_t \preceq \overline{m}_{\eta,\beta} I_k, \qquad t \ge 0. \tag{89}$$

Proof. **Step 1: lower bound.** Substituting $Z_t = Q_t W_t^{1/2}$ and $\Lambda Q_t = Q_t H_t + R_t$ into (37), and using $Q_t^\top R_t = 0$, gives the exact radial recursion

$$W_{t+1} = \Big((1-\eta)W_t^{1/2} + \eta W_t^{-1/2}\Big)^2 + \eta^2 W_t^{-1/2}(R_t H_t^{-1})^\top (R_t H_t^{-1}) W_t^{-1/2}. \tag{90}$$

The residual correction in (90) is positive semidefinite, so

$$W_{t+1} \succeq \Big((1-\eta)W_t^{1/2} + \eta W_t^{-1/2}\Big)^2. \tag{91}$$

For every scalar $w > 0$,

$$\left((1-\eta)\sqrt{w} + \frac{\eta}{\sqrt{w}}\right)^2 \ge 4\eta(1-\eta). \tag{92}$$

For $\eta \le 1/2$, the scalar map $r \mapsto (1-\eta)r + \eta/r$ is increasing on $[1,\infty)$ and equals 1 at $r = 1$. Thus $W_t \succeq I_k$ implies $W_{t+1} \succeq I_k$. For $\eta > 1/2$, the displayed arithmetic–geometric mean bound gives $W_{t+1} \succeq 4\eta(1-\eta)I_k$. Starting from $W_0 \succeq I_k$, these bounds also ensure full rank at every step.

Step 2: upper bound. Lemma A.5 gives $I_k + (R_t H_t^{-1})^\top (R_t H_t^{-1}) \preceq \kappa I_k$. Using this in (90) gives

$$W_{t+1} \preceq (1-\eta)^2 W_t + 2\eta(1-\eta)I_k + \eta^2 \kappa W_t^{-1} \tag{93}$$

$$\preceq (1-\eta)^2 W_t + \left(2\eta(1-\eta) + \frac{\eta^2\kappa}{\underline{m}_\eta}\right) I_k. \tag{94}$$

The second term divided by $1 - (1-\eta)^2 = \eta(2-\eta)$ is the stated scalar fixed point. Induction from $W_0 \preceq \beta I_k$ proves the upper bound. □

B.2 Proof of Lemma 5.1

Write $\Psi(Q) = \log(\det(Q^\top \Lambda Q))$. The next lemma establishes determinant ascent. Its quantitative form, combined with the upper bound on W_t from Lemma B.1, gives angular summability.

Lemma B.2 (Full-range determinant–Ritz improvement)**.** *Fix $0 < \eta \le 1$ and $t \ge 0$. Let Z_t have full column rank, let Z_{t+1} be its update under (37), and let Q_{t+1} be any orthonormal basis of* $\mathrm{range}(Z_{t+1})$. *Write $A_t := (R_t H_t^{-1})^\top (R_t H_t^{-1})$. Then*

$$\Psi(Q_{t+1}) - \Psi(Q_t) \ge \log\Big(\det\big(I_k + \eta A_t\big((1-\eta)W_t + \eta I_k\big)^{-1}\big)\Big) \ge 0. \tag{95}$$

The determinant inequality is strict unless $R_t = 0$. Moreover, for $0 < \eta < 1$ and initialization $I_k \preceq W_0 \preceq \beta I_k$, with $\overline{m}_{\eta,\beta}$ as in Lemma B.1,

$$\Psi(Q_{t+1}) - \Psi(Q_t) \ge \frac{\eta}{(1-\eta)\overline{m}_{\eta,\beta} + \eta\kappa} \|R_t H_t^{-1}\|_F^2. \tag{96}$$

Proof. If $k = d$, then $R_t = 0$ and the claim is immediate. Assume $k < d$.

Step 1: determinant ratio identity. The whitened update gives

$$Z_{t+1} = Y_t W_t^{-1/2}, \qquad Y_t := Q_t\big((1-\eta)W_t + \eta I_k\big) + \eta R_t H_t^{-1}. \tag{97}$$

Thus $\mathrm{range}(Z_{t+1}) = \mathrm{range}(Y_t)$. Moreover, $Q_t^\top Y_t = (1-\eta)W_t + \eta I_k \succ 0$, so Y_t has full column rank. For the polar factor Q_{t+1}, the identities $Z_{t+1} = Y_t W_t^{-1/2} = Q_{t+1} W_{t+1}^{1/2}$ give

$$Y_t = Q_{t+1} W_{t+1}^{1/2} W_t^{1/2}. \tag{98}$$

Since $Q_{t+1}^\top Q_{t+1} = I_k$, the factor $\det(W_t)\det(W_{t+1})$ cancels between $\det(Y_t^\top \Lambda Y_t)$ and $\det(Y_t^\top Y_t)$, giving

$$\det(Q_{t+1}^\top \Lambda Q_{t+1}) = \frac{\det(Y_t^\top \Lambda Y_t)}{\det(Y_t^\top Y_t)}. \tag{99}$$

The identity holds for any orthonormal basis of $\mathrm{range}(Y_t)$, since an orthogonal change of basis preserves the determinant. We will lower bound the numerator and upper bound the denominator.

Step 2: lower bound on the numerator. Complete Q_t to an orthogonal matrix $[Q_t, Q_{t,\perp}]$ and write

$$[Q_t, Q_{t,\perp}]^\top \Lambda [Q_t, Q_{t,\perp}] = \begin{bmatrix} H_t & C^\top \\ C & L \end{bmatrix} = \begin{bmatrix} Q_t^\top \Lambda Q_t & Q_t^\top \Lambda Q_{t,\perp} \\ Q_{t,\perp}^\top \Lambda Q_t & Q_{t,\perp}^\top \Lambda Q_{t,\perp} \end{bmatrix}. \tag{100}$$

The Schur complement gives $L \succeq C H_t^{-1} C^\top$. This separates a positive-semidefinite remainder from a term whose determinant factors explicitly. Since $R_t = Q_{t,\perp} C$ and $Q_{t,\perp}^\top Q_{t,\perp} = I_{d-k}$, we have $A_t = H_t^{-1} C^\top C H_t^{-1}$. Thus

$$(R_t H_t^{-1})^\top \Lambda (R_t H_t^{-1}) = H_t^{-1} C^\top L C H_t^{-1} \succeq H_t^{-1} C^\top (C H_t^{-1} C^\top) C H_t^{-1} = A_t H_t A_t. \tag{101}$$

Using this inequality and $Q_t^\top \Lambda R_t H_t^{-1} = H_t A_t$ in the expansion of $Y_t^\top \Lambda Y_t$ gives

$$Y_t^\top \Lambda Y_t \succeq \big((1-\eta)W_t + \eta I_k + \eta A_t\big) H_t \big((1-\eta)W_t + \eta I_k + \eta A_t\big). \tag{102}$$

Hence

$$\det(Y_t^\top \Lambda Y_t) \ge \det(H_t)\det\big((1-\eta)W_t + \eta I_k + \eta A_t\big)^2. \tag{103}$$

Step 3: upper bound on the denominator. Orthogonality yields

$$Y_t^\top Y_t = \big((1-\eta)W_t + \eta I_k\big)^2 + \eta^2 A_t \succ 0. \tag{104}$$

Since $(1-\eta)W_t + \eta I_k \succeq \eta I_k$, spectral calculus gives

$$\eta^2\big((1-\eta)W_t + \eta I_k\big)^{-2} \preceq \eta\big((1-\eta)W_t + \eta I_k\big)^{-1}. \tag{105}$$

Factoring the squared matrix out of $Y_t^\top Y_t$, then using Sylvester's determinant identity and this inequality, yields

$$\det(Y_t^\top Y_t) = \det\big((1-\eta)W_t + \eta I_k\big)^2 \det\Big(I_k + \eta^2 A_t^{1/2}\big((1-\eta)W_t + \eta I_k\big)^{-2} A_t^{1/2}\Big) \tag{106}$$

$$\leq \det\big((1-\eta)W_t + \eta I_k\big)\det\big((1-\eta)W_t + \eta I_k + \eta A_t\big). \tag{107}$$

Step 4: combining the bounds. Substituting these estimates into (99) gives

$$\det(Q_{t+1}^\top \Lambda Q_{t+1}) \geq \det(H_t)\frac{\det\big((1-\eta)W_t + \eta I_k + \eta A_t\big)}{\det\big((1-\eta)W_t + \eta I_k\big)}. \tag{108}$$

This proves (95). By Sylvester's identity, the determinant in its logarithm equals that of $I_k + \eta A_t^{1/2}((1-\eta)W_t + \eta I_k)^{-1}A_t^{1/2}$. It is at least 1, and strictly greater than 1 when $A_t \neq 0$, equivalently when $R_t \neq 0$.

Step 5: quantitative improvement. By Lemma B.1, we have $W_t \preceq \overline{m}_{\eta,\beta} I_k$, so

$$\eta\big((1-\eta)W_t + \eta I_k\big)^{-1} \succeq p I_k, \qquad p := \frac{\eta}{(1-\eta)\overline{m}_{\eta,\beta} + \eta}. \tag{109}$$

Using Sylvester's identity again and the bound $A_t \preceq (\kappa - 1)I_k$ from Lemma A.5, we obtain

$$\log\Big(\det\big(I_k + \eta A_t\big((1-\eta)W_t + \eta I_k\big)^{-1}\big)\Big) \geq \log(\det(I_k + pA_t)) \tag{110}$$

$$\geq \frac{p}{1 + p(\kappa - 1)}\operatorname{tr} A_t \tag{111}$$

$$= \frac{\eta}{(1-\eta)\overline{m}_{\eta,\beta} + \eta\kappa}\|R_t H_t^{-1}\|_F^2. \tag{112}$$

The second inequality follows by applying $\log(1+x) \geq x/(1+x)$ to the eigenvalues of pA_t. □

Corollary B.3 (Angular residual summability)**.** *For $0 < \eta < 1$,*

$$\sum_{t=0}^{\infty}\|R_t H_t^{-1}\|_F^2 \leq \left(\frac{(1-\eta)\overline{m}_{\eta,\beta} + \eta\kappa}{\eta}\right) k\log(\kappa). \tag{113}$$

In particular, $R_t H_t^{-1} \to 0$ and $R_t \to 0$.

Proof. For every integer $T \geq 0$, summing (96) from $t = 0$ to T gives

$$\frac{\eta}{(1-\eta)\overline{m}_{\eta,\beta} + \eta\kappa}\sum_{t=0}^{T}\|R_t H_t^{-1}\|_F^2 \leq \sum_{t=0}^{T}\big(\Psi(Q_{t+1}) - \Psi(Q_t)\big) \tag{114}$$

$$= \Psi(Q_{T+1}) - \Psi(Q_0) \tag{115}$$

$$\leq \log(\det\Lambda_k) - \Psi(Q_0). \tag{116}$$

The last inequality follows from Poincaré separation. Moreover, $H_0 \succeq \lambda_d I_k$ implies $\Psi(Q_0) \geq k\log\lambda_d$, so

$$\log(\det\Lambda_k) - \Psi(Q_0) \leq \sum_{i=1}^{k}\log\left(\frac{\lambda_i}{\lambda_d}\right) \leq k\log(\kappa). \tag{117}$$

Letting $T \to \infty$ proves (113), since the summands are nonnegative. In particular, $\|R_t H_t^{-1}\|_F^2 \to 0$. Finally, $H_t \preceq \lambda_1 I_k$ converts $R_t H_t^{-1} \to 0$ into $R_t \to 0$. □

B.3 Proof of Lemma 5.2

We prove summability of $\|W_t - I_k\|_F$. The next subsection establishes summability of $\|R_t H_t^{-1}\|_F$ and combines the two estimates to obtain finite total movement of Z_t.

Lemma B.4 (Radial stabilization)**.** *Let $0 < \eta < 1$ and $I_k \preceq W_0 \preceq \beta I_k$. With the bounds from Lemma B.1, define*

$$\rho_{\eta,\beta} := \max_{w\in[\underline{m}_\eta,\overline{m}_{\eta,\beta}]}\left|\frac{(1-\eta)w - \eta}{(1-\eta)w + \eta}\right| < 1. \tag{118}$$

Then

$$\sum_{t=0}^{\infty}\|W_t - I_k\|_F \leq \frac{\overline{m}_{\eta,\beta}}{1-\rho_{\eta,\beta}}\left(\sqrt{k}\log(\beta) + \frac{\eta\big((1-\eta)\overline{m}_{\eta,\beta} + \eta\kappa\big)}{\underline{m}_\eta^2}\,k\log(\kappa)\right). \tag{119}$$

In particular, $W_t \xrightarrow[t\to\infty]{} I_k$.

Proof. **Step 1: bound on the correction to $\Phi_\eta(W_t)$.** Write (90) as

$$W_{t+1} = \Phi_\eta(W_t) + E_t, \qquad E_t := \eta^2 W_t^{-1/2} A_t W_t^{-1/2} \succeq 0. \tag{120}$$

Since $A_t = (R_t H_t^{-1})^\top (R_t H_t^{-1}) \succeq 0$, the lower matrix bound gives

$$\|E_t\|_F \le \frac{\eta^2}{\underline{m}_\eta} \|A_t\|_F \le \frac{\eta^2}{\underline{m}_\eta} \operatorname{tr}(A_t) = \frac{\eta^2}{\underline{m}_\eta} \|R_t H_t^{-1}\|_F^2. \tag{121}$$

Step 2: contraction of $\log(\Phi_\eta(W_t))$. The weighted Heron update contracts in logarithmic coordinates. Indeed, set

$$F_\eta(x) := 2\log\big((1-\eta)e^{x/2} + \eta e^{-x/2}\big). \tag{122}$$

Then $F_\eta(0) = 0$ and, for $x \in [\log(\underline{m}_\eta), \log(\overline{m}_{\eta,\beta})]$,

$$|F_\eta'(x)| = \left| \frac{(1-\eta)e^x - \eta}{(1-\eta)e^x + \eta} \right| \le \rho_{\eta,\beta}. \tag{123}$$

This interval contains zero. Applying the mean value theorem to each eigenvalue of $\log(W_t)$ therefore yields

$$\|\log(\Phi_\eta(W_t))\|_F \le \rho_{\eta,\beta} \|\log(W_t)\|_F. \tag{124}$$

Step 3: summability of $\|W_t - I_k\|_F$. The scalar lower bounds in the proof of Lemma B.1 also give $\Phi_\eta(W_t) \succeq \underline{m}_\eta I_k$. The integral representation of the matrix logarithm and the resolvent identity thus imply

$$\|\log(W_{t+1}) - \log(\Phi_\eta(W_t))\|_F \le \|E_t\|_F \int_0^\infty \frac{ds}{(\underline{m}_\eta + s)^2} = \frac{\|E_t\|_F}{\underline{m}_\eta}. \tag{125}$$

Combining this estimate with (121) and (124) gives

$$\|\log(W_{t+1})\|_F \le \rho_{\eta,\beta} \|\log(W_t)\|_F + \frac{\eta^2}{\underline{m}_\eta^2} \|R_t H_t^{-1}\|_F^2. \tag{126}$$

For every integer $T \ge 1$, summing this recursion and dropping the nonpositive terminal term gives

$$(1-\rho_{\eta,\beta}) \sum_{t=0}^{T} \|\log(W_t)\|_F \le \|\log(W_0)\|_F + \frac{\eta^2}{\underline{m}_\eta^2} \sum_{t=0}^{T-1} \|R_t H_t^{-1}\|_F^2. \tag{127}$$

Finally, $|w-1| \le \overline{m}_{\eta,\beta} |\log(w)|$ on the spectral interval, so $\|W_t - I_k\|_F \le \overline{m}_{\eta,\beta} \|\log(W_t)\|_F$. The initialization bounds place every eigenvalue of $\log(W_0)$ in $[0, \log(\beta)]$, hence $\|\log(W_0)\|_F \le \sqrt{k}\log(\beta)$. Using this bound, letting $T \to \infty$, and applying Corollary B.3 proves (119). Summability implies $W_t \xrightarrow[t\to\infty]{} I_k$. □

B.4 Proof of Proposition 5.3

Lemma B.5 (Explicit multi-level determinant PL inequality). *Suppose $k < d$. Let*

$$\mathcal{C}_{\det} := \big\{ \log\big(\textstyle\prod_{i\in I} \lambda_i\big) : I \subseteq \{1,\dots,d\},\ |I| = k \big\}, \tag{128}$$

where repeated values are retained only once. If this set is a singleton, $\Psi(Q)$ equals its sole element for every $Q^\top Q = I_k$. Otherwise, set

$$\gamma_{\det} := \min\{|c - c'| : c, c' \in \mathcal{C}_{\det},\ c \ne c'\} > 0. \tag{129}$$

In this nontrivial case, every $Q^\top Q = I_k$ satisfies

$$\operatorname{dist}\big(\Psi(Q), \mathcal{C}_{\det}\big) \le \frac{\|RH^{-1}\|_F^2}{1 - e^{-\gamma_{\det}/2}}. \tag{130}$$

Proof. If $\mathcal{C}_{\det}$ is a singleton, Cauchy–Binet expresses $\det(Q^\top \Lambda Q)$ as a convex combination of identical products, which proves the assertion. In the remaining case $\gamma_{\det} > 0$.

Step 1: the residual controls the logarithmic determinant discrepancy. The identity $Q^\top \Lambda^2 Q = H^2 + R^\top R$ gives

$$\begin{aligned}\log\left(\frac{\det(Q^\top \Lambda^2 Q)}{\det(H)^2}\right) &= \log\big(\det\big(I_k + (RH^{-1})^\top (RH^{-1})\big)\big) \\ &\le \|RH^{-1}\|_F^2,\end{aligned} \tag{131}$$

where the inequality follows by applying $\log(1+s) \le s$ to the nonnegative eigenvalues of $(RH^{-1})^\top (RH^{-1})$.

Step 2: compare with the adjacent critical values. For each k-element index set I, let Q_I denote the submatrix formed by the rows indexed by I, and write $p_I := \det(Q_I)^2$ and $v_I := \prod_{i\in I} \lambda_i$. All sums below run over these index sets. Cauchy–Binet gives

$$\sum_I p_I = 1, \qquad \sum_I p_I v_I = \det(H), \qquad \sum_I p_I v_I^2 = \det(Q^\top \Lambda^2 Q). \tag{132}$$

If $\Psi(Q) \in \mathcal{C}_{\det}$, the lemma is immediate. Otherwise, $\det(H)$ is a convex combination of the v_I, so $\Psi(Q)$ lies between two adjacent values $a < b$ of $\mathcal{C}_{\det}$. Every v_I lies outside (e^a, e^b), hence

$$\begin{aligned}0 \le \sum_I p_I (v_I - e^a)(v_I - e^b) = \det(Q^\top \Lambda^2 Q) - \det(H)^2 \\ - (\det(H) - e^a)(e^b - \det(H)).\end{aligned} \tag{133}$$

Define the two logarithmic distances

$$\ell := \Psi(Q) - a > 0, \qquad r := b - \Psi(Q) > 0. \tag{134}$$

Then $\ell + r = b - a \ge \gamma_{\det}$ and $\min\{\ell, r\} = \operatorname{dist}(\Psi(Q), \mathcal{C}_{\det})$. Dividing the preceding bound by $\det(H)^2$ and using (131) yields

$$\|RH^{-1}\|_F^2 \ge \log\big(1 + (1 - e^{-\ell})(e^r - 1)\big)\,. \tag{135}$$

Step 3: control the smaller logarithmic distance. Since $\max\{\ell, r\} \ge \gamma_{\det}/2$,

$$\begin{aligned}(1 - e^{-\ell})(e^r - 1) &\ge (1 - e^{-\max\{\ell,r\}})(e^{\min\{\ell,r\}} - 1) \\ &\ge (1 - e^{-\gamma_{\det}/2})(e^{\min\{\ell,r\}} - 1).\end{aligned} \tag{136}$$

The first inequality is an equality when $\ell \ge r$; otherwise, it follows by writing $1 - e^{-\ell} = e^{-\ell}(e^\ell - 1)$ and using $e^{-\ell} \ge e^{-r}$. For $0 \le c \le 1$ and $x \ge 0$, convexity of the exponential gives

$$1 + c(e^x - 1) = (1 - c)e^0 + ce^x \ge e^{cx}. \tag{137}$$

Applying this with $c = 1 - e^{-\gamma_{\det}/2}$ and $x = \min\{\ell, r\}$ in (135) gives

$$\|RH^{-1}\|_F^2 \ge (1 - e^{-\gamma_{\det}/2}) \min\{\ell, r\} = (1 - e^{-\gamma_{\det}/2}) \operatorname{dist}\big(\Psi(Q), \mathcal{C}_{\det}\big), \tag{138}$$

which proves (130). □

Lemma B.6 (Convergence to one critical point for $0 < \eta < 1$)**.** *Let $0 < \eta < 1$ and $I_k \preceq W_0 \preceq \beta I_k$. Then*

$$\sum_{t=0}^{\infty} \|R_t H_t^{-1}\|_F < \infty, \qquad \sum_{t=0}^{\infty} \|Z_{t+1} - Z_t\|_F < \infty. \tag{139}$$

Consequently, $Z_t \xrightarrow[t\to\infty]{} Z_\star$ for some full-rank $Z_\star$ satisfying $\nabla h(Z_\star) = 0$.

Proof. **Step 1: convergence of the determinant to a critical value.** If $k = d$ or Λ is a scalar multiple of I_d, then $R_t = 0$ for all t, so $\sum_t \|R_t H_t^{-1}\|_F = 0$ and we proceed directly to Step 3. Otherwise, Lemma B.2 gives

$$\Psi(Q_{t+1}) - \Psi(Q_t) \ge c_{\mathrm{ang}} \|R_t H_t^{-1}\|_F^2, \qquad c_{\mathrm{ang}} := \frac{\eta}{(1-\eta)\overline{m}_{\eta,\beta} + \eta\kappa} > 0. \tag{140}$$

Since $\Psi(Q_t) \le \log(\det \Lambda_k)$, this increasing sequence converges to a finite value Ψ_∞. The orthonormal matrices Q_t have an accumulation point Q_∞. By Corollary B.3, $R_t \to 0$, so $\mathrm{range}(Q_\infty)$ is invariant under Λ. Continuity therefore gives

$$\Psi_\infty \in \mathcal{C}_{\det} := \left\{ \log\left(\prod_{i\in I} \lambda_i\right) : I \subseteq \{1,\dots,d\},\ |I| = k \right\}. \tag{141}$$

Step 2: summability of $\|R_t H_t^{-1}\|_F$. The pointwise inequality in Lemma B.5 gives

$$\mathrm{dist}(\Psi(Q_t), \mathcal{C}_{\det}) \le \frac{\|R_t H_t^{-1}\|_F^2}{1 - e^{-\gamma_{\det}/2}}. \tag{142}$$

Since $\mathcal{C}_{\det}$ is finite and $\Psi(Q_t) \to \Psi_\infty$, there is an integer $T \ge 0$ such that $0 \le \Psi_\infty - \Psi(Q_t) < \gamma_{\det}/2$ for every $t \ge T$. Thus Ψ_∞ is the unique closest critical value. This uses separation of distinct critical values, not convergence of Q_t. Hence

$$e_t := \Psi_\infty - \Psi(Q_t) = \mathrm{dist}(\Psi(Q_t), \mathcal{C}_{\det}) \le \frac{\|R_t H_t^{-1}\|_F^2}{1 - e^{-\gamma_{\det}/2}}, \qquad t \ge T. \tag{143}$$

Combining this bound with (140) yields

$$e_{t+1} \le e_t - c_{\mathrm{ang}} \|R_t H_t^{-1}\|_F^2 \le q e_t, \qquad q := 1 - c_{\mathrm{ang}}(1 - e^{-\gamma_{\det}/2}) \in (0,1). \tag{144}$$

Here $0 < c_{\mathrm{ang}} < 1$ because $\overline{m}_{\eta,\beta} \ge 1$ and $\kappa \ge 1$. Moreover, $c_{\mathrm{ang}} \|R_t H_t^{-1}\|_F^2 \le e_t - e_{t+1} \le e_t$, so

$$\|R_t H_t^{-1}\|_F^2 \le \frac{e_T}{c_{\mathrm{ang}}} q^{t-T}, \qquad \sum_{t=T}^{\infty} \|R_t H_t^{-1}\|_F \le \frac{\sqrt{e_T / c_{\mathrm{ang}}}}{1 - \sqrt{q}} < \infty. \tag{145}$$

If $e_T = 0$, the ascent bound forces $R_t = 0$ for every $t \ge T$, so the same conclusion holds. Adding the finite prefix proves the first assertion of (139).

Step 3: finite total movement and criticality of the limit. For all $k \le d$, the gradient identity in Lemma A.6 and the lower matrix bound imply

$$\|Z_{t+1} - Z_t\|_F = \eta \|\nabla h(Z_t)\|_F \le \frac{\eta}{\sqrt{\overline{m}_\eta}} \left(\|W_t - I_k\|_F + \|R_t H_t^{-1}\|_F \right). \tag{146}$$

The terms $\|W_t - I_k\|_F$ are summable by Lemma B.4, and $\|R_t H_t^{-1}\|_F$ is summable by the preceding argument. Summing proves the second assertion of (139). Hence (Z_t) is Cauchy and converges to a matrix $Z_\star$. The uniform lower bound on W_t ensures that $Z_\star$ has full column rank. Finally, $\|\nabla h(Z_t)\|_F = \eta^{-1} \|Z_{t+1} - Z_t\|_F \to 0$, so continuity gives $\nabla h(Z_\star) = 0$. □

B.5 Proof of Proposition 5.4

We use the following criterion, restated from Cheridito et al. (2024, Proposition 2.5).

Proposition B.7 (General saddle-avoidance criterion). *Let $m \ge 1$, let $T : \mathbb{R}^m \to \mathbb{R}^m$, and let $\mathcal{V} \subseteq \mathcal{U} \subseteq \mathbb{R}^m$ be open sets whose complements have Lebesgue measure zero. Suppose that:*

(i) T is continuously differentiable on $\mathcal{U}$, and its Jacobian DT is locally Lipschitz on $\mathcal{U}$;

(ii) $\det DT(x) \ne 0$ for every $x \in \mathcal{V}$;

(iii) $\mathcal{S} \subseteq \{x \in \mathcal{U} : T(x) = x\}$, and for every $x \in \mathcal{S}$, the matrix $DT(x)$ is symmetric and has an eigenvalue of absolute value strictly greater than 1.

Then the set of initial points $x_0 \in \mathbb{R}^m$ whose iterates $T^t(x_0)$ converge to a point in $\mathcal{S}$ has Lebesgue measure zero.

The criterion applies to the entire set $\mathcal{S}$, even if it is uncountable, and permits singular Jacobians at its points. We apply it to the transformed ScaledGD map in the following lemma.

Lemma B.8 (Full-range saddle avoidance). *Fix $0 < \eta < 1$. The set of full-column-rank initial factors $Z_0 \in \mathbb{R}^{d\times k}$ whose iterates converge to a point in $\mathrm{Crit}(h) \setminus \mathcal{Z}_\star$ has Lebesgue measure zero. Consequently, any absolutely continuous law of Z_0 assigns probability zero to convergence to a nonminimizing critical point.*

Proof. If $k = d$, there are no nonminimizing critical points, so assume $k < d$. We apply Proposition B.7 to $T_\eta(Z) := Z - \eta\nabla h(Z)$ on $\mathcal{U} := \{Z \in \mathbb{R}^{d\times k} : \mathrm{rank}(Z) = k\}$, identifying matrices with $\mathbb{R}^{dk}$, so $m = dk$. The three steps verify regularity, almost-everywhere invertibility, and instability of the fixed points.

Step 1: regularity and an invariant full-measure domain. The set $\mathcal{U}$ is open and its complement is null, since $\det(Z^\top Z)$ is a nonzero polynomial. The map T_η is real analytic on $\mathcal{U}$, so it is continuously differentiable with a locally Lipschitz Jacobian. For the polar decomposition $Z = QW^{1/2}$,

$$Q^\top T_\eta(Z) = (1-\eta)W^{1/2} + \eta W^{-1/2} \succ 0. \tag{147}$$

Thus $T_\eta(\mathcal{U}) \subseteq \mathcal{U}$. To obtain a map on all of $\mathbb{R}^{dk}$ as required by the cited proposition, extend T_η by zero outside $\mathcal{U}$. This extension does not affect any trajectory starting in $\mathcal{U}$.

Step 2: almost-everywhere invertibility of the Jacobian. The function $Z \mapsto \det DT_\eta(Z)$ is real analytic on $\mathcal{U}$. For $Z \in \mathcal{U}$ and $r > 0$, Lemma A.7 and $(rZ)^\top\Lambda(rZ) = r^2 Z^\top\Lambda Z$ give

$$DT_\eta(rZ) = I_{dk} - \eta\nabla^2 h(rZ) \succeq \left(1 - \eta - \frac{\eta\lambda_1}{r^2\lambda_{\min}(Z^\top\Lambda Z)}\right) I_{dk}. \tag{148}$$

Since $\eta < 1$, the Jacobian $DT_\eta(rZ)$ is positive definite, hence invertible, whenever

$$r^2 > \frac{\eta\lambda_1}{(1-\eta)\lambda_{\min}(Z^\top\Lambda Z)}. \tag{149}$$

For each fixed Z, the continuous curve $r \mapsto rZ$, $r > 0$, preserves full column rank and therefore stays in the connected component of Z in $\mathcal{U}$. Thus every component contains a point with an invertible Jacobian, so the analytic determinant is not identically zero on any component.

A real-analytic function that is not identically zero on a connected open set has a zero set of Lebesgue measure zero (Mityagin, 2020). Applying this result on each of the at most countably many connected components of $\mathcal{U}$ shows that $\{Z \in \mathcal{U} : \det DT_\eta(Z) = 0\}$ is null. Consequently, $\mathcal{V} := \{Z \in \mathcal{U} : \det DT_\eta(Z) \neq 0\}$ is open and has full Lebesgue measure in $\mathbb{R}^{dk}$: openness follows from continuity, and its complement is null by this argument and Step 1.

Step 3: instability of every nonminimizing fixed point. Every $Z_\dagger \in \mathrm{Crit}(h) \setminus \mathcal{Z}_\star$ is a fixed point of T_η, and

$$DT_\eta(Z_\dagger) = I_{dk} - \eta\nabla^2 h(Z_\dagger) \tag{150}$$

is symmetric. By Lemma A.8, the subspace $\mathrm{col}(Z_\dagger)$ is invariant under Λ and is not a leading k-dimensional eigenspace. Hence there exist a selected eigenvalue λ_i, with an eigenvector in $\mathrm{col}(Z_\dagger)$, and an omitted eigenvalue $\lambda_j > \lambda_i$, with an eigenvector in $\mathrm{col}(Z_\dagger)^\perp$. By Lemma A.9, the perturbation rotating the selected direction toward the omitted one has Hessian eigenvalue $1 - \lambda_j/\lambda_i < 0$, so $DT_\eta(Z_\dagger)$ has eigenvalue $1 + \eta(\lambda_j/\lambda_i - 1) > 1$.

All hypotheses of Proposition B.7 now hold with the open sets $\mathcal{V} \subseteq \mathcal{U}$ and fixed-point set $\mathrm{Crit}(h)\setminus\mathcal{Z}_\star$. It follows that the initial points whose iterates converge to a point in this set form a Lebesgue-null set. The criterion permits singular Jacobians at the fixed points, so the argument also covers $\eta = 1/2$. Absolute continuity gives the probabilistic conclusion; Lemma A.2 verifies it for the prescribed initialization. □

B.6 Objective contraction near the limiting minimizer

Theorem B.9 (Global convergence and eventual linear contraction). *Under the assumptions and initialization of §3, fix $0 < \eta < 1$. Recall that ρ_η is defined in (7). Almost surely, every iterate is full column rank and $Z_t \xrightarrow[t\to\infty]{} Z_\star \in \mathcal{Z}_\star$. Moreover, there is an almost-surely finite time T such that (24) holds for every $t \geq T$.*

Proof. **Step 1: almost-sure convergence to a minimizer.** Lemma B.6 gives convergence to one critical point, and Lemma B.8 excludes all nonleading critical manifolds almost surely. Hence $Z_t \xrightarrow[t\to\infty]{} Z_\star \in \mathcal{Z}_\star$. Lemma B.1 gives full column rank at every step.

Step 2: local descent and objective contraction. Fix $2 < L < 2/\eta$ and $\overline{K} > K_\star$, where $K_\star$ is defined in Lemma A.10. By Lemma A.9, the Hessian at the limiting minimizer is bounded above by $2I_{dk}$. Continuity gives a neighborhood of that minimizer on which $\nabla^2 h \preceq L I_{dk}$. Since Z_t converges, every update segment eventually lies in this neighborhood. Taylor's formula then gives

$$h(Z_{t+1}) \le h(Z_t) - c_{\eta,L}\|\nabla h(Z_t)\|_F^2, \qquad c_{\eta,L} := \eta\left(1 - \frac{\eta L}{2}\right) > 0. \tag{151}$$

The trajectory also eventually lies in the local PL neighborhood from Lemma A.10. Combining these inequalities shows that, almost surely, there is a finite $T_{\overline{K},L}$ such that

$$h(Z_{t+1}) - h_\star \le q_{\overline{K},L}\big(h(Z_t) - h_\star\big), \qquad q_{\overline{K},L} := 1 - \frac{c_{\eta,L}}{\overline{K}}, \qquad t \ge T_{\overline{K},L}. \tag{152}$$

Moreover, $c_{\eta,L} < \eta(1-\eta) \le 1/4$ and $\overline{K} > K_\star \ge 1/4$, so $q_{\overline{K},L} \in (0,1)$. Absorbing the finite prefix into a random prefactor proves the geometric error bound $h(Z_t) - h_\star \le C q_{\overline{K},L}^t$ for all $t \ge 0$.

Step 3: the stated contraction factor. Taking $L = 1 + \eta^{-1}$ and $\overline{K} = \lambda_k/(\lambda_k - \lambda_{k+1})$ gives $\overline{K} > K_\star$ also when $k = d$, since then $\overline{K} = 1 > K_\star = 1/4$. These choices yield $q_{\overline{K},L} = \rho_\eta^2$ and prove (24). □

B.7 Proof of Theorem 3.1

Proof. **Step 1: convergence and objective rate.** By Lemmas B.6 and B.8, almost surely Z_t is well-defined and converges to one $Z_\star \in \mathcal{Z}_\star$. Thus $X_t = U\Lambda^{1/2}Z_t$ converges to $X_\star := U\Lambda^{1/2}Z_\star \in \mathcal{X}_\star$, and every X_t has full column rank. Set

$$\delta_k := \frac{\lambda_k - \lambda_{k+1}}{\lambda_k} \in (0,1], \qquad L = 1 + \eta^{-1}, \qquad \overline{K} = \delta_k^{-1}. \tag{153}$$

Since $0 < \eta < 1$, we have $2 < L < 2/\eta$. If $k < d$, then $\overline{K} = 2K_\star > K_\star$; if $k = d$, then $\overline{K} = 1 > K_\star = 1/4$. These choices in (152) give

$$c_{\eta,L} = \frac{\eta(1-\eta)}{2}, \qquad q_{\overline{K},L} = 1 - \frac{\eta(1-\eta)\delta_k}{2} = \rho_\eta^2. \tag{154}$$

Consequently, almost surely there is a finite $C_h > 0$ such that $h(Z_t) - h_\star \le C_h \rho_\eta^{2t}$ for all $t \ge 0$.

Step 2: convergence rate of X_t. For all sufficiently large t, the sufficient-decrease estimate (151) and $h(Z_{t+1}) \ge h_\star$ give

$$c_{\eta,L}\|\nabla h(Z_t)\|_F^2 \le h(Z_t) - h(Z_{t+1}) \le h(Z_t) - h_\star \le C_h \rho_\eta^{2t}. \tag{155}$$

Using $Z_{t+1} - Z_t = -\eta \nabla h(Z_t)$ and summing the geometric tail therefore yields

$$\begin{aligned} \|X_t - X_\star\|_F &\le \sqrt{\lambda_1} \sum_{j=t}^{\infty} \|Z_{j+1} - Z_j\|_F \\ &\le \frac{\eta\sqrt{\lambda_1 C_h / c_{\eta,L}}}{1 - \rho_\eta}\rho_\eta^t. \end{aligned} \tag{156}$$

Enlarging the finite random constant C_X includes the initial finite prefix.

Step 3: reconstruction rate. Since $X_\star X_\star^\top = M_k$ and $\|X_\star\|_2 = \sqrt{\lambda_1}$,

$$\|X_t X_t^\top - M_k\|_F \le (\|X_t\|_2 + \sqrt{\lambda_1})\|X_t - X_\star\|_F. \tag{157}$$

The spectral annulus gives $\|X_t\|_2 \le \sqrt{\lambda_1 \overline{m}_{\eta,\beta}}$ uniformly in t. Taking $C_M = (\sqrt{\lambda_1 \overline{m}_{\eta,\beta}} + \sqrt{\lambda_1})C_X$ proves the reconstruction bound and completes the proof. □

B.8 Removing the normalization of the initialization

The normalization in §3 simplifies the bounds used in the proof. It is not needed for convergence from a random initialization supported on the range of the target.

Corollary B.10 (Unnormalized random range initialization)**.** *Retain the spectral assumptions of §3 and fix $0 < \eta < 1$. Suppose $X_0 = UA_0$, where $A_0 \in \mathbb{R}^{d\times k}$ has an absolutely continuous distribution. Then all conclusions of Theorem 3.1 hold, with the same ρ_η. In particular, when $M \succ 0$, this covers every absolutely continuous law of $X_0 \in \mathbb{R}^{n\times k}$, including unnormalized Gaussian initialization.*

Proof. The change of variables $Z_0 = \Lambda^{-1/2}A_0$ preserves absolute continuity, and Z_0 has full column rank almost surely. Fix any such realization and set $W_0 = Z_0^\top Z_0 \succ 0$. Define

$$m := \min\{\lambda_{\min}(W_0), 4\eta(1-\eta)\} > 0, \qquad L := \max\left\{1, \lambda_{\max}(W_0), \frac{2(1-\eta)+\eta\kappa/m}{2-\eta}\right\}. \tag{158}$$

The radial identity (90) and the arithmetic–geometric mean bound give $W_{t+1} \succeq 4\eta(1-\eta)I_k$. The upper-bound argument of Lemma B.1 gives

$$W_{t+1} \preceq (1-\eta)^2 W_t + \big(2\eta(1-\eta) + \eta^2\kappa/m\big)\, I_k. \tag{159}$$

Induction therefore yields $mI_k \preceq W_t \preceq LI_k$ for every $t \geq 0$, and all iterates are full rank. These bounds depend on the realization of X_0 but are finite and positive.

The pointwise determinant inequality in Lemma B.2 does not require a normalized initialization. With the upper bound L, its quantitative form becomes

$$\Psi(Q_{t+1}) - \Psi(Q_t) \geq c\|R_t H_t^{-1}\|_F^2, \qquad c := \frac{\eta}{(1-\eta)L + \eta\kappa} > 0. \tag{160}$$

Thus $\sum_t \|R_t H_t^{-1}\|_F^2 \leq k\log(\kappa)/c$. Since $m \leq 1 \leq L$, the radial logarithmic contraction argument also applies on $[m, L]$, with

$$r := \max_{w\in[m,L]} \left|\frac{(1-\eta)w - \eta}{(1-\eta)w + \eta}\right| < 1. \tag{161}$$

The scalar lower bound also ensures $\Phi_\eta(W_t) \succeq mI_k$. Repeating the matrix-log estimate from Lemma B.4 therefore yields

$$\sum_{t=0}^{\infty} \|W_t - I_k\|_F \leq \frac{L}{1-r}\left(\|\log(W_0)\|_F + \frac{\eta^2}{m^2 c}\, k\log(\kappa)\right) < \infty. \tag{162}$$

The only change to the initial radial budget is to retain the finite quantity $\|\log(W_0)\|_F$ instead of bounding it by $\sqrt{k}\log\beta$.

The multi-level determinant inequality and the positive ascent constant c now give $\sum_t \|R_t H_t^{-1}\|_F < \infty$ by Steps 1–2 of Lemma B.6. Its Step 3, using m in the update-length bound, proves convergence to a full-rank critical point. Absolute continuity and Lemma B.8 exclude nonminimizing limits almost surely. Finally, the local argument in Theorem B.9 and §B.7 depends only on the limiting minimizer, η, and the spectrum. It gives the same ρ_η, with the finite prefix absorbed into random constants. No bound on the initial scale or additional radial randomization is needed. □

General ambient random initialization. When $d < n$, removing the range restriction is distinct from removing the normalization. For a general X_0, let $P = UU^\top$ and write $X_t = U\Lambda^{1/2}Z_t + N_t$, where $N_t = (I_n - P)X_t$. As long as the iterates are well-defined, their exact recursions are

$$\begin{aligned} N_t &= (1-\eta)^t N_0, \\ Z_{t+1} &= (1-\eta)Z_t + \eta\Lambda Z_t\big(Z_t^\top \Lambda Z_t + N_t^\top N_t\big)^{-1}. \end{aligned} \tag{163}$$

The extra term $N_t^\top N_t$ decays geometrically. We expect the proof to extend to general absolutely continuous random initialization by controlling this perturbation. This requires bounds on the inverse Gram matrices along the perturbed trajectory, which are not supplied by the range-supported argument above; the corollary does not assert this further extension.

C PROOFS :: RANK-k APPROXIMATION FOR ASYMMETRIC MATRICES

This appendix proves Theorem 6.1 through the more general Theorem C.1 below. We first prove decay of D_t conditional on a uniform positive lower bound on G_t. We then establish this bound and verify the initialization. Descent of g gives convergence to one critical point; saddle avoidance identifies the limit as a minimizer, and local contraction gives the convergence rate. The rank-one results are proved in §D.

The dilation objective separates as

$$g(Z) = \frac{1}{4}\|M - 2YX^\top\|_F^2 + \frac{1}{4}\|X^\top X - Y^\top Y\|_F^2 + \frac{1}{4}\|M\|_F^2. \tag{164}$$

Throughout, use the notation of §6: $M = U\Sigma V^\top$ has rank $r \geq k$, $\sigma_k > \sigma_{k+1}$, $\sigma_{r+1} := 0$, and $\kappa = \sigma_1/\sigma_r$. The range conditions $\mathrm{range}(X_t) \subseteq \mathrm{range}(V)$ and $\mathrm{range}(Y_t) \subseteq \mathrm{range}(U)$ are preserved by (26). Define the signed coordinates by

$$C_t = \frac{1}{\sqrt{2}}\Sigma^{-1/2}(V^\top X_t + U^\top Y_t), \qquad D_t = \frac{1}{\sqrt{2}}\Sigma^{-1/2}(V^\top X_t - U^\top Y_t). \tag{165}$$

The inverse linear change is

$$X_t = \frac{1}{\sqrt{2}}V\Sigma^{1/2}(C_t + D_t), \qquad Y_t = \frac{1}{\sqrt{2}}U\Sigma^{1/2}(C_t - D_t). \tag{166}$$

Their update is

$$\begin{aligned} C_{t+1} &= (1-\eta)C_t + \eta\Sigma C_t G_t^{-1}, \\ D_{t+1} &= (1-\eta)D_t - \eta\Sigma D_t G_t^{-1}, \end{aligned} \qquad G_t = C_t^\top \Sigma C_t + D_t^\top \Sigma D_t. \tag{167}$$

Denote this map by $\mathcal{T}$, and write $H_t := C_t^\top C_t - D_t^\top D_t$. When $D_0 = 0$, all $D_t = 0$ and the C_t update is the symmetric whitened iteration (37) with $\Lambda = \Sigma$.

Theorem C.1 (Asymmetric convergence and linear rate for the iterates)**.** *Fix $0 < \mu < 1$ and a deterministic $R > 0$. Assume range initialization $\mathrm{range}(X_0) \subseteq \mathrm{range}(V)$, $\mathrm{range}(Y_0) \subseteq \mathrm{range}(U)$ satisfying*

$$H_0 \succeq \mu I_k, \qquad R_0^2 := \sqrt{k}\left(\|\mathcal{H}(M)\|_F + 2\sqrt{g(Z_0)}\right) \leq R^2. \tag{168}$$

Choose a fixed step size $0 < \eta < \eta_{\max}$, where

$$\eta_{\max} := \min\left\{\frac{2\mu(1-\mu)}{\kappa + \mu(2-\mu)}, \quad \frac{\mu\sigma_r}{12R^2 + \sigma_1}\right\}. \tag{169}$$

Then every iterate is well-defined, the stacked iterate has full column rank, and the iterates converge to one critical point of g. If the initialization has a density relative to $\mathrm{range}(V) \times \mathrm{range}(U)$ and the preceding bounds hold almost surely, then, almost surely, the limits satisfy $2Y_\star X_\star^\top = M_k$ and

$$X_t \xrightarrow[t\to\infty]{} X_\star, \qquad Y_t \xrightarrow[t\to\infty]{} Y_\star. \tag{170}$$

With q_η defined in (28), there are almost-surely finite random constants $C_{XY}, C_M > 0$ such that, for every $t \geq 0$, the bounds (29) hold.

C.1 DECAY OF THE NEGATIVE COMPONENT

The next estimate controls D_t whenever G_t has a uniform positive lower bound.

Lemma C.2 (Geometric decay of the negative component)**.** *Let $0 < \eta < 1$ and $\delta > 0$. Suppose that $G_t \succeq \delta I_k$ for all $t \geq 0$ and that*

$$\eta\sigma_1 \leq 2(1-\eta)\delta. \tag{171}$$

Then the negative-component recursion in (167) satisfies

$$\|D_{t+1}\|_F \leq (1-\eta)\|D_t\|_F, \qquad \|D_t\|_F \leq (1-\eta)^t\|D_0\|_F. \tag{172}$$

Proof. **Step 1: expansion and control of the quadratic term.** Expansion of the update gives

$$\begin{aligned}\|D_{t+1}\|_F^2 = (1-\eta)^2\|D_t\|_F^2 - 2\eta(1-\eta)\operatorname{tr}(D_t^\top \Sigma D_t G_t^{-1}) \\ + \eta^2 \operatorname{tr}(G_t^{-1} D_t^\top \Sigma^2 D_t G_t^{-1}).\end{aligned} \tag{173}$$

Using $\Sigma^2 \preceq \sigma_1 \Sigma$, cyclicity of the trace, and $G_t^{-2} \preceq \delta^{-1} G_t^{-1}$, we obtain

$$\operatorname{tr}(G_t^{-1} D_t^\top \Sigma^2 D_t G_t^{-1}) \le \frac{\sigma_1}{\delta} \operatorname{tr}(D_t^\top \Sigma D_t G_t^{-1}). \tag{174}$$

Step 2: contraction of D_t. Substitution gives

$$\begin{aligned}\|D_{t+1}\|_F^2 &\le (1-\eta)^2\|D_t\|_F^2 - \eta\left(2(1-\eta) - \frac{\eta\sigma_1}{\delta}\right)\operatorname{tr}(D_t^\top \Sigma D_t G_t^{-1}) \\ &\le (1-\eta)^2\|D_t\|_F^2.\end{aligned} \tag{175}$$

The last inequality follows from (171) and nonnegativity of the trace. Taking square roots and iterating proves the claim. □

C.2 Proof of Theorem 6.1: well-posedness and initialization

We show that $H_0 \succeq \mu I_k$ yields a uniform lower bound on G_t, then verify this condition for the initialization in (27). The lower-bound argument does not use the preceding decay lemma.

Lemma C.3 (Preservation of $H_t \succeq \mu I_k$ and well-posedness)**.** *Fix $0 < \mu < 1$, assume the range conditions above and $H_0 \succeq \mu I_k$, and suppose*

$$0 < \eta \le \frac{2\mu(1-\mu)}{\kappa + 2\mu - \mu^2}. \tag{176}$$

Then every update is well-defined and, for every $t \ge 0$,

$$H_t \succeq \mu I_k, \qquad G_t \succeq \mu\sigma_r I_k, \qquad C_t^\top C_t \succeq \mu I_k. \tag{177}$$

Proof. **Step 1: recursion for H_t.** Write $\alpha := 1 - \eta$. Expanding (167) and subtracting the identities for $C_{t+1}^\top C_{t+1}$ and $D_{t+1}^\top D_{t+1}$ gives

$$H_{t+1} = \alpha^2 H_t + 2\alpha\eta I_k + \eta^2 G_t^{-1}(C_t^\top \Sigma^2 C_t - D_t^\top \Sigma^2 D_t) G_t^{-1}. \tag{178}$$

Since $\Sigma^2 \preceq \sigma_1 \Sigma$,

$$C_t^\top \Sigma^2 C_t - D_t^\top \Sigma^2 D_t \succeq -\sigma_1 D_t^\top \Sigma D_t \succeq -\sigma_1 G_t. \tag{179}$$

Step 2: preservation of the lower bound. If $H_t \succeq \mu I_k$, then

$$G_t \succeq \sigma_r(C_t^\top C_t + D_t^\top D_t) = \sigma_r(H_t + 2D_t^\top D_t) \succeq \mu\sigma_r I_k. \tag{180}$$

Consequently,

$$H_{t+1} \succeq \left((1-\eta)^2\mu + 2\eta(1-\eta) - \frac{\eta^2\kappa}{\mu}\right) I_k. \tag{181}$$

The scalar coefficient is at least μ precisely when $\eta(\kappa + 2\mu - \mu^2) \le 2\mu(1-\mu)$, as assumed. The lower bound at $t = 0$ therefore starts an induction: G_t is invertible and $H_{t+1} \succeq \mu I_k$ at every step. The third bound follows from $C_t^\top C_t = H_t + D_t^\top D_t$. □

Lemma C.4 (Initialization without singular values)**.** *Let $M \neq 0$ have rank $r \ge k$. Draw a standard Gaussian matrix $\Omega \in \mathbb{R}^{m\times k}$ and, independently, $A_X \in \mathbb{R}^{m\times k}$ and $A_Y \in \mathbb{R}^{n\times k}$ from absolutely continuous distributions supported in their unit Frobenius balls. Set*

$$Q = M^\top \Omega(\Omega^\top M M^\top \Omega)^{-1/2}, \qquad X_0 = Q + \frac{M^\top A_X}{4\|M\|_F}, \qquad Y_0 = M\left(Q + \frac{1}{4}A_Y\right). \tag{182}$$

Almost surely, Q is well-defined, $Q^\top Q = I_k$, and

$$\operatorname{range}(X_0) \subseteq \operatorname{range}(V), \qquad \operatorname{range}(Y_0) \subseteq \operatorname{range}(U), \qquad H_0 \succeq \frac{7}{8} I_k. \tag{183}$$

The pair (X_0, Y_0) has a density relative to the two ranges and satisfies

$$\|Z_0\|_F \le \left(\sqrt{k} + \frac{1}{4}\right)\sqrt{1 + \|M\|_F^2} =: L_0. \tag{184}$$

In particular, the initialization condition $H_0 \succeq \mu I_k$ holds with $\mu = 1/2$, without singular-value information.

Proof. **Step 1: well-posedness and the lower bound on H_0.** Since $U^\top\Omega$ is an $r \times k$ standard Gaussian matrix and $r \geq k$, $M^\top\Omega = V\Sigma U^\top\Omega$ has full column rank almost surely. Hence Q is well-defined, $Q^\top Q = I_k$, and $Q \in \mathrm{range}(V)$. The range conditions on X_0, Y_0 follow immediately. For iterates in these ranges, the signed-coordinate definitions give

$$H_0 = X_0^\top M^\dagger Y_0 + Y_0^\top (M^\dagger)^\top X_0. \tag{185}$$

The matrix $M^\dagger M$ is the orthogonal projector onto $\mathrm{range}(V)$. Since X_0 lies in that range, substitution of Y_0 gives the exact identity

$$H_0 = X_0^\top \left(Q + \frac{1}{4}A_Y\right) + \left(Q + \frac{1}{4}A_Y\right)^\top X_0. \tag{186}$$

Now $\|X_0 - Q\|_2 \leq \|M\|_2\|A_X\|_F/(4\|M\|_F) \leq 1/4$, so $\|X_0\|_2 \leq 5/4$. Using $Q^\top Q = I_k$ yields

$$\|H_0 - 2I_k\|_2 \leq 2\|X_0 - Q\|_2 + \frac{1}{2}\|X_0\|_2\|A_Y\|_2 \leq \frac{1}{2} + \frac{5}{8} = \frac{9}{8}. \tag{187}$$

Thus $H_0 \succeq (2 - 9/8)I_k = \frac{7}{8}I_k$.

Step 2: density on the two ranges. Conditional on Q, the maps $A_X \mapsto M^\top A_X/(4\|M\|_F)$ and $A_Y \mapsto MA_Y/4$ are surjective onto the respective ranges. Their inputs are independent and absolutely continuous, so the conditional joint law of (X_0, Y_0) has a density there. Integrating over Q preserves absolute continuity.

Step 3: a deterministic norm bound. The triangle inequality and $\|M\|_2 \leq \|M\|_F$ give

$$\|X_0\|_F \leq \sqrt{k} + \frac{1}{4}, \qquad \|Y_0\|_F \leq \|M\|_F\left(\sqrt{k} + \frac{1}{4}\right). \tag{188}$$

Squaring and adding proves the stated bound on $\|Z_0\|_F$. □

This bounded initialization permits a fixed step size for the entire law. Indeed, the triangle inequality gives

$$R_0^2 := \sqrt{k}\left(\|\mathcal{H}(M)\|_F + 2\sqrt{g(Z_0)}\right) \leq \sqrt{k}\left(2\|\mathcal{H}(M)\|_F + L_0^2\right) =: R^2. \tag{189}$$

Thus $\mu = 1/2$ and this deterministic R satisfy the initialization hypotheses of Theorem C.1. The construction itself uses no pseudoinverse or singular values.

C.3 Proof of Theorem 6.1: convergence to a critical point

We first quantify the perturbation of the symmetric map. We then prove convergence directly through descent of g and use the symmetric critical-point characterization to identify the limit.

The step-size bound (176) implies $\eta \leq 2\mu/(\kappa + 2\mu)$, since $(1 - \mu)(\kappa + 2\mu) \leq \kappa + 2\mu - \mu^2$. Thus Lemma C.3 supplies (171) with $\delta = \mu\sigma_r$, so Lemma C.2 applies along the trajectory.

For full-column-rank C, define

$$T_\eta(C) := (1 - \eta)C + \eta\Sigma C(C^\top\Sigma C)^{-1}. \tag{190}$$

This is the symmetric whitened map, with $\Lambda = \Sigma$.

Proposition C.5 (Geometrically perturbed symmetric dynamics)**.** *Assume the hypotheses of Lemma C.3. Then*

$$C_{t+1} = T_\eta(C_t) + R_t, \qquad R_t = -\eta\Sigma C_t(C_t^\top\Sigma C_t)^{-1}D_t^\top\Sigma D_t G_t^{-1}. \tag{191}$$

If additionally $\sup_t \|C_t\|_F < \infty$*, then*

$$\|R_t\|_F \leq c_R(1 - \eta)^{2t}, \qquad c_R := \frac{\eta\sigma_1^2}{\mu^2\sigma_r^2}\left(\sup_t \|C_t\|_F\right)\|D_0\|_F^2. \tag{192}$$

In particular, $\sum_{t=0}^\infty \|R_t\|_F < \infty$.

Proof. **Step 1: difference from the symmetric update.** Since $G_t = C_t^\top \Sigma C_t + D_t^\top \Sigma D_t$, the inverse identity gives

$$G_t^{-1} - (C_t^\top \Sigma C_t)^{-1} = -(C_t^\top \Sigma C_t)^{-1} D_t^\top \Sigma D_t G_t^{-1}. \tag{193}$$

Substitution into the C_t update proves (191).

Step 2: geometric decay of the difference. Lemma C.3 gives $C_t^\top \Sigma C_t \succeq \mu\sigma_r I_k$ and $G_t \succeq \mu\sigma_r I_k$. Hence

$$\|R_t\|_F \le \frac{\eta\sigma_1^2}{\mu^2\sigma_r^2} \|C_t\|_F \|D_t\|_F^2. \tag{194}$$

Apply Lemma C.2 and the assumed bound on C_t. The resulting geometric bound also proves summability. □

The next proposition establishes boundedness and convergence using the objective g; the preceding comparison identifies the limiting signed coordinates.

Proposition C.6 (Bounded descent and convergence to one critical point)**.** *Assume the range conditions and $H_0 \succeq \mu I_k$, where $0 < \mu < 1$. Let R_0 be as in (189), choose $R \ge R_0$, and fix*

$$0 < \eta < \min\left\{ \frac{2\mu(1-\mu)}{\kappa + 2\mu - \mu^2}, \frac{\mu\sigma_r}{12R^2 + \sigma_1} \right\}. \tag{195}$$

Then every iterate is full column rank, $\|Z_t\|_F \le R$, and

$$\sum_{t=0}^{\infty} \|Z_{t+1} - Z_t\|_F < \infty, \qquad Z_t \xrightarrow[t\to\infty]{} Z_\star, \qquad \nabla g(Z_\star) = 0. \tag{196}$$

Writing $(C_\star, D_\star)$ for its signed coordinates, we have $D_\star = 0$, $C_\star^\top C_\star = I_k$, and $\Sigma C_\star = C_\star(C_\star^\top \Sigma C_\star)$.

Proof. **Step 1: bounded descent.** Write $\delta := \mu\sigma_r$ and $L := 12R^2 + \sigma_1$. By Lemma C.3, $G_t \succeq \delta I_k$ for all t. On $\|Z\|_F \le 2R$, differentiation of the quartic objective gives

$$\nabla^2 g(Z)[E] = E(Z^\top Z) + ZE^\top Z + ZZ^\top E - \mathcal{H}(M)E, \qquad \|\nabla^2 g(Z)\|_2 \le L. \tag{197}$$

Here the operator norm is induced by the Frobenius norm. Suppose inductively that $g(Z_t) \le g(Z_0)$. Then

$$\|Z_t\|_F^2 \le \sqrt{k}\|Z_t Z_t^\top\|_F \le \sqrt{k}\left(\|\mathcal{H}(M)\|_F + 2\sqrt{g(Z_0)}\right) = R_0^2 \le R^2. \tag{198}$$

Moreover, $\|\nabla g(Z_t)\|_F \le (R^2 + \sigma_1)R$. Consequently,

$$\|Z_{t+1} - Z_t\|_F \le \frac{\eta}{\delta}(R^2 + \sigma_1)R < R, \tag{199}$$

since $\eta < \delta/L$ and $R^2 + \sigma_1 < L$. Thus the entire update segment lies in the ball of radius $2R$, before descent is invoked. Taylor's inequality yields

$$\begin{aligned} g(Z_{t+1}) &\le g(Z_t) - \eta\|\nabla g(Z_t) G_t^{-1/2}\|_F^2 + \frac{\eta^2 L}{2}\|\nabla g(Z_t) G_t^{-1}\|_F^2 \\ &\le g(Z_t) - \frac{\eta}{2}\|\nabla g(Z_t) G_t^{-1/2}\|_F^2. \end{aligned} \tag{200}$$

This closes the induction and proves boundedness.

Step 2: finite length and a single limit. Put $\Delta Z_t := Z_{t+1} - Z_t$. Since $\Delta Z_t = -\eta \nabla g(Z_t) G_t^{-1}$ and $\delta I_k \preceq G_t \preceq R^2 I_k$, (200) implies

$$g(Z_t) - g(Z_{t+1}) \ge \frac{\delta}{2\eta}\|\Delta Z_t\|_F^2. \tag{201}$$

The Lipschitz gradient bound on the update segment also gives

$$\|\nabla g(Z_{t+1})\|_F \le \left(L + \frac{R^2}{\eta}\right)\|\Delta Z_t\|_F. \tag{202}$$

These verify conditions H1 and H2 of Attouch et al. (2013, Theorem 2.9) with fixed constants $a = \delta/(2\eta) > 0$ and $b = L + R^2/\eta > 0$, respectively; in this smooth setting the limiting subdifferential is $\{\nabla g\}$. The polynomial g is finite, continuous, and bounded below, so it is proper and lower semicontinuous and has the Kurdyka–Łojasiewicz property. Boundedness supplies a subsequence converging to $\widehat{Z}$, and continuity gives $g(Z_{t_j}) \to g(\widehat{Z})$, which is H3. That theorem therefore gives finite length and convergence to one critical point. The uniform lower bound on G_t passes to the limit, so $Z_\star$ has full column rank.

Step 3: identification in signed coordinates. By Lemma C.2, $D_\star = 0$. Boundedness of Z_t implies boundedness of C_t by the invertible linear change of coordinates. Thus Proposition C.5 gives $R_t \to 0$, and passing to the limit yields $C_\star = T_\eta(C_\star)$. The bound $C_\star^\top C_\star \succeq \mu I_k$ makes $C_\star$ full rank. Apply Lemma A.8 with $\Lambda = \Sigma$ to conclude. □

C.4 Proof of Theorem 6.1: saddle avoidance

We apply the general criterion in Proposition B.7 to the full signed map $\mathcal{T}$. The symmetric Hessian supplies an unstable direction at every nonminimizing critical point.

Proposition C.7 (Almost-sure avoidance of nonminimizing limits)**.** *Fix $0 < \eta < 1$. The initial signed pairs whose iterates converge to a nonminimizing full-rank critical point of g form a Lebesgue-null set in $\mathbb{R}^{r\times k} \times \mathbb{R}^{r\times k}$.*

Proof. **Step 1: regularity on a full-measure domain.** Consider $\mathcal{T}$ on the open set $\mathcal{U} := \{(C, D) : C^\top \Sigma C + D^\top \Sigma D \succ 0\}$. Its complement is null, since the determinant is a nonzero polynomial. Extend $\mathcal{T}$ by zero outside $\mathcal{U}$ to obtain a map on the whole space. On $\mathcal{U}$ it is real analytic, hence continuously differentiable with locally Lipschitz derivative.

Step 2: almost-everywhere invertibility of the Jacobian. For any $\xi = (C, D) \in \mathcal{U}$ and $a > 0$, scaling the inverse in the update and differentiating gives

$$\mathrm{D}\mathcal{T}(a\xi) = (1-\eta)I_{2rk} + a^{-2}\big(\mathrm{D}\mathcal{T}(\xi) - (1-\eta)I_{2rk}\big). \tag{203}$$

Thus its determinant tends to $(1-\eta)^{2rk} \neq 0$ as $a \to \infty$. Positive scaling stays in the same connected component of $\mathcal{U}$. On each component, the determinant is therefore a real-analytic function that is not identically zero. There are at most countably many components. The analytic zero-set theorem (Mityagin, 2020) then shows that $\mathrm{D}\mathcal{T}$ is nonsingular on an open set with null complement in the whole space.

Step 3: instability of nonminimizing fixed points. At a full-rank critical point, the signed fixed-point equations are $\Sigma C = CG$ and $\Sigma D = -DG$. Taking the Frobenius inner product of the second equation with D gives $D = 0$, since $\Sigma, G \succ 0$. The first equation then identifies C as a critical point of the symmetric hidden objective $h(C) = \frac{1}{2}\|C\|_F^2 - \frac{1}{2}\log(\det(C^\top \Sigma C))$. At $(C_\star, 0)$ the cross derivatives vanish, and the two diagonal blocks are

$$E \longmapsto E - \eta\nabla^2 h(C_\star)[E], \qquad F \longmapsto (1-\eta)F - \eta\Sigma F(C_\star^\top \Sigma C_\star)^{-1}. \tag{204}$$

Both blocks are self-adjoint for the Frobenius inner product: the first is the identity minus a Hessian, and the second uses left and right multiplication by symmetric matrices. The cited criterion is applied in these signed Euclidean coordinates; no symmetry of the physical-coordinate Jacobian is required. If the critical point is nonminimizing, $\mathrm{range}(C_\star)$ contains an eigenvector for some σ_i while omitting an eigenvector for a larger σ_j. The corresponding positive-block multiplier is $1 + \eta(\sigma_j/\sigma_i - 1) > 1$ by Lemma A.9. All hypotheses of Proposition B.7 therefore hold. That proposition gives the asserted null set. The invertible linear change of coordinates transfers this conclusion to any absolutely continuous initialization on the two physical ranges. □

C.5 Proof of Theorem 6.1: local rate and conclusion

The preceding propositions give almost-sure convergence to a minimizer. We now compute the derivative there, establish local contraction, and convert it into the iterate and reconstruction bounds.

Lemma C.8 (Local multipliers and contraction toward minimizers)**.** *Let $\nu_k := (\sigma_k - \sigma_{k+1})/\sigma_k$ and $\mathcal{S}_\star := \{(E_k R, 0) : R \in O(k)\}$, where $E_k = (I_k, 0)^\top \in \mathbb{R}^{r\times k}$. The normal multipliers of the signed*

update at this manifold are

$$1-2\eta, \qquad 1-\eta\left(1-\frac{\sigma_j}{\sigma_i}\right) \quad (i\le k<j), \qquad 1-\eta\left(1+\frac{\sigma_j}{\sigma_i}\right) \quad (i\le k,\ j\le r). \tag{205}$$

All normal multipliers have magnitude strictly below 1 *precisely when*

$$0<\eta<\frac{2}{1+\sigma_1/\sigma_k}. \tag{206}$$

Under the stronger condition $0<\eta<1/(1+\kappa)$*, there is a forward-invariant neighborhood* $\mathcal{N}$ *of* $\mathcal{S}_\star$ *such that*

$$\operatorname{dist}\big(\mathcal{T}(C,D),\mathcal{S}_\star\big)\le\left(1-\frac{\eta\nu_k}{2}\right)\operatorname{dist}\big((C,D),\mathcal{S}_\star\big), \qquad (C,D)\in\mathcal{N}. \tag{207}$$

Here $\operatorname{dist}(\xi,\mathcal{S}_\star):=\inf_{\zeta\in\mathcal{S}_\star}\|\xi-\zeta\|_F$*, with the product Frobenius norm on pairs.*

Proof. **Step 1: normal multipliers and the stability range.** At $(E_k,0)$, (204) and Lemma A.9 give the multipliers for perturbations of C. Since $E_k^\top\Sigma E_k=\Sigma_k$, the negative block acts on its (j,i) entry by $1-\eta(1+\sigma_j/\sigma_i)$. The remaining multipliers equal 1 and correspond exactly to the tangent directions $(E_k\Omega,0)$ with $\Omega^\top=-\Omega$. The multiplier $1-2\eta$ has magnitude less than 1 for $0<\eta<1$. The multipliers involving omitted positive eigenvalues also lie in $(-1,1)$ throughout this interval. The most restrictive negative mode has $j=1,i=k$, giving (206), whose upper endpoint is at most 1.

Step 2: contraction near the minimizing manifold. Now assume $0<\eta<1/(1+\kappa)$. All normal multipliers are nonnegative and at most $1-\eta\nu_k$. The derivative is self-adjoint by (204), so its operator norm restricted to the normal space is at most $1-\eta\nu_k$, since $0<\nu_k\le 1$ and $\sigma_j/\sigma_i\le\kappa$. This includes $k=r$, where there are no omitted positive directions and $\nu_k=1$. For $\xi=(C,D)$ sufficiently close to the compact smooth manifold $\mathcal{S}_\star$, let $\pi(\xi)$ be its unique closest point and put $e:=\xi-\pi(\xi)$. This displacement is normal at $\pi(\xi)$. Every point of $\mathcal{S}_\star$ is fixed, so Taylor expansion gives

$$\mathcal{T}(\xi)-\pi(\xi)=\mathrm{D}\mathcal{T}(\pi(\xi))[e]+O(\|e\|_F^2). \tag{208}$$

The bound on the remainder is uniform along the compact manifold. Right-orthogonal equivariance transfers the normal derivative bound from $(E_k,0)$ to every $\pi(\xi)$. Therefore

$$\operatorname{dist}\big(\mathcal{T}(\xi),\mathcal{S}_\star\big)\le\|\mathcal{T}(\xi)-\pi(\xi)\|_F\le(1-\eta\nu_k)\|e\|_F+O(\|e\|_F^2). \tag{209}$$

Choose a sufficiently small tubular neighborhood so the quadratic term is at most $(\eta\nu_k/2)\|e\|_F$. This proves the contraction. A smaller tube lies in the full-rank domain, and the contraction makes it forward invariant. □

The multipliers associated with D explain the restriction on η: when $\sigma_1>\sigma_k$, some step sizes in $(0,1)$ already cause local instability. The global bound (195) also guarantees well-posedness and descent along the entire trajectory.

Proof of Theorems 6.1 and C.1. **Step 1: almost-sure convergence to a minimizer.** Fix η satisfying (195), with the deterministic R in the theorem. For the initialization in (27), Lemma C.4 and (189) supply these hypotheses with $\mu=1/2$. By Proposition C.6, every trajectory converges to one full-rank critical point. Under the density assumption, Proposition C.7 identifies it almost surely as a minimizer. Thus, for some random $R_\star\in O(k)$,

$$C_t\xrightarrow[t\to\infty]{}E_kR_\star, \qquad D_t\xrightarrow[t\to\infty]{}0, \tag{210}$$

and (166) gives the stated limits $X_\star,Y_\star$, with $2Y_\star X_\star^\top=M_k$.

Step 2: local contraction and convergence rate. The first bound in (195) is strictly less than $1/(1+\kappa)$ because

$$\kappa+2\mu-\mu^2-2\mu(1-\mu)(1+\kappa)=\kappa\big((1-\mu)^2+\mu^2\big)+\mu^2>0. \tag{211}$$

Thus Lemma C.8 applies. Write $q_\eta = 1 - \eta\nu_k/2$ as in (28), $\xi_t := (C_t, D_t)$, and $\xi_\star := (E_k R_\star, 0)$. There is an almost-surely finite entrance time T into $\mathcal{N}$, and

$$\operatorname{dist}(\xi_t, \mathcal{S}_\star) \le q_\eta^{t-T} \operatorname{dist}(\xi_T, \mathcal{S}_\star), \qquad t \ge T. \tag{212}$$

Because $\mathcal{T}$ fixes $\mathcal{S}_\star$ and has bounded derivative on a smaller tubular neighborhood, there is a finite $L_\mathcal{T}$ such that

$$\|\xi_{t+1} - \xi_t\|_F \le L_\mathcal{T} \operatorname{dist}(\xi_t, \mathcal{S}_\star), \qquad t \ge T, \tag{213}$$

enlarging T if needed. Summing the geometric tail gives

$$\|\xi_t - \xi_\star\|_F \le \frac{L_\mathcal{T}}{1 - q_\eta} \operatorname{dist}(\xi_T, \mathcal{S}_\star) q_\eta^{t-T}. \tag{214}$$

Step 3: iterate and reconstruction rates. The linear coordinate map (166) has operator norm at most $\sqrt{\sigma_1}$ between the product Frobenius spaces. Cauchy–Schwarz therefore gives

$$\|X_t - X_\star\|_F + \|Y_t - Y_\star\|_F \le \sqrt{2\sigma_1}\|\xi_t - \xi_\star\|_F. \tag{215}$$

The preceding estimate bounds this by a finite constant times q_η^{t-T}. Finally,

$$\|2Y_t X_t^\top - M_k\|_F \le 2\|Y_t\|_2\|X_t - X_\star\|_F + 2\|X_\star\|_2\|Y_t - Y_\star\|_F. \tag{216}$$

Boundedness of the iterates proves the reconstruction rate. Absorbing the finite initial segment and q_η^{-T} into the constants gives both bounds for every $t \ge 0$.

In addition, the gradient of g vanishes on the minimizing manifold, and its Hessian is bounded nearby. Taylor expansion about a closest minimizer therefore gives

$$g(Z_t) - g(Z_\star) \le C_g q_\eta^{2t}, \qquad t \ge 0, \tag{217}$$

after increasing the finite constant C_g to cover the initial segment. □

As in Theorem 3.1, the constants absorb the finite entrance time T. The proof does not provide a uniform bound on T over the initialization law.

D PROOFS :: RANK-1 APPROXIMATION

We prove the gap-dependent and gap-free rates in Theorems 7.1 and 7.2 in four separate subsections. Each rate is established in a proposition using coordinate-ratio estimates stated as lemmas. The symmetric gap-dependent proof uses the whitened iterate; the gap-free proofs use the original symmetric iterate or the stacked rectangular iterate. The two gap-free propositions share the estimates in §D.2. Throughout this appendix, the scalar version of (87) is denoted by

$$\phi_\eta(w) := \left((1-\eta)\sqrt{w} + \frac{\eta}{\sqrt{w}}\right)^2, \qquad w > 0. \tag{218}$$

D.1 PROOF OF THEOREM 7.1(I): GAP-DEPENDENT RATE

Assume $\lambda_1 > \lambda_2 \ge \cdots \ge \lambda_d > 0$, $d \ge 2$, $0 < \eta < 1$, $1 < \|z_0\|^2 < \beta$ and $z_{0,1} \ne 0$. Define

$$\rho := 1 - \frac{\min\{(1-\eta)\underline{m}_\eta, \eta/2\}}{(1-\eta)\overline{m}_{\eta,\beta} + \eta} \frac{\lambda_1 - \lambda_2}{\lambda_1} \in (0,1), \tag{219}$$

Here $\underline{m}_\eta, \overline{m}_{\eta,\beta}$ are as in Lemma 4.1. Write $\kappa = \lambda_1/\lambda_d$ and

$$w_t := \|z_t\|^2, \qquad q_t := z_t/\|z_t\|, \qquad e_t := \frac{\sum_{i=2}^d z_{t,i}^2}{z_{t,1}^2}. \tag{220}$$

For $k = 1$, write the quantities from §4 as $h_t = q_t^\top \Lambda q_t$ and $r_t = (I_d - q_t q_t^\top)\Lambda q_t$. Here h_t denotes the Rayleigh quotient; the objective value is written $h(z_t)$. Lemma B.1 gives

$$\underline{m}_\eta \le w_t \le \overline{m}_{\eta,\beta}, \qquad t \ge 0. \tag{221}$$

We control e_t and $(\log(w_t))^2$ using $2\rho - 1 \in (0, \rho)$ as a common bound for the intermediate contraction factors. The larger factor ρ allows us to combine these estimates despite their coupling.

Lemma D.1 (Contraction of the squared coordinate ratios)**.** *Under the preceding assumptions,* $z_{t,1} \neq 0$ *and*

$$e_{t+1} \le (2\rho - 1)^2 e_t, \qquad e_t \le (2\rho - 1)^{2t} e_0 \quad (t \ge 0). \tag{222}$$

Proof. All coordinate multipliers in the update are positive. For $i > 1$,

$$\left|\frac{z_{t+1,i}}{z_{t+1,1}}\right| = \left(1 - \frac{\eta(\lambda_1 - \lambda_i)}{(1-\eta) z_t^\top \Lambda z_t + \eta\lambda_1}\right) \left|\frac{z_{t,i}}{z_{t,1}}\right|. \tag{223}$$

Since $z_t^\top \Lambda z_t \le \lambda_1 \overline{m}_{\eta,\beta}$, the multiplier is at most

$$1 - \frac{\eta}{(1-\eta)\overline{m}_{\eta,\beta} + \eta} \frac{\lambda_1 - \lambda_2}{\lambda_1} \le 2\rho - 1 < 1. \tag{224}$$

Squaring, summing over $i > 1$, and iterating proves the claim. □

Lemma D.2 (Control of $(\log(\|z_t\|^2))^2$)**.** *Under the same assumptions, define*

$$D := \frac{2\eta^2(\kappa - 1)^2}{\underline{m}_\eta^2} \max\{|\log(\underline{m}_\eta)|, \log(\overline{m}_{\eta,\beta})\}. \tag{225}$$

Then, for every $t \ge 0$*,*

$$(\log(w_{t+1}))^2 \le (2\rho - 1)^2 (\log(w_t))^2 + D e_t. \tag{226}$$

Proof. **Step 1: the unperturbed update.** Use ϕ_η from (218). The estimate (124) and the definition of ρ give

$$\rho_{\eta,\beta} \le 1 - \frac{2\min\{(1-\eta)\underline{m}_\eta, \eta\}}{(1-\eta)\overline{m}_{\eta,\beta} + \eta} \le 2\rho - 1, \qquad |\log(\phi_\eta(w_t))| \le (2\rho - 1)|\log(w_t)|. \tag{227}$$

Here $(\lambda_1 - \lambda_2)/\lambda_1 \le 1$ and $0 < 2\rho - 1 < 1$.

Step 2: the correction to $\phi_\eta(w_t)$**.** Specializing (90) to $k = 1$ gives

$$w_{t+1} - \phi_\eta(w_t) = \frac{\eta^2}{w_t} \frac{\|r_t\|^2}{h_t^2} \ge 0. \tag{228}$$

Since the weighted mean h_t minimizes squared deviations,

$$\|r_t\|^2 = \sum_i q_{t,i}^2 (\lambda_i - h_t)^2 \le \sum_i q_{t,i}^2 (\lambda_i - \lambda_1)^2 \le (\lambda_1 - \lambda_d)^2 e_t. \tag{229}$$

Thus $w_{t+1} - \phi_\eta(w_t) \le \eta^2(\kappa - 1)^2 e_t / \underline{m}_\eta$. Both arguments belong to $[\underline{m}_\eta, \overline{m}_{\eta,\beta}]$ by the proof of Lemma B.4. On this interval the absolute derivative of $(\log(w))^2$ is at most $2\max\{|\log(\underline{m}_\eta)|, \log(\overline{m}_{\eta,\beta})\}/\underline{m}_\eta$. The mean value theorem and Step 1 prove the claim. □

The following proposition also controls $h(z_t) - h_\star$. To obtain this additional conclusion, we combine the two preceding estimates with the decomposition of h in the potential $\mathcal{L}$ below.

Proposition D.3 (Symmetric rank-one approximation with a strict gap)**.** *Under the preceding assumptions, the bounds in* (31) *hold with the explicit constants in* (240)–(241)*. Moreover, there are explicit* $a, b > 0$ *such that*

$$\mathcal{L}(z) := h(z) - h_\star + a(\log(\|z\|^2))^2 + b\frac{\sum_{i=2}^d z_i^2}{z_1^2}, \qquad \mathcal{L}(z_{t+1}) \le \rho^2 \mathcal{L}(z_t). \tag{230}$$

In particular,

$$h(z_t) - h_\star \le C\rho^{2t}(h(z_0) - h_\star), \qquad C := 1 + \max\left\{4a, \frac{2b\lambda_1(1 + e_0)}{\lambda_1 - \lambda_2}\right\}. \tag{231}$$

Proof. **Step 1: comparison with e_t and $(\log(w_t))^2$.** The exact decomposition is

$$h(z_t) - h_\star = \frac{1}{2}(w_t - \log(w_t) - 1) + \frac{1}{2}\log\left(\frac{\lambda_1}{h_t}\right). \tag{232}$$

Taylor's formula for e^u on the interval between 0 and $\log(w_t)$, and $w_t \le \overline{m}_{\eta,\beta}$, bound the first term by $\overline{m}_{\eta,\beta}(\log(w_t))^2/4$. For the second, $\log(\lambda_1/h_t) \le (\lambda_1 - h_t)/h_t \le (\kappa - 1)e_t$. Consequently,

$$0 \le h(z_t) - h_\star \le \frac{\overline{m}_{\eta,\beta}}{4}(\log(w_t))^2 + \frac{\kappa - 1}{2}e_t. \tag{233}$$

Step 2: contraction of $\mathcal{L}$. Choose

$$a := 1 + \frac{\overline{m}_{\eta,\beta}(2\rho - 1)^2}{4(\rho^2 - (2\rho - 1)^2)}, \qquad b := 1 + \frac{(\overline{m}_{\eta,\beta}/4 + a)D + (\kappa - 1)(2\rho - 1)^2/2}{\rho^2 - (2\rho - 1)^2}. \tag{234}$$

The denominator is positive. Combining Step 1 with Lemmas D.1 and D.2 yields

$$\mathcal{L}(z_{t+1}) \le (\overline{m}_{\eta,\beta}/4 + a)(2\rho - 1)^2(\log(w_t))^2 \tag{235}$$

$$+ \big((\overline{m}_{\eta,\beta}/4 + a)D + ((\kappa - 1)/2 + b)(2\rho - 1)^2\big)e_t \tag{236}$$

$$\le \rho^2\big(a(\log(w_t))^2 + be_t\big) \le \rho^2\mathcal{L}(z_t). \tag{237}$$

Since $w_0 > 1$, the two terms in (232) give

$$h(z_0) - h_\star \ge \frac{1}{4}(\log(w_0))^2 + \frac{\lambda_1 - \lambda_2}{2\lambda_1}\frac{e_0}{1 + e_0}. \tag{238}$$

Hence $\mathcal{L}(z_0) \le C(h(z_0) - h_\star)$ with C as stated, which proves (231).

Step 3: iterate and reconstruction bounds. The leading coordinate keeps its sign. For $z_\star = \operatorname{sign}(z_{0,1})e_1$, $\|q_t - z_\star\|^2 = 2(1 - |q_{t,1}|) \le e_t$ and $|\sqrt{w_t} - 1| \le \sqrt{\overline{m}_{\eta,\beta}}|\log(w_t)|/2$. The contraction of $\mathcal{L}$ therefore gives

$$\|z_t - z_\star\| \le \left(\frac{\sqrt{\overline{m}_{\eta,\beta}}}{2\sqrt{a}} + \frac{1}{\sqrt{b}}\right)\sqrt{\mathcal{L}(z_0)}\rho^t. \tag{239}$$

Apply $x_t = U\Lambda^{1/2}z_t$, whose operator norm is $\sqrt{\lambda_1}$, to obtain

$$K_X := \sqrt{\lambda_1\mathcal{L}(z_0)}\left(\frac{\sqrt{\overline{m}_{\eta,\beta}}}{2\sqrt{a}} + \frac{1}{\sqrt{b}}\right), \qquad x_\star = \operatorname{sign}(z_{0,1})\sqrt{\lambda_1}u_1. \tag{240}$$

Finally, $\|x_t\| \le \sqrt{\lambda_1\overline{m}_{\eta,\beta}}$ and $\|x_\star\| = \sqrt{\lambda_1}$ imply

$$\|x_tx_t^\top - x_\star x_\star^\top\|_F \le (\|x_t\| + \|x_\star\|)\|x_t - x_\star\| \le K_M\rho^t, \qquad K_M := \sqrt{\lambda_1}(\sqrt{\overline{m}_{\eta,\beta}} + 1)K_X. \tag{241}$$

This proves Theorem 7.1(i). □

D.2 Proof of Theorem 7.1(ii): gap-free rate

The next two lemmas also apply to the dilation $\mathcal{H}(M)$ and will be reused in §D.4.

Coordinates for the gap-free estimates. In this subsection, z_t denotes the iterate of the general update below, not the whitened iterate of §D.1. For the symmetric problem we set $A = M$ and $z_t = x_t$; for the rectangular problem we set $A = \mathcal{H}(M)$ and $z_t = (y_t^\top, x_t^\top)^\top$. Let $A \in \mathbb{S}^N$ satisfy $\lambda := \lambda_{\max}(A) = \|A\|_2 > 0$. Fix $0 < \eta < 1$ and consider

$$z_{t+1} = \alpha z_t + \eta\frac{Az_t}{\|z_t\|^2}, \qquad \alpha := 1 - \eta. \tag{242}$$

Let $P_\star$ be the orthogonal projector onto the full eigenspace of A associated with λ. The top eigenvalue need not be simple. Assume $P_\star z_0 \neq 0$ and put

$$a_0 := \|P_\star z_0\| > 0, \qquad s_0 := \frac{\|z_0\|^2}{\lambda}, \qquad e_0 := \frac{\|(I_N - P_\star)z_0\|^2}{a_0^2}. \tag{243}$$

Define the bounds on $\|z_t\|^2/\lambda$ by

$$\underline{s} := \min\left\{s_0, \frac{4\alpha\eta}{1 + e_0}\right\}, \qquad \overline{s} := \max\left\{s_0, 1, \frac{2\alpha\eta + \eta^2/\underline{s}}{1 - \alpha^2}\right\}. \tag{244}$$

Lemma D.4 (Coordinate ratios and the Rayleigh quotient)**.** *Under the preceding assumptions, every iterate is well-defined and $\underline{s} \le \|z_t\|^2/\lambda \le \overline{s}$ for all $t \ge 0$. Set*

$$c := \frac{\eta}{\alpha\overline{s}+\eta}, \qquad q_- := \max\left\{\frac{\alpha\overline{s}}{\alpha\overline{s}+\eta}, \frac{|\alpha\underline{s}-\eta|}{\alpha\underline{s}+\eta}\right\}. \tag{245}$$

Both constants lie in $(0,1)$. For an orthonormal eigenbasis $Ap_j = \lambda\mu_j p_j$, put $a_t = \|P_\star z_t\|$ and $b_{j,t} = p_j^\top z_t$. Then

$$\frac{|b_{j,t}|}{a_t} \le \frac{|b_{j,0}|}{a_0}\begin{cases}(1-c(1-\mu_j))^t, & 0 \le \mu_j < 1,\\ q_-^t, & -1 \le \mu_j < 0.\end{cases} \tag{246}$$

Consequently, with $\varepsilon_t := 1 - z_t^\top A z_t/(\lambda\|z_t\|^2)$, we have

$$0 \le \varepsilon_t \le \frac{C_{\mathrm{ang}}}{t+1}, \qquad C_{\mathrm{ang}} := e_0\left(\frac{1}{c} + \frac{2}{1-q_-^2}\right). \tag{247}$$

Proof. **Step 1: nonzero projection and bounds on s_t.** For every well-defined iterate put

$$a_t := \|P_\star z_t\|, \qquad s_t := \frac{\|z_t\|^2}{\lambda}, \qquad e_t := \frac{\|(I_N - P_\star)z_t\|^2}{a_t^2}. \tag{248}$$

Projection of (242) onto the leading eigenspace gives

$$P_\star z_{t+1} = \left(\alpha + \frac{\eta}{s_t}\right) P_\star z_t, \qquad a_{t+1} = \left(\alpha + \frac{\eta}{s_t}\right) a_t > 0. \tag{249}$$

Induction proves $a_t > 0$, hence $\|z_t\| > 0$, for all t. Let $\{p_j\}$ be an orthonormal eigenbasis of A, write $Ap_j = \lambda\mu_j p_j$ with $\mu_j \in [-1,1]$, and put $b_{j,t} := p_j^\top z_t$. For each direction outside the leading eigenspace, $\mu_j < 1$ and

$$\left|\frac{b_{j,t+1}}{a_{t+1}}\right| = \frac{|\alpha s_t + \eta\mu_j|}{\alpha s_t + \eta}\left|\frac{b_{j,t}}{a_t}\right| \le \left|\frac{b_{j,t}}{a_t}\right|. \tag{250}$$

Therefore $e_t \le e_0$. Using $\|z_t\|^2 = a_t^2(1+e_t)$ gives

$$\frac{a_{t+1}}{\sqrt{\lambda}} = \alpha\frac{a_t}{\sqrt{\lambda}} + \frac{\eta}{(a_t/\sqrt{\lambda})(1+e_t)}, \qquad \frac{a_{t+1}^2}{\lambda} \ge \frac{4\alpha\eta}{1+e_t} \ge \frac{4\alpha\eta}{1+e_0}. \tag{251}$$

Consequently $s_t \ge \underline{s}$ for all $t \ge 0$. Since $z_t^\top A z_t \le \lambda\|z_t\|^2$ and $\|Az_t\|^2 \le \lambda^2\|z_t\|^2$,

$$s_{t+1} \le \alpha^2 s_t + 2\alpha\eta + \frac{\eta^2}{s_t} \le \alpha^2 s_t + 2\alpha\eta + \frac{\eta^2}{\underline{s}}. \tag{252}$$

The definition of $\overline{s}$ closes an induction proving

$$\underline{s} \le s_t \le \overline{s}, \qquad t \ge 0. \tag{253}$$

Step 2: a gap-free Rayleigh-quotient bound. Define the normalized spectral deficit

$$\varepsilon_t := 1 - \frac{z_t^\top A z_t}{\lambda\|z_t\|^2} = \sum_{j:\mu_j<1}(1-\mu_j)\frac{b_{j,t}^2}{\|z_t\|^2} \ge 0. \tag{254}$$

For $0 \le \mu_j < 1$, (250) and (253) imply

$$\frac{b_{j,t}^2}{a_t^2} \le \frac{b_{j,0}^2}{a_0^2}\big(1-c(1-\mu_j)\big)^{2t}. \tag{255}$$

For $0 \le d \le 1$ and $t \ge 0$, maximizing $x(1-x)^{2t}$ on $[0,1]$ with $x = cd$ gives

$$d(1-cd)^{2t} \le \frac{1}{c(2t+1)}. \tag{256}$$

For $-1 \le \mu_j < 0$, the multiplier in (250) is at most q_-. Indeed, its maximum over $\mu_j \in [-1,0]$ occurs at an endpoint. The endpoint 0 gives a ratio increasing with s_t, hence bounded by $\alpha\overline{s}/(\alpha\overline{s}+\eta)$.

The endpoint -1 gives a ratio decreasing up to $s_t = \eta/\alpha$ and increasing afterwards. Its value at $\overline{s}$ is either no larger than its value at $\underline{s}$, or bounded by the endpoint-0 ratio. Thus $b_{j,t}^2/a_t^2 \le q_-^{2t} b_{j,0}^2/a_0^2$ for negative modes. Using $\|z_t\|^2 \ge a_t^2$ and $1-\mu_j \le 2$, we obtain

$$\varepsilon_t \le \frac{e_0}{c(2t+1)} + 2e_0 q_-^{2t} \le \frac{C_{\mathrm{ang}}}{t+1}. \tag{257}$$

The last inequality uses $(t+1)q_-^{2t} \le \sum_{j=0}^{t} q_-^{2j} \le (1-q_-^2)^{-1}$.

□

For the next estimate, let ρ denote the contraction factor of $\log(\phi_\eta)$ in logarithmic coordinates on the interval above:

$$\rho := \max_{u\in[\underline{s},\overline{s}]} \frac{|\alpha u - \eta|}{\alpha u + \eta} < 1. \tag{258}$$

$$C_{\log} := \frac{|\log(s_0)|}{1-\rho} + \frac{2\eta}{\underline{s}}\left(\alpha + \frac{\eta}{\underline{s}}\right) C_{\mathrm{ang}} \left(\frac{1}{1-\rho} + \frac{1}{(1-\rho)^2}\right). \tag{259}$$

Lemma D.5 (Control of $\log(\|z_t\|^2/\lambda)$)**.** *Under the assumptions of Lemma D.4, write $s_t = \|z_t\|^2/\lambda$. For every $t \ge 0$,*

$$|\log(s_t)| \le \frac{C_{\log}}{t+1}, \qquad |s_t - 1| \le \frac{\overline{s} C_{\log}}{t+1}. \tag{260}$$

Proof. **Step 1: recursion for $\log(s_t)$.** Write $s_t = \|z_t\|^2/\lambda$ and ε_t as in Lemma D.4. Use ϕ_η from (218). Expanding the update gives

$$s_{t+1} = \alpha^2 s_t + 2\alpha\eta(1-\varepsilon_t) + \frac{\eta^2 \|Az_t\|^2}{\lambda^2 \|z_t\|^2 s_t}. \tag{261}$$

The inequalities $0 \le 1-\mu_j^2 \le 2(1-\mu_j)$ imply $0 \le 1 - \|Az_t\|^2/(\lambda^2\|z_t\|^2) \le 2\varepsilon_t$. Consequently,

$$0 \le \phi_\eta(s_t) - s_{t+1} \le 2\eta\left(\alpha + \frac{\eta}{\underline{s}}\right)\varepsilon_t. \tag{262}$$

Both s_{t+1} and $\phi_\eta(s_t)$ are at least $\underline{s}$: for the latter, $\phi_\eta(s_t) \ge 4\alpha\eta \ge \underline{s}$. The logarithm is therefore $1/\underline{s}$-Lipschitz between them. Also, the derivative of $x \mapsto \log(\phi_\eta(e^x))$ is $(\alpha e^x - \eta)/(\alpha e^x + \eta)$; its absolute value is at most ρ on $e^x \in [\underline{s}, \overline{s}]$. This interval contains 1 and $\phi_\eta(1) = 1$, so

$$|\log(s_{t+1})| \le \rho |\log(s_t)| + \frac{2\eta}{\underline{s}}\left(\alpha + \frac{\eta}{\underline{s}}\right)\frac{C_{\mathrm{ang}}}{t+1}. \tag{263}$$

Step 2: summing the recursion. For $t \ge 1$ the convolution obeys

$$\sum_{j=0}^{t-1} \frac{\rho^{t-1-j}}{j+1} = \sum_{k=0}^{t-1} \frac{\rho^k}{t-k} \le \frac{1}{t+1}\sum_{k=0}^{\infty}(k+2)\rho^k = \frac{1}{t+1}\left(\frac{1}{1-\rho} + \frac{1}{(1-\rho)^2}\right). \tag{264}$$

Here $(t+1)/(t-k) \le k+2$ for $0 \le k < t$. Together with $(t+1)\rho^t \le (1-\rho)^{-1}$, iteration of (263) gives

$$|\log(s_t)| \le \frac{C_{\log}}{t+1}, \qquad |s_t - 1| \le \overline{s}|\log(s_t)| \le \frac{\overline{s}\, C_{\log}}{t+1}, \qquad t \ge 0. \tag{265}$$

□

For the reconstruction bounds, set

$$C_1 := 2\lambda^2 \overline{s}\, C_{\mathrm{ang}}, \qquad C_2 := \lambda^2 \overline{s}^{\,2} C_{\log}^2. \tag{266}$$

These constants depend only on $\eta, \lambda, \|z_0\|$ and $\|P_\star z_0\|$.

Proposition D.6 (Symmetric rank-one approximation without a strict gap)**.** *Let $0 \neq M \succeq 0$ in dimension d, fix $0 < \eta < 1$, and assume $P_\star x_0 \neq 0$, where $P_\star$ projects onto the leading eigenspace. The rank-one ScaledGD iterates converge to $x_\star = \sqrt{\lambda_1} P_\star x_0 / \|P_\star x_0\|$ and satisfy*

$$0 \leq \|M - x_t x_t^\top\|_F^2 - \|M - x_\star x_\star^\top\|_F^2 \leq \frac{C_1}{t+1} + \frac{C_2}{(t+1)^2}, \qquad t \geq 0. \tag{267}$$

The constants in (266) are evaluated with $A = M$, $\lambda = \lambda_1$ and $z_0 = x_0$; they depend only on $\eta, \lambda_1, \|x_0\|$ and $\|P_\star x_0\|$.

Proof. **Step 1: convergence of the iterates.** Apply Lemmas D.4 and D.5 with $A = M$ and $z_t = x_t$. Every coordinate outside the leading eigenspace decays relative to $\|P_\star x_t\|$, while $P_\star x_t$ remains a positive multiple of $P_\star x_0$. Since $\|x_t\|^2 \to \lambda_1$, $x_t \to x_\star$. This also holds when every eigenvalue equals λ_1, since there are then no coordinates outside the leading eigenspace. The limit satisfies $\|M - x_\star x_\star^\top\|_F^2 = \|M\|_F^2 - \lambda_1^2$.

Step 2: reconstruction error. For the general matrix A in the two lemmas, direct expansion gives

$$\|A - z_t z_t^\top\|_F^2 - (\|A\|_F^2 - \lambda^2) = \lambda^2\big((s_t - 1)^2 + 2 s_t \varepsilon_t\big) \leq \frac{C_1}{t+1} + \frac{C_2}{(t+1)^2}. \tag{268}$$

The inequality uses $s_t \leq \bar{s}$ and the estimates in the two lemmas. Substituting $A = M$ proves the claim and, in particular, (32). □

D.3 Proof of Theorem 7.2(i): gap-dependent rate

Fix $M \neq 0$, $k = 1$, $0 < \eta < 1$ and $\sigma_1 > \sigma_2$, with $\sigma_2 = 0$ if $r = 1$. Assume $a_0 = (u_1^\top y_0 + v_1^\top x_0)/\sqrt{2} \neq 0$. Put $\alpha := 1 - \eta$, $p_1 := 2^{-1/2}(u_1^\top, v_1^\top)^\top$, and $z_t := (y_t^\top, x_t^\top)^\top$. In this subsection z_t is the stacked rectangular iterate, and w_t measures its squared leading coordinate, normalized by σ_1:

$$a_t := p_1^\top z_t, \qquad e_t := \frac{\|z_t - a_t p_1\|^2}{a_t^2}, \qquad w_t := \frac{a_t^2}{\sigma_1}. \tag{269}$$

The assumption $a_0 \neq 0$ makes these quantities well-defined at $t = 0$. We first give the constants used below:

$$\delta := \min\left\{\|z_0\|^2, \frac{4\alpha\eta\sigma_1}{1 + e_0}\right\}, \qquad B := \max\left\{\|z_0\|^2, \sigma_1, \frac{2\alpha\eta\sigma_1 + \eta^2\sigma_1^2/\delta}{1 - \alpha^2}\right\}, \tag{270}$$

$$\underline{w} := \min\left\{1, w_0, \frac{4\alpha\eta}{1 + e_0}\right\}, \qquad \overline{w} := B/\sigma_1. \tag{271}$$

Define

$$\rho := 1 - \frac{\min\{(1-\eta)\underline{w}, \eta/2\}}{(1-\eta)\overline{w} + \eta}\,\frac{\sigma_1 - \sigma_2}{\sigma_1} \in (0, 1), \tag{272}$$

As in §D.1, the intermediate estimates use $2\rho - 1 \in (0, \rho)$, and their combination yields the stated rate ρ^t.

Lemma D.7 (Contraction of dilation coordinate ratios)**.** *With the preceding assumptions and constants, every iterate is well-defined, a_t keeps the sign of a_0, and for all $t \geq 0$,*

$$\delta \leq \|z_t\|^2 \leq B, \qquad \underline{w} \leq w_t \leq \overline{w}, \qquad e_{t+1} \leq (2\rho - 1)^2 e_t \leq (2\rho - 1)^{2(t+1)} e_0. \tag{273}$$

Proof. **Step 1: coordinate ratios and well-posedness.** Let $\{p_j\}_{j=1}^{m+n}$ be an orthonormal eigenbasis of $\mathcal{H}(M)$, with eigenvalues $\ell_1 = \sigma_1$ and $\ell_j \in [-\sigma_1, \sigma_2]$ for $j > 1$. This interval also contains all zero eigenvalues when M is rank deficient. Write $b_{j,t} := p_j^\top z_t$, so $b_{1,t} = a_t$. The update gives

$$b_{j,t+1} = \left(\alpha + \frac{\eta\ell_j}{\|z_t\|^2}\right) b_{j,t}, \qquad a_{t+1} = \left(\alpha + \frac{\eta\sigma_1}{\|z_t\|^2}\right) a_t. \tag{274}$$

The leading multiplier is positive. Inductively $a_t \neq 0$ and $\|z_t\|^2 > 0$ for every t, so all updates are well-defined and the sign of a_t is fixed. For $j > 1$,

$$\left|\frac{b_{j,t+1}}{a_{t+1}}\right| = \frac{|\alpha\|z_t\|^2 + \eta\ell_j|}{\alpha\|z_t\|^2 + \eta\sigma_1}\left|\frac{b_{j,t}}{a_t}\right|, \qquad \frac{|\alpha\|z_t\|^2 + \eta\ell_j|}{\alpha\|z_t\|^2 + \eta\sigma_1} < 1. \tag{275}$$

The strict inequality uses $\alpha\|z_t\|^2 > 0$, $\ell_j \geq -\sigma_1$, and $\ell_j < \sigma_1$. Squaring and summing yields $e_{t+1} \leq e_t \leq e_0$.

Step 2: bounds on $\|z_t\|^2$ and w_t. Since $\|z_t\|^2 = a_t^2(1 + e_t)$ and the sign of a_t is fixed,

$$|a_{t+1}| = \alpha|a_t| + \frac{\eta\sigma_1}{|a_t|(1 + e_t)}, \qquad a_{t+1}^2 \geq \frac{4\alpha\eta\sigma_1}{1 + e_t} \geq \frac{4\alpha\eta\sigma_1}{1 + e_0}. \tag{276}$$

Thus $\|z_t\|^2 \geq \delta$ for every $t \geq 0$. Expanding the squared norm of the update and using $z_t^\top \mathcal{H}(M) z_t \leq \sigma_1\|z_t\|^2$ and $\|\mathcal{H}(M)z_t\|^2 \leq \sigma_1^2\|z_t\|^2$ gives

$$\|z_{t+1}\|^2 \leq \alpha^2\|z_t\|^2 + 2\alpha\eta\sigma_1 + \frac{\eta^2\sigma_1^2}{\|z_t\|^2} \leq \alpha^2\|z_t\|^2 + 2\alpha\eta\sigma_1 + \frac{\eta^2\sigma_1^2}{\delta}. \tag{277}$$

The definition of B closes an induction proving $\|z_t\|^2 \leq B$. Together with (276), this proves

$$\delta \leq \|z_t\|^2 \leq B, \qquad \underline{w} \leq w_t \leq \overline{w}, \qquad t \geq 0. \tag{278}$$

Step 3: uniform contraction of the coordinate ratios. For fixed $s > 0$, the largest value of $|\alpha s + \eta\ell|/(\alpha s + \eta\sigma_1)$ over $\ell \in [-\sigma_1, \sigma_2]$ occurs at an endpoint. At $\ell = \sigma_2$ this ratio increases with s, and is bounded by its value at B. At $\ell = -\sigma_1$ it decreases up to $s = \eta\sigma_1/\alpha$ and then increases, so its maximum on $[\delta, B]$ occurs at δ or B. If $\alpha B < \eta\sigma_1$, its value at B is no larger than its value at δ; otherwise its value at B is bounded by the $\ell = \sigma_2$ ratio there. These endpoint bounds are no larger than $2\rho - 1$. Indeed, for $u \in [\underline{w}, \overline{w}]$,

$$\frac{|\alpha u - \eta|}{\alpha u + \eta} \leq 1 - \frac{2\min\{\alpha\underline{w}, \eta\}}{\alpha\overline{w} + \eta} \leq 2\rho - 1, \tag{279}$$

while the positive endpoint is at most $1 - \eta(\sigma_1 - \sigma_2)/(\sigma_1(\alpha\overline{w} + \eta)) \leq 2\rho - 1$. Here $\underline{w} \leq \delta/\sigma_1 \leq \overline{w}$ and $(\sigma_1 - \sigma_2)/\sigma_1 \leq 1$. Squaring and summing the coordinate ratios gives

$$e_{t+1} \leq (2\rho - 1)^2 e_t, \qquad e_t \leq (2\rho - 1)^{2t} e_0. \tag{280}$$

□

Proposition D.8 (Asymmetric rank-one approximation with a strict gap)**.** *Under the assumptions of Lemma D.7, the iterates converge to a pair $(x_\star, y_\star)$ satisfying $2y_\star x_\star^\top = M_1$, and*

$$\left(\|x_t - x_\star\|^2 + \|y_t - y_\star\|^2\right)^{1/2} \leq K\rho^t, \qquad \|2y_t x_t^\top - M_1\|_F \leq K_M\rho^t \quad (t \geq 0), \tag{281}$$

where

$$K := \sqrt{B}\left(\frac{|\log(w_0)|}{2} + \frac{e_0}{1 - \rho} + \sqrt{e_0}\right). \tag{282}$$

and $K_M = 2(\sqrt{B} + \sqrt{\sigma_1/2})K$.

Proof. **Step 1: recursion for $\log(w_t)$.** By Lemma D.7, $e_t \leq e_0(2\rho - 1)^{2t}$ and $\underline{w} \leq w_t \leq \overline{w}$. The leading-coordinate recursion gives

$$w_{t+1} = \left(\alpha\sqrt{w_t} + \frac{\eta}{\sqrt{w_t}(1 + e_t)}\right)^2. \tag{283}$$

Use ϕ_η from (218). For $w \in [\underline{w}, \overline{w}]$, the derivative with respect to $x = \log(w)$ of $\log(\phi_\eta(e^x))$ is

$$\frac{\alpha e^x - \eta}{\alpha e^x + \eta}. \tag{284}$$

Its absolute value is at most $(2\rho - 1)$ throughout this interval. Since $\phi_\eta(1) = 1$ and the interval contains 1, the mean-value theorem gives $|\log(\phi_\eta(w))| \leq (2\rho - 1)|\log(w)|$. Moreover,

$$\log(w_{t+1}) - \log(\phi_\eta(w_t)) = 2\log\left(\frac{\alpha w_t + \eta/(1 + e_t)}{\alpha w_t + \eta}\right) \in [-2\log(1 + e_t), 0]. \tag{285}$$

It follows that

$$|\log(w_{t+1})| \le (2\rho-1)|\log(w_t)| + 2e_0(2\rho-1)^{2t}. \tag{286}$$

Iterating this inequality and using $(2\rho-1)^2 < \rho$ yields

$$|\log(w_t)| \le (2\rho-1)^t|\log(w_0)| + 2e_0\sum_{j=0}^{t-1}(2\rho-1)^{t-1-j}(2\rho-1)^{2j} \tag{287}$$

$$\le \rho^t\left(|\log(w_0)| + \frac{2e_0}{1-\rho}\right), \qquad t \ge 0. \tag{288}$$

The sum is empty at $t = 0$; for $t \ge 1$ it is bounded by $\sum_{j=0}^{t-1}(2\rho-1)^{t-1-j}\rho^j = (\rho^t - (2\rho-1)^t)/(1-\rho)$.

Step 2: iterate and reconstruction bounds. Set $z_\star := \operatorname{sign}(a_0)\sqrt{\sigma_1}p_1$. Because $w_t \le \overline{w}$ and $\overline{w} \ge 1$, the mean-value theorem for $x \mapsto e^{x/2}$ gives

$$|\sqrt{w_t} - 1| \le \frac{\sqrt{\overline{w}}}{2}|\log(w_t)|. \tag{289}$$

The orthogonal component satisfies $\|z_t - a_t p_1\| = |a_t|\sqrt{e_t} \le \sqrt{B}\sqrt{e_0}(2\rho-1)^t$. Hence

$$\|z_t - z_\star\| \le \sqrt{\sigma_1}|\sqrt{w_t} - 1| + \|z_t - a_t p_1\| \tag{290}$$

$$\le \sqrt{B}\left(\frac{|\log(w_0)|}{2} + \frac{e_0}{1-\rho} + \sqrt{e_0}\right)\rho^t = K\rho^t. \tag{291}$$

This proves the iterate bound and the stated limits in Theorem 7.2(i). Since $2y_\star x_\star^\top = M_1$, $\|y_t\| \le \sqrt{B}$, and $\|x_\star\| = \sqrt{\sigma_1/2}$,

$$\|2y_t x_t^\top - M_1\|_F \le K_M\rho^t, \qquad K_M := 2\left(\sqrt{B} + \sqrt{\sigma_1/2}\right)K, \qquad t \ge 0. \tag{292}$$

Finally, $a_0 = 0$ defines a proper hyperplane in the initialization space, which has probability zero under any law with a density. □

D.4 Proof of Theorem 7.2(ii): gap-free rate

The estimates from §D.2 apply directly to $A = \mathcal{H}(M)$ and the stacked iterate $z_t = (y_t^\top, x_t^\top)^\top$; only the conversion to the rectangular reconstruction error remains.

Proposition D.9 (Asymmetric rank-one approximation without a strict gap)**.** *Let $M \ne 0$, fix $0 < \eta < 1$, and write $z_0 = (y_0^\top, x_0^\top)^\top$. Suppose $P_\star z_0 \ne 0$, where $P_\star$ projects onto the leading positive eigenspace of $\mathcal{H}(M)$. Then the rank-one rectangular iterates converge to a pair $(x_\star, y_\star)$ such that $2y_\star x_\star^\top$ is a best rank-one approximation, and*

$$0 \le \|M - 2y_t x_t^\top\|_F^2 - \|M - 2y_\star x_\star^\top\|_F^2 \le \frac{C_1}{t+1} + \frac{C_2}{(t+1)^2}, \qquad t \ge 0. \tag{293}$$

The constants in (266), evaluated with $A = \mathcal{H}(M)$ and $\lambda = \sigma_1$, depend only on $\eta, \sigma_1, \|z_0\|$ and $\|P_\star z_0\|$.

Proof. **Step 1: convergence of the iterates.** Apply Lemmas D.4 and D.5 with $A = \mathcal{H}(M)$ and $z_t = (y_t^\top, x_t^\top)^\top$. Every coordinate outside the leading positive eigenspace decays relative to $\|P_\star z_t\|$, including those associated with negative eigenvalues. The projected vector keeps its direction and $\|z_t\|^2 \to \sigma_1$, hence

$$z_t \xrightarrow[t\to\infty]{} z_\star := \sqrt{\sigma_1}\frac{P_\star z_0}{\|P_\star z_0\|}. \tag{294}$$

Let $U_\star, V_\star$ contain matching singular vectors for σ_1. Every unit vector in the leading positive dilation eigenspace has the form $(U_\star a, V_\star a)/\sqrt{2}$ with $\|a\| = 1$. The blocks of $z_\star$ therefore satisfy

$$2y_\star x_\star^\top = \sigma_1 U_\star a a^\top V_\star^\top, \qquad \|M - 2y_\star x_\star^\top\|_F^2 = \|M\|_F^2 - \sigma_1^2. \tag{295}$$

Step 2: reconstruction error. Define $\mathcal{E}_t := \|M - 2y_t x_t^\top\|_F^2 - (\|M\|_F^2 - \sigma_1^2)$. Direct expansion gives

$$\mathcal{E}_t = \|\mathcal{H}(M) - z_t z_t^\top\|_F^2 - (\|\mathcal{H}(M)\|_F^2 - \sigma_1^2) - (\|x_t\|^2 - \|y_t\|^2)^2. \tag{296}$$

Dropping the final nonpositive term and applying (268) proves the upper bound. The error is nonnegative by optimality of $2y_\star x_\star^\top$. This proves (34). □

E FURTHER COMPARISONS WITH EXISTING CONVERGENCE GUARANTEES

This appendix supplements §2 and Table 1 with details of the convergence hypotheses and the guarantees being compared.

E.1 GENERAL CONVERGENCE THEORY AND THE SCALEDGD TRAJECTORY

For a proper lower-semicontinuous KL objective, Attouch et al. (2013) prove finite trajectory length and convergence to a critical point for bounded sequences satisfying sufficient decrease, a relative-error bound, and a subsequential continuity condition. Analyticity establishes the KL property locally, but does not supply these trajectory conditions. For our hidden objective, the gradient has no global Lipschitz bound near rank deficiency, so a smoothness-based step-size estimate first requires bounds keeping the iterates away from that boundary. Such an estimate does not directly identify the full spectrum-independent interval $0 < \eta < 1$. Determinant ascent and the matrix Heron recursion establish finite trajectory length throughout this interval without requiring global descent of the hidden objective. For the rectangular proof, we instead verify the decrease and relative-error conditions and invoke KL theory.

Lee et al. (2016) exclude convergence to a fixed strict saddle under an L-Lipschitz gradient and a step size below $1/L$; their countable-saddle corollary uses a countable union of null sets. Rotational symmetry can produce uncountable critical manifolds. We apply Cheridito et al. (2024, Proposition 2.5), which handles uncountable sets of unstable fixed points and requires an invertible Jacobian almost everywhere on the relevant domain, while allowing singular Jacobians at the fixed points. This avoids requiring global invertibility, but convergence to a limit must still be proved separately.

Quantitative results for perturbed GD (Jin et al., 2017) depend on smoothness, Hessian-Lipschitz, and accuracy parameters to reach approximate second-order stationary points. They do not directly establish guarantees for unperturbed ScaledGD throughout $0 < \eta < 1$. Likewise, obtaining a linear rate from a local PL inequality requires identifying its constant and a region of sufficient decrease. Our spectral analysis supplies the relative-eigengap dependence and eventual contraction, but does not quantify the general-rank entrance time.

E.2 INTERPRETING THE CLOSEST APPROXIMATION GUARANTEES

The "linear" entries in Table 1 concern distinct statements. The analysis of Jin et al. (2023) includes local linear convergence after entry into a suitable neighborhood. In our proof, the displayed rate follows from eventual local contraction, with the finite prefix absorbed into random constants. Thus the table compares convergence guarantees, not matched bounds on global iteration complexity; the differing initialization and spectral assumptions remain material.

The scalar Heron argument of Jedra & Shah (2025) cannot simply be applied to each column of our simultaneous update: the shared inverse Gram matrix couples their evolution. The matrix recursion (90) contains an additional positive semidefinite correction determined by the subspace residual. Controlling this term through determinant ascent is what allows the radial argument to extend beyond rank one without sequential deflation.

For the weak-optimality condition studied by Li et al. (2025), take any index set $I \subseteq \{1, \dots, d\}$ with $|I| = k$ and let U_I and Λ_I denote the corresponding eigenvectors and eigenvalues. For any $R \in O(k)$,

$$X = U_I \Lambda_I^{1/2} R \quad \Longrightarrow \quad X^\top M^\dagger X = I_k. \tag{297}$$

Thus this condition also holds at nonleading invariant subspaces. The additional subspace-selection argument in our analysis is needed to identify an optimal approximation.

F NUMERICAL PROTOCOLS AND VERIFICATION

This appendix gives the protocols for Figure 1, including the matrix constructions, random seeds, and numerical verification. All updates are evaluated in the physical coordinates of (4) and (26). Initialization uses independent standard Gaussian entries, without normalization or products with the target matrix. The symmetric targets are positive definite, so Corollary B.10 covers their Gaussian initialization. The rectangular experiments explore behavior beyond the stated initialization guarantees; finite runs do not establish convergence or nonconvergence in general.

F.1 SYMMETRIC CONVERGENCE

We set $d = 50$, $k = 4$, and $(\lambda_k - \lambda_{k+1})/\lambda_k = 0.2$. For each $\kappa \in \{10, 10^4\}$, the leading eigenvalues are geometrically spaced from κ to $\sqrt{\kappa/0.8}$, and the remaining eigenvalues from $\sqrt{0.8\kappa}$ to 1. We draw an orthogonal eigenbasis by Gaussian QR, then draw $X_0 \in \mathbb{R}^{50\times 4}$ with standard Gaussian entries. For each $\eta \in \{0.2, 0.5, 0.8\}$, we run 1000 iterations with seeds $10000 + 97s$, $s = 0, \dots, 5$. Reusing seeds pairs the initializations across step sizes. The plotted error is $\|X_t X_t^\top - M_k\|_F / \|M_k\|_F$, evaluated directly in physical coordinates; lines show medians and shading shows minima and maxima. Final errors in the 36 runs lie between 5.7×10^{-16} and 1.7×10^{-15}. The axis extends to 10^{-16}; errors are not clipped or replaced at the floating-point floor. Both the target construction and iteration use double precision, so the final plateau should not be interpreted as an exact rate.

F.2 RECTANGULAR TRAJECTORIES ACROSS THE STABILITY THRESHOLD

We generate a 9×7 rank-3 matrix with singular values $(6, 2, 1)$ and random orthogonal singular-vector bases using seed 0. Set $k = 2$, and draw $X_0 \in \mathbb{R}^{7\times 2}$ and $Y_0 \in \mathbb{R}^{9\times 2}$ independently with standard Gaussian entries. These iterates are not projected onto the ranges of $M^\top$ and M. The local stability threshold in (30) is $2/(1+6/2) = 1/2$. We run (26) for 2000 iterations with $\eta \in \{0.2, 0.4, 0.51, 0.6\}$, using the same initial pair across step sizes, and plot $\|2Y_t X_t^\top - M_k\|_F / \|M_k\|_F$ for seed 123.

We repeat the test for seeds $123, \dots, 128$. Across the final 500 iterations and all six seeds, the errors for $\eta \in \{0.2, 0.4\}$ are below 10^{-15}. For $\eta = 0.51$ and 0.6, the errors remain above 3×10^{-3}, with large excursions. The iterate increments $(\|X_{t+1} - X_t\|_F^2 + \|Y_{t+1} - Y_t\|_F^2)^{1/2}$ also remain above 0.018 in these tails. Thus the observed failure concerns the iterates as well as reconstruction. This finite-horizon evidence supports the local instability calculation; it does not prove unboundedness or exclude convergence under exceptional initializations.

F.3 MULTIPLE ANGLE REVERSALS FOR RANK THREE IN DIMENSION FIVE

We search 2500 candidates with $M = \Lambda \in \mathbb{R}^{5\times 5}$ and $k = 3$. Using NumPy's default random generator with seed 8741, we draw all spectra first, all step sizes second, and all initial iterates third. Spectra are obtained by sorting five independent 10^U values with U uniform on $[-1, 2]$ in decreasing order; step sizes are 10^V with V uniform on $[-2.5, -0.03]$; each X_0 has independent standard Gaussian entries. Candidate 2193 (zero-based) is selected to illustrate multiple reversals. Its spectrum is approximately $(51.686, 3.976, 0.814, 0.576, 0.173)$ and $\eta \approx 0.020342$; the JSON output records the full numerical data.

We compute an orthonormal basis Q_t for $\mathrm{range}(Z_t)$, where $Z_t = \Lambda^{-1/2} X_t$, by an SVD. The plotted quantity is $\sin\theta_{\max,t} = \|(I_5 - E_3 E_3^\top) Q_t\|_2$, and the second curve is $\Psi(Q_t) = \log(\det(Q_t^\top \Lambda Q_t))$. At $t = 0, 3, 12, 28$, the sine values are approximately 0.9999573, 0.9999985, 0.9999396, and 0.9999999. The two increases are approximately 4.12×10^{-5} and 6.03×10^{-5}, separated by a decrease of 5.89×10^{-5}. The left axis is zoomed near one to make these small changes visible; the example does not demonstrate large oscillations. The potential increases at every displayed step. The plot focuses on $t \le 40$ to show the reversals; at $t = 1500$ the sine is about 2.94×10^{-4}. An independent mpmath replay at 80 decimal digits, using a polar decomposition, checks both quantities at $t = 0, 3, 12, 28, 40$ and agrees with the double-precision sines within 10^{-12}. This is a numerical verification, not an exact-arithmetic certificate. Selection makes this an illustrative example, not an estimate of how frequently such reversals occur.