

The Schwarzian Bridge in a Single Sigmoid Neuron

Sharp Empirical Complexity and Global Spherical Saturation

Lluis Eriksson
Independent Researcher
lluiseriksson@gmail.com

August 2026

Abstract

A sigmoid is increasing on scalars, yet its bend controls two different geometries. The classical Schwarzian criterion implies that a negative-Schwarzian activation is not matrix-monotone of order two. We give a closed logistic formula showing every distinct two-point Löwner determinant to be negative, a dimension-minimal 2×2 positive-semidefinite perturbation witness certified with interval arithmetic, and the corresponding classical serial-chain consequence.

Our central link is that the same local invariant determines the leading quadratic onset coefficient of saturation anisotropy. At a centered inflection of an increasing activation g , for a sensitivity $h_p = (g')^p$ and isotropic spherically symmetric input X ,

$$\kappa_{p,X}(r) = 1 - q_X p Sg(0)r^2 + O(r^4), \quad q_X = \frac{\mathbb{E}\|X\|^4}{d(d+2)}.$$

The Gaussian coefficient is the special case $q_X = 1$. The adjacent-point Löwner determinant is $g'(0)^2 Sg(0)\delta^2/6 + O(\delta^3)$. Negative Schwarzian therefore simultaneously gives local noncommutative order failure and positive initial radial-tangential anisotropy. We turn the equality of limits into a finite-scale inequality with explicit constants determined by derivative suprema and input moments.

For isotropic Gaussian inputs in dimension $d \geq 2$, population curvature at a teacher of norm r has one radial eigenvalue and $d - 1$ equal tangential eigenvalues. For any nonnegative sensitivity profile h with finite defining expectations, these eigenvalues are exactly

$$\alpha_h(r) = \mathbb{E}h(rZ), \quad \beta_h(r) = \mathbb{E}[Z^2 h(rZ)], \quad Z \sim N(0, 1),$$

and, if its zeroth and second Lebesgue moments are finite and positive,

$$\alpha_h(r) \sim \frac{\varphi(0)m_0}{r}, \quad \beta_h(r) \sim \frac{\varphi(0)m_2}{r^3}.$$

Thus saturation suppresses magnitude information by two additional powers: the radial-tangential ratio obeys $\alpha_h(r)/\beta_h(r) \sim (m_0/m_2)r^2$. For every real $p > 0$ and $h_p = \sigma'^p$, we prove strict monotonicity and the unified identities

$$m_0(p) = \frac{\Gamma(p)^2}{\Gamma(2p)}, \quad \kappa_p(r) \sim \frac{r^2}{2\psi_1(p)}, \quad \kappa_p(r) = 1 + \frac{p}{2}r^2 - \frac{p}{8}r^4 + \frac{p(p+1)}{16}r^6 + O(r^8).$$

For n samples we prove an explicit two-sided relative Löwner bound for the weighted empirical Gram matrix. A conditional-Gaussian lower bound shows that its sample complexity is sharply

$$\Theta_p\left(\frac{r[d + \log(1/\delta)]}{\varepsilon^2}\right).$$

The factor rd is intrinsic: a tangential off-diagonal column has $d - 2$ coordinates and normalized variance of order r/n . We also give a uniform full-dyadic-shell extension and concrete sufficient constants $C_1 = 22929$ and $C_2 = 294162$.

The quadratic saturation law is global beyond Gaussian inputs. If $X = RU$ is isotropic spherical, $\mathbb{E}R^{-1} < \infty$, then

$$\kappa_h(r) \sim Q_R \frac{m_0}{m_2} r^2, \quad Q_R = \frac{\mathbb{E}R}{(d-1)\mathbb{E}R^{-1}}.$$

Gaussian, fixed-sphere and isotropic Student inputs give respectively $Q_R = 1$, $d/(d-1)$ and $(\nu-2)/(\nu-1)$. For radii regularly varying at zero, we prove that the inverse-radius condition is sharp through a three-regime phase transition, and give an isotropic Rademacher counterexample. These claims refine, but do not replace, prior one-sided and bilateral Hessian results and recent labeled-estimation theory. Deterministic replays, a fixed-seed finite-sample phase diagram and an immutable public source record accompany the proofs.

Keywords. single neuron; sigmoid; matrix monotonicity; Löwner order; Schwarzian derivative; Fisher information; spherical design; weighted sample covariance; finite-sample concentration.

1 Question, result, and claim boundary

The scalar logistic curve is

$$\sigma(s) = \frac{1}{1 + e^{-s}}, \quad \sigma'(s) = \frac{1}{4} \operatorname{sech}^2(s/2).$$

It is strictly increasing, so larger scalar scores produce larger outputs. We first ask whether this order survives spectral functional calculus on noncommuting matrices. It does not, at the smallest possible matrix size and on every interval. We then return to the ordinary teacher neuron $\sigma(w_\star^\top X)$ with $X \sim N(0, I_d)$. There, a perturbation perpendicular to $w_\star$ rotates the separating hyperplane, while a parallel perturbation changes its confidence scale and must be detected through samples near the transition layer. Our goal is to compute both failures exactly.

The main contributions are:

1. a closed logistic specialization of the classical negative-Schwarzian obstruction, proving directly that every nontrivial two-point Löwner matrix is indefinite;
2. a dimension-minimal 2×2 witness $A \preceq B$ but $\sigma(A) \not\preceq \sigma(B)$, with outward-rounded eigenvalue certification, and a Löwner-order reformulation of the known serial-composition consequence;
3. a spherical-input Schwarzian bridge equating the local order defect with the leading radial-tangential anisotropy, including its exact fourth-moment multiplier q_X ;
4. an explicit finite-scale bridge bounding the discrepancy between the normalized Löwner determinant and anisotropy at nonzero (δ, r) ;
5. a profile-generic two-eigenvalue refinement under Gaussian design, with finite- r brackets and, for every $p > 0$, closed moments, strict ratio monotonicity and exact endpoint laws for $h_p = \sigma'^p$;
6. a two-sided relative concentration theorem for the empirical sensitivity matrix, with a matching $r[d + \log(1/\delta)]/\varepsilon^2$ lower bound for bilateral relative Loewner loss, explicit constants for $p = 1, 2$, and a uniform full-dyadic-shell extension;

7. an exact global spherical saturation theorem with the radius factor Q_R , an elliptical whitening corollary, and a sharp inverse-radius phase transition for radii regularly varying at zero; and
8. exact local consequences for stationary fixed-step optimization and estimation, plus a one-command replay of the identities, interval witness, inequalities, and figure.

The matrix-order theorem concerns spectral functional calculus $\sigma(A)$, not entrywise activation of a vector. The radial theorem is not a claim about deep networks, global optimization, implicit bias, or arbitrary inputs. Its local coefficient and its large-saturation law extend from Gaussian to spherical inputs under different, stated moment hypotheses. The two questions meet at a precise boundary: scalar monotonicity controls neither noncommuting PSD perturbations nor the conditioning of parameter directions.

2 A sigmoid is not an order-preserving operator

For Hermitian matrices, $A \preceq B$ means that $B - A$ is positive semidefinite. A scalar function f is *matrix-monotone of order two* on an interval J if $f(A) \preceq f(B)$ for every pair of 2×2 Hermitian matrices with spectra in J and $A \preceq B$, where $f(A)$ is defined by spectral functional calculus. If A and B commute, scalar monotonicity is sufficient because they are simultaneously diagonalizable. The question is what happens without commutativity.

For distinct $x, y \in J$, define the two-point Löwner matrix

$$L_f(x, y) = \begin{pmatrix} f'(x) & \frac{f(x) - f(y)}{x - y} \\ \frac{f(x) - f(y)}{x - y} & f'(y) \end{pmatrix}. \quad (1)$$

The fixed-order Löwner criterion says that matrix monotonicity of order two requires $L_f(x, y) \succeq 0$ for every pair $[1, 2]$. The equivalent local criterion $Sf \geq 0$ for increasing smooth functions is classical [3, Lemma 10]; the result below is an explicit logistic specialization, not a new qualitative matrix-monotonicity criterion.

Cook, Hammerlindl and Tucker define

$$\chi_f(x, y) = f'(x)f'(y) \left(\frac{x - y}{f(x) - f(y)} \right)^2$$

and study the class $\chi_f \leq 1$, including sigmoid activations and serial compositions [4]. If $D_f(x, y) = (f(x) - f(y))/(x - y)$, then

$$\det L_f(x, y) = D_f(x, y)^2 [\chi_f(x, y) - 1]. \quad (2)$$

Thus their two-point inequality is algebraically the nonpositive Löwner-determinant inequality. Our contribution in this subsection is the strict closed logistic identity and the certified matrix witness, not the underlying Schwarzian/composition mechanism.

Theorem 2.1 (closed logistic two-point determinant). *Let*

$$s_{\alpha, b}(x) = \frac{1}{1 + e^{-\alpha(x-b)}}, \quad \alpha > 0.$$

For every $x \neq y$, the Löwner matrix $L_{s_{\alpha, b}}(x, y)$ has negative determinant. Consequently the logistic activation is not matrix-monotone of order two on any nondegenerate interval.

Proof. Put $u = e^{\alpha(x-b)}$, $v = e^{\alpha(y-b)}$, and $r = v/u = e^{\alpha(y-x)}$. Direct cancellation gives the exact identity

$$\det L_{s_{\alpha,b}}(x, y) = \frac{\alpha^2 u^2}{(1+u)^2(1+v)^2} \left[r - \frac{(r-1)^2}{(\log r)^2} \right]. \quad (3)$$

For $r \neq 1$, write $t = \log r$. Strict convexity of $\sinh$ on $(0, \infty)$ gives

$$|r - 1| = 2e^{t/2} |\sinh(t/2)| > e^{t/2} |t| = \sqrt{r} |\log r|.$$

The bracket in (3) is therefore strictly negative. Every nondegenerate interval contains a distinct pair, so the Löwner criterion completes the proof. $\square$

This is stronger than finding one bad scale: changing the slope α , shifting the bias b , or restricting the score range never removes the obstruction. It also identifies the precise failure. The diagonal entries in (1) are positive, but the divided difference is too large for their geometric mean.

Proposition 2.2 (dimension-minimal explicit PSD witness). *For the standard logistic sigmoid, let*

$$A = \begin{pmatrix} -1/2 & 0 \\ 0 & 1/2 \end{pmatrix}, \quad B = A + \frac{1}{100} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}. \quad (4)$$

Then $A \preceq B$ but $\sigma(A) \not\preceq \sigma(B)$. At 256-bit outward-rounded Arb precision,

$$\text{spec}(B) \subset [-0.4901000, -0.4900999] \cup [0.5100999, 0.5101000], \quad (5)$$

$$\text{spec}(\sigma(B) - \sigma(A)) \subset [-9.9150, -9.9147] 10^{-5} \cup [0.0047990, 0.0047992], \quad (6)$$

$$\det(\sigma(B) - \sigma(A)) \in [-4.7583, -4.7581] 10^{-7}. \quad (7)$$

Proof. $B - A$ has eigenvalues 0 and $1/50$. If $a = 1/2$, $\varepsilon = 1/100$, and $d = (a^2 + \varepsilon^2)^{1/2}$, then B has eigenvalues $\lambda_{\pm} = \varepsilon \pm d$ and

$$\sigma(B) = mI + n(B - \varepsilon I), \quad m = \frac{\sigma(\lambda_+) + \sigma(\lambda_-)}{2}, \quad n = \frac{\sigma(\lambda_+) - \sigma(\lambda_-)}{2d}.$$

Substitution reduces the three enclosures to scalar exponential and square-root balls. The replay evaluates them with directed Arb rounding; the negative eigenvalue in (6) proves the claim. $\square$

The witness is illustrative, while Theorem 2.1 is the proof for every interval. The derivative form also explains why an indefinite Löwner matrix produces witnesses: for diagonal $A = \text{diag}(x, y)$, the Fréchet derivative in the all-ones PSD direction is precisely the Schur product with $L_f(x, y)$.

2.1 Serial depth does not repair the order

For a C^3 function with nonzero derivative, its Schwarzian is

$$Sf = \frac{f'''}{f'} - \frac{3}{2} \left(\frac{f''}{f'} \right)^2. \quad (8)$$

The logistic sigmoid and a scaled hyperbolic tangent satisfy

$$Ss_{\alpha,b} = -\frac{\alpha^2}{2}, \quad S \tanh(\alpha x + \beta) = -2\alpha^2. \quad (9)$$

Corollary 2.3 (serial spectral reformulation). *Let F be a finite scalar composition of positive-slope affine maps and at least one increasing logistic or hyperbolic-tangent activation. Then $SF < 0$ everywhere, and F is not matrix-monotone of order two on any nondegenerate interval.*

Proof. The composition identity

$$S(f \circ g) = (Sf \circ g)(g')^2 + Sg \quad (10)$$

and (9) give $SF < 0$, since positive affine maps have zero Schwarzian and all derivatives in the chain are positive. For $h \rightarrow 0$, Taylor expansion of the two-point determinant gives

$$\det L_F(x, x + h) = \frac{F'(x)^2}{6} SF(x) h^2 + O(h^3). \quad (11)$$

It is negative for all sufficiently small nonzero h . Every interval contains such a pair. $\square$

The corollary restates, in spectral Löwner order, the serial closure proved in the nowhere-coexpanding framework [4]. It is deliberately about *serial scalar chains*. It does not cover sums of neurons, skip connections, multivariate architectures, or entrywise coordinate order. Its content is that composing more of the same S-shaped bend cannot turn the resulting scalar function into a PSD-order-preserving spectral activation.

3 Two models, one sensitivity matrix

Write $r = \|w_\star\|$ and $u = w_\star/r$ when $r > 0$. We use two standard observation models.

Squared-output regression. For the realizable population loss

$$\mathcal{L}_{\text{sq}}(w) = \frac{1}{2} \mathbb{E}[\sigma(w^\top X) - \sigma(w_\star^\top X)]^2, \quad (12)$$

the residual vanishes at $w_\star$, so its Hessian is the Gauss–Newton matrix

$$H_{\text{sq}}(w_\star) = \mathbb{E}[\sigma'(w_\star^\top X)^2 X X^\top]. \quad (13)$$

The same matrix, divided by a known output-noise variance τ^2 , is the Fisher information of $Y = \sigma(w_\star^\top X) + \varepsilon$, $\varepsilon \sim N(0, \tau^2)$.

Bernoulli logistic neuron. If $Y \mid X \sim \text{Bernoulli}(\sigma(w_\star^\top X))$, the expected negative-log-likelihood Hessian and the Fisher information coincide:

$$H_{\text{B}}(w_\star) = \mathbb{E}[\sigma'(w_\star^\top X) X X^\top]. \quad (14)$$

Both are instances of the sensitivity matrix

$$H_h(r, u) := \mathbb{E}[h(r u^\top X) X X^\top], \quad h \geq 0. \quad (15)$$

The power of the reduction below is that no high-dimensional integration remains.

Amari, Karakida and Oizumi already derived the exact one-unit Gaussian Fisher decomposition for $h = \phi'^2$, including the radial–tangential split and the bias block [21]. We do not claim that decomposition or the squared-output branch as new. For the Bernoulli profile $h = \sigma'$, Chen and Mazumdar recently identified the same radial and orthogonal Hessian functions and proved their r^{-3} and r^{-1} orders in a finite-sample analysis of logistic regression [14]. Their follow-up proves the minimax norm-estimation rate $\Theta(\sqrt{r^3/n})$ [15]. We do not claim those exponents or the Bernoulli optimization penalty as new. The purpose of the next sections is to isolate the profile-generic mechanism, give finite-radius brackets and exact logistic constants, prove strict monotonicity of both eigenvalue ratios, and connect their small- r coefficients to the Schwarzian order defect.

4 Universal Gaussian saturation theorem

Let $\varphi(z) = (2\pi)^{-1/2}e^{-z^2/2}$ and define the even sensitivity moments

$$m_j(h) := \int_{\mathbb{R}} t^j h(t) dt \quad (j = 0, 2, 4),$$

whenever finite. Evenness of h is not needed for the spectral reduction or leading asymptotics, but it holds in our sigmoid applications and makes the interpretation symmetric.

Theorem 4.1 (exact spectrum and finite- r brackets). *Let $d \geq 2$, $X \sim N(0, I_d)$, $r > 0$, $u \in S^{d-1}$, and let $h : \mathbb{R} \rightarrow [0, \infty)$ be measurable with $\mathbb{E}[(1 + Z^2)h(rZ)] < \infty$. Then*

$$H_h(r, u) = \alpha_h(r)(I - uu^\top) + \beta_h(r)uu^\top, \quad (16)$$

where

$$\alpha_h(r) = \mathbb{E}h(rZ), \quad \beta_h(r) = \mathbb{E}[Z^2 h(rZ)], \quad Z \sim N(0, 1). \quad (17)$$

If $m_0, m_2 < \infty$ and are positive, then

$$\alpha_h(r) \sim \frac{\varphi(0)m_0}{r}, \quad \beta_h(r) \sim \frac{\varphi(0)m_2}{r^3}, \quad \kappa_h(r) := \frac{\alpha_h(r)}{\beta_h(r)} \sim \frac{m_0}{m_2}r^2. \quad (18)$$

If $m_0, m_2, m_4 < \infty$, the following nonasymptotic brackets hold for every $r > 0$:

$$\frac{\varphi(0)}{r} \left(m_0 - \frac{m_2}{2r^2} \right) \leq \alpha_h(r) \leq \frac{\varphi(0)m_0}{r}, \quad (19)$$

$$\frac{\varphi(0)}{r^3} \left(m_2 - \frac{m_4}{2r^2} \right) \leq \beta_h(r) \leq \frac{\varphi(0)m_2}{r^3}. \quad (20)$$

Negative lower endpoints are interpreted as valid but uninformative bounds.

Proof. Decompose $X = Zu + Y$, where $Z \sim N(0, 1)$, $Y \sim N(0, I - uu^\top)$, and Z, Y are independent. The mixed terms vanish and $\mathbb{E}YY^\top = I - uu^\top$, giving (16)–(17). Changing variables $t = rz$ gives the exact one-dimensional formulae

$$\alpha_h(r) = \frac{\varphi(0)}{r} \int_{\mathbb{R}} h(t) e^{-t^2/(2r^2)} dt, \quad \beta_h(r) = \frac{\varphi(0)}{r^3} \int_{\mathbb{R}} t^2 h(t) e^{-t^2/(2r^2)} dt. \quad (21)$$

Dominated convergence proves (18). Finally, $1 - x \leq e^{-x} \leq 1$ for $x \geq 0$ inserted into (21) proves (19)–(20). $\square$

5 Global saturation beyond Gaussian design

The Gaussian density at the transition hyperplane is not the essential constant. Let $X = RU$ be spherically symmetric, where U is uniform on S^{d-1} and independent of the radius R . Isotropy is exactly $\mathbb{E}R^2 = d$. Write

$$c_d = \frac{\Gamma(d/2)}{\sqrt{\pi} \Gamma((d-1)/2)}, \quad Q_R = \frac{\mathbb{E}R}{(d-1)\mathbb{E}R^{-1}}.$$

Theorem 5.1 (exact global spherical saturation). *Let $d \geq 2$, $R > 0$ almost surely, $\mathbb{E}R^2 = d$ and $\mathbb{E}R^{-1} < \infty$. Let $h \geq 0$ have $0 < m_0, m_2 < \infty$. For $d = 2$, additionally assume $h(t) \leq Ce^{-a|t|}$; no additional condition is needed when $d \geq 3$. Then*

$$\alpha_{h,X}(r) \sim \frac{c_d \mathbb{E}R}{d-1} \frac{m_0}{r}, \quad \beta_{h,X}(r) \sim c_d \mathbb{E}R^{-1} \frac{m_2}{r^3}, \quad (22)$$

$$\kappa_{h,X}(r) \sim Q_R \frac{m_0}{m_2} r^2. \quad (23)$$

For $h_p = \sigma^p$ and every $p > 0$,

$$\kappa_{p,X}(r) \sim \frac{Q_R}{2\psi_1(p)} r^2. \quad (24)$$

Proof. Put $T = u^\top U$. Its density is $c_d(1 - t^2)^{(d-3)/2} \mathbf{1}_{(-1,1)}(t)$, and conditional on $T = t$ every unit transverse coordinate has second moment $(1 - t^2)/(d - 1)$. Changing variables $s = rRT$ gives the exact identities

$$r\alpha_{h,X}(r) = \frac{c_d}{d-1} \mathbb{E} \left[R \int_{|s| < rR} h(s) \left(1 - \frac{s^2}{r^2 R^2} \right)^{(d-1)/2} ds \right], \quad (25)$$

$$r^3 \beta_{h,X}(r) = c_d \mathbb{E} \left[R^{-1} \int_{|s| < rR} s^2 h(s) \left(1 - \frac{s^2}{r^2 R^2} \right)^{(d-3)/2} ds \right]. \quad (26)$$

The first integrand is dominated by Rm_0 . For $d \geq 3$ the second is dominated by $R^{-1}m_2$. For $d = 2$, splitting its inner integral at $|s| = rR/2$ gives a uniform bound from exponential localization; the boundary contribution is bounded by a constant times $(rR)^3 e^{-arR/2}$. Dominated convergence proves (22). The ratio and the identity $m_0(p)/m_2(p) = 1/[2\psi_1(p)]$ prove the rest. $\square$

The factor Q_R is the conditional transverse variance on the transition hyperplane. It is distribution sensitive and satisfies $0 < Q_R \leq d/(d - 1)$. Three closed examples are

isotropic spherical law	radius	Q_R
$N(0, I_d)$	χ_d	1
uniform on $\sqrt{d}S^{d-1}$	$\sqrt{d}$	$d/(d - 1)$
isotropic Student $t_\nu, \nu > 2$	$\sqrt{\nu - 2} \chi_d / \chi_\nu$	$(\nu - 2)/(\nu - 1)$

The last line follows from the adjacent chi-moment identities. If $X = AS$ with A invertible and S isotropic spherical, then exactly

$$H_{h,X}(w) = AH_{h,S}(\|A^\top w\|, A^\top w / \|A^\top w\|) A^\top.$$

Thus the theorem transfers to elliptical inputs after whitening; in Euclidean coordinates it is a generalized eigenvalue statement relative to $AA^\top$, not generally a literal radial–tangential eigendecomposition.

Proposition 5.2 (sharp boundary and failure of isotropy alone). *Suppose the spherical radius has density $f_R(\rho) \sim c\rho^{a-1}$ as $\rho \downarrow 0$, with its tail chosen so that $\mathbb{E}R^2 = d$, and let $h \geq 0$ be exponentially localized with $m_0, m_2 > 0$. Then the tangential eigenvalue remains of order r^{-1} , whereas*

$$\beta_{h,X}(r) \asymp \begin{cases} r^{-(a+2)}, & 0 < a < 1, \\ r^{-3} \log r, & a = 1, \\ r^{-3}, & a > 1. \end{cases}$$

Consequently κ has order r^{a+1} , $r^2/\log r$, or r^2 , respectively. Moreover, covariance isotropy without spherical or transition-density structure is insufficient: for independent Rademacher coordinates, $u = e_1$ and even h , one has $H_h(r, u) = h(r)I_d$ and hence $\kappa = 1$.

Proof. Insert $\rho = q/r$ in the radius integral behind (26). For $a < 1$ the resulting positive constant is

$$c_d c \int_0^\infty q^{a+1} \int_{-1}^1 t^2 h(qt) (1 - t^2)^{(d-3)/2} dt dq,$$

which is finite because its large- q integrand is asymptotic to $m_2 q^{a-2}$. At $a = 1$ this integral diverges logarithmically; at $a > 1$, $\mathbb{E}R^{-1} < \infty$ and Theorem 5.1 applies. The tangential claim follows from (25). For Rademacher input, $X_1^2 = 1$, all mixed expectations vanish by sign symmetry, and $h(rX_1) = h(r)$. $\square$

The radius–direction representation is classical [7]; non-Gaussian reference exponents also appear in regular-design logistic theory [16]. The contribution here is the exact Q_R constant, its all- p logistic specialization, and the sharp inverse-radius boundary within the stated regularly varying class, not the spherical representation or the first occurrence of non-Gaussian r^{-1}/r^{-3} scales.

6 Spherical and finite-scale Schwarzian bridges

The matrix-order and saturation calculations meet exactly at an inflection point. Gaussian independence makes the full saturation curve one-dimensional, but it is not responsible for the local bridge: rotational symmetry and suitable moments suffice. The fourth moment identifies the coefficient below; the stated sixth moment controls its $O(r^4)$ remainder.

Theorem 6.1 (spherical local order–anisotropy bridge). *Let $d \geq 2$ and let $X \in \mathbb{R}^d$ be isotropic and spherically symmetric with $\mathbb{E}\|X\|^6 < \infty$. Let $g \in C^5(\mathbb{R})$ satisfy $g' > 0$ and $g''(0) = 0$. For $p > 0$, suppose $h_p = (g')^p$ has bounded fourth derivative. For $u \in S^{d-1}$ put $Z = u^\top X$ and*

$$H_{p,X}(r, u) = \mathbb{E}[h_p(rZ)XX^\top] = \alpha_{p,X}(r)(I - uu^\top) + \beta_{p,X}(r)uu^\top, \quad \kappa_{p,X} = \frac{\alpha_{p,X}}{\beta_{p,X}}.$$

Then, with

$$q_X := \frac{\mathbb{E}\|X\|^4}{d(d+2)},$$

one has

$$\lim_{r \downarrow 0} \frac{\kappa_{p,X}(r) - 1}{r^2} = -q_X p Sg(0) = -\frac{6q_X p}{g'(0)^2} \lim_{\delta \rightarrow 0} \frac{\det L_g(0, \delta)}{\delta^2}. \quad (27)$$

In particular, $Sg(0) < 0$ makes $L_g(0, \delta)$ indefinite and $\kappa_{p,X}(r) > 1$ at all sufficiently small nonzero scales. Standard Gaussian input has $q_X = 1$.

Proof. Rotations fixing u force the displayed two-eigenspace form, with

$$\beta_{p,X}(r) = \mathbb{E}[Z^2 h_p(rZ)], \quad \alpha_{p,X}(r) = \frac{1}{d-1} \mathbb{E}[(\|X\|^2 - Z^2) h_p(rZ)].$$

Choose a unit $v \perp u$. Spherical symmetry and isotropy give

$$\mathbb{E}Z^4 = 3q_X, \quad \mathbb{E}[Z^2(v^\top X)^2] = q_X.$$

Taylor's formula, cancellation of odd terms under $X \stackrel{d}{=} -X$, and the sixth-moment assumption therefore give

$$\alpha_{p,X}(r) = h_p(0) + \frac{q_X h_p''(0)}{2} r^2 + O(r^4), \quad \beta_{p,X}(r) = h_p(0) + \frac{3q_X h_p''(0)}{2} r^2 + O(r^4).$$

Taking the ratio yields $\kappa_{p,X}(r) = 1 - q_X h_p''(0) r^2 / h_p(0) + O(r^4)$. Since $\log h_p = p \log g'$ and $g''(0) = 0$,

$$\frac{h_p''(0)}{h_p(0)} = p \frac{g'''(0)}{g'(0)} = p Sg(0).$$

The adjacent-point expansion (11) gives the remaining equality. $\square$

The equality of limits has a computable finite-scale version. The constants below are deliberately explicit rather than optimized.

Theorem 6.2 (explicit finite-scale bridge). *Under Gaussian input, retain the hypotheses of Theorem 6.1 and write*

$$a = h_p(0), \quad b = h_p''(0), \quad M_4 = \|h_p^{(4)}\|_\infty.$$

Choose $r_0 > 0$ so that

$$\frac{3|b|}{2}r_0^2 + \frac{5M_4}{8}r_0^4 \leq \frac{a}{2}, \quad (28)$$

and set

$$C_h = \frac{3M_4}{2a} + 3\left(\frac{b}{a}\right)^2 + \frac{5|b|M_4}{4a^2}r_0^2.$$

For a chosen $\delta_0 > 0$, put

$$A = g'(0), \quad c = g'''(0), \quad d_4 = g^{(4)}(0), \quad M_5 = \sup_{|x| \leq \delta_0} |g^{(5)}(x)|,$$

$$K = \frac{|c|}{6} + \frac{|d_4|\delta_0}{24} + \frac{M_5\delta_0^2}{120}, \quad D_g = \frac{|Ad_4|}{12} + \frac{7AM_5\delta_0}{120} + \delta_0 K^2.$$

Then for $0 < r \leq r_0$ and $0 < |\delta| \leq \delta_0$,

$$|\kappa_p(r) - 1 + pSg(0)r^2| \leq C_h r^4, \quad (29)$$

$$\left| \det L_g(0, \delta) - \frac{A^2 Sg(0)}{6} \delta^2 \right| \leq D_g |\delta|^3, \quad (30)$$

$$\left| \frac{\kappa_p(r) - 1}{r^2} + \frac{6p}{A^2} \frac{\det L_g(0, \delta)}{\delta^2} \right| \leq C_h r^2 + \frac{6pD_g}{A^2} |\delta|. \quad (31)$$

Thus the order defect predicts anisotropy at finite, independently chosen scales with a certified error bar.

Proof. Gaussian Taylor remainders give

$$\alpha_p = a + \frac{b}{2}r^2 + R_\alpha, \quad |R_\alpha| \leq \frac{M_4}{8}r^4, \quad \beta_p = a + \frac{3b}{2}r^2 + R_\beta, \quad |R_\beta| \leq \frac{5M_4}{8}r^4.$$

Condition (28) makes $\beta_p \geq a/2$. Subtracting $1 - (b/a)r^2$ before dividing by β_p yields

$$\left| \kappa_p - 1 + \frac{b}{a}r^2 \right| \leq \left[\frac{3M_4}{2a} + 3\left(\frac{b}{a}\right)^2 + \frac{5|b|M_4}{4a^2}r_0^2 \right] r^4,$$

which is (29) because $b/a = pSg(0)$.

Taylor expansion at zero gives

$$g'(\delta) = A + \frac{c}{2}\delta^2 + \frac{d_4}{6}\delta^3 + R_1, \quad \frac{g(\delta) - g(0)}{\delta} = A + \frac{c}{6}\delta^2 + \frac{d_4}{24}\delta^3 + R_2,$$

where $|R_1| \leq M_5|\delta|^4/24$ and $|R_2| \leq M_5|\delta|^4/120$. Writing the determinant as $A(g'(\delta) - A) - 2A(V - A) - (V - A)^2$, with $V = (g(\delta) - g(0))/\delta$, gives (30) with the displayed D_g . Finally, divide (29) by r^2 , divide (30) by δ^2 , multiply the latter by $6p/A^2$, and use the triangle inequality. $\square$

For the standard logistic sigmoid, $S\sigma = -1/2$. Theorem 6.1 therefore predicts the exact initial coefficient $p/2$ for every $h_p = \sigma'^p$, recovered globally below. Equation (31) strengthens the conceptual coefficient match into a falsifiable finite-scale comparison.

For a general h , $\kappa_h = \alpha_h/\beta_h$ in (18) is an anisotropy ratio; it is the spectral condition number only after the eigenvalue ordering is known. The two powers are geometric. At large r , only a score slab of width $O(r^{-1})$ remains unsaturated, producing $\alpha \asymp r^{-1}$. A radial perturbation carries an additional factor $Z^2 = O(r^{-2})$ inside that slab, producing $\beta \asymp r^{-3}$. Both modes are supported by the near-boundary slab; the radial mode is weaker because it carries the additional Z^2 factor.

7 All-power logistic saturation laws

For every real $p > 0$ set $h_p(t) = \sigma'(t)^p$. The cases $p = 1$ and $p = 2$ are Bernoulli Fisher and squared-output curvature, respectively, but the analytic law is not restricted to those observation models. The Gamma transform below is classical generalized-logistic distribution theory [23, eqs. (2.3)–(2.5)]; our use of it is geometric.

Lemma 7.1 (closed sensitivity moments for every $p > 0$). *Let ψ_j denote the polygamma function of order j . Then*

$$\widehat{h}_p(k) := \int_{\mathbb{R}} e^{ikt} h_p(t) dt = \frac{\Gamma(p+ik)\Gamma(p-ik)}{\Gamma(2p)}, \quad (32)$$

and

$$m_0(p) = \frac{\Gamma(p)^2}{\Gamma(2p)}, \quad m_2(p) = 2\psi_1(p)m_0(p), \quad m_4(p) = (12\psi_1(p)^2 + 2\psi_3(p))m_0(p). \quad (33)$$

In particular,

profile	m_0	m_2	m_4
$h_1 = \sigma'$	1	$\pi^2/3$	$7\pi^4/15$
$h_2 = \sigma'^2$	1/6	$(\pi^2 - 6)/18$	$7\pi^4/90 - 2\pi^2/3$

Proof. With $x = e^t$,

$$h_p(t) = \frac{x^p}{(1+x)^{2p}}, \quad \widehat{h}_p(k) = \int_0^\infty \frac{x^{p+ik-1}}{(1+x)^{2p}} dx = B(p+ik, p-ik),$$

which proves (32). Since $h_p(t) = O(e^{-p|t|})$, all polynomial moments are finite and differentiation under the Fourier integral is justified. Differentiating its logarithm at zero gives $(\log \widehat{h}_p)''(0) = -2\psi_1(p)$ and $(\log \widehat{h}_p)^{(4)}(0) = 2\psi_3(p)$. Since $m_2 = -\widehat{h}_p''(0)$ and $m_4 = \widehat{h}_p^{(4)}(0)$, (33) follows. The table uses the standard integer polygamma values. $\square$

Theorem 7.2 (strict all-power logistic anisotropy). *For every $p > 0$, $\kappa_p(0) = 1$, $\kappa_p(r) > 1$ for $r > 0$, and κ_p is strictly increasing on $(0, \infty)$. At the two endpoints,*

$$\kappa_p(r) = 1 + \frac{p}{2}r^2 - \frac{p}{8}r^4 + \frac{p(p+1)}{16}r^6 + O(r^8) \quad (r \downarrow 0), \quad (34)$$

and

$$\kappa_p(r) = \frac{r^2}{2\psi_1(p)} + 1 + \frac{\psi_3(p)}{4\psi_1(p)^2} + O(r^{-2}) \quad (r \rightarrow \infty). \quad (35)$$

Thus $\kappa_1(r) \sim 3r^2/\pi^2$ and $\kappa_2(r) \sim 3r^2/(\pi^2 - 6)$.

Proof. Normalize the density $q_{p,r}(z) \propto h_p(rz)\varphi(z)$. Then $\beta_p(r)/\alpha_p(r) = \mathbb{E}_{q_{p,r}} Z^2$. If $r_2 > r_1$, the derivative on $z > 0$ of the log likelihood ratio is

$$\frac{d}{dz} \log \frac{h_p(r_2 z)}{h_p(r_1 z)} = -p[r_2 \tanh(r_2 z/2) - r_1 \tanh(r_1 z/2)] < 0.$$

Thus the distribution of $|Z|$ under $q_{p,r}$ decreases strictly in monotone-likelihood-ratio order as r increases, and its strictly increasing statistic Z^2 has decreasing expectation. Hence α_p/β_p increases strictly. At $r = 0$ isotropy gives one. Lemma 7.1 inserted into Theorem 4.1, with one further term in $e^{-t^2/(2r^2)}$, gives the large- r statement. More explicitly, $m_6(p) < \infty$ by the exponential tail, and Taylor remainder bounds for e^{-x} give

$$\alpha_p(r) = \frac{\varphi(0)}{r} \left(m_0 - \frac{m_2}{2r^2} + \frac{m_4}{8r^4} + O(r^{-6}) \right), \quad \beta_p(r) = \frac{\varphi(0)}{r^3} \left(m_2 - \frac{m_4}{2r^2} + O(r^{-4}) \right).$$

Their ratio yields (35). Finally,

$$h_p(t) = 4^{-p} \left[1 - \frac{p}{4} t^2 + \frac{p(1+3p)}{96} t^4 - \frac{p(15p^2+15p+4)}{5760} t^6 + O(t^8) \right].$$

For each fixed $p > 0$, every derivative of h_p is bounded. Taylor's theorem therefore bounds the displayed remainder globally by $\|h_p^{(8)}\|_\infty |t|^8/8!$; after setting $t = rZ$, this is integrable in both α_p and β_p because the Gaussian has a finite tenth moment. Inserting the Gaussian moments through $\mathbb{E}Z^8 = 105$ in the coefficients and taking the ratio gives (34). $\square$

Figure 1 shows the two statistically canonical members of the all-power family. The solid and dashed curves are direct Gaussian integrals; the dotted curves are the asymptotic constants in (35).




Figure 1: Exact curvature/information eigenvalues and their ratio. Saturation leaves $d - 1$ tangential directions at order r^{-1} but pushes the radial direction to order r^{-3} . The quadratic condition-number law is not a fitted exponent; it follows from Theorem 4.1.

The cancellation of every p^2 contribution in the fourth-order ratio coefficient is worth noting: the full all-power family has the correction $-pr^4/8$, followed by the first nonlinear-in- p term $p(p+1)r^6/16$.

8 Finite-sample resolution of the weighted Gram matrix

The population eigenvalues alone do not say when a sample resolves the radial mode. We now answer that question for the fixed-oracle matrix

$$\widehat{H}_{h,n}(r, u) = \frac{1}{n} \sum_{i=1}^n h(ru^\top X_i) X_i X_i^\top, \quad X_i \stackrel{\text{iid}}{\sim} N(0, I_d).$$

Prior self-concordant analyses already give bilateral constant-factor empirical-Hessian sandwiches, including a uniform result for Gaussian logistic regression [17, Eq. (92)]; fixed-point standardized-Hessian concentration is also available under generic matrix-Bernstein conditions [18]. Chardon, Lerasle and Mourtada prove a uniform one-sided lower bound for the empirical Bernoulli Hessian ($p = 1$) under the sufficient condition $n \gtrsim r(d + t)$ [16, Theorem 6]. We do not claim the first bilateral concentration or the first finite-sample lower bound. The result below instead gives tunable two-sided relative error for a fixed teacher, applies to a general bounded profile and hence every $p > 0$, separates radial, tangential and cross errors, and returns an eigenspace angle.

Put $m = d - 1$, $Z = u^\top X$, $Y = X - Zu$, and $W = h(rZ)$. Define

$$\alpha = \mathbb{E}W, \quad \beta = \mathbb{E}[WZ^2], \quad \gamma_j = \mathbb{E}[W^2 Z^{2j}] \quad (j = 0, 1, 2),$$

and the envelopes

$$H = \|h\|_\infty, \quad K_1 = \sup_z z^2 h(rz)^2, \quad K_2 = \sup_z z^2 h(rz).$$

Theorem 8.1 (explicit relative empirical spectrum). *Assume $0 \leq h \leq H$, $K_1, K_2 < \infty$, and $\alpha, \beta > 0$. Let $0 < \delta < 1$ and set*

$$t = \log(12/\delta), \quad u_0 = m \log 9 + t,$$

$$S_0^+ = n\gamma_0 + \sqrt{2nH^2\gamma_0 t} + \frac{H^2 t}{3}, \quad S_1^+ = n\gamma_1 + \sqrt{2nK_1\gamma_1 t} + \frac{K_1 t}{3},$$

$$e_T = \sqrt{\frac{2\gamma_0 t}{n}} + \frac{Ht}{3n} + \frac{4}{n}(\sqrt{S_0^+ u_0} + H u_0), \tag{36}$$

$$e_R = \sqrt{\frac{2\gamma_2 t}{n}} + \frac{K_2 t}{3n}, \quad q_n = \frac{\sqrt{S_1^+}}{n}(\sqrt{m} + \sqrt{2t}). \tag{37}$$

Then, with probability at least $1 - \delta$, in the radial-tangential basis

$$\widehat{H}_{h,n} = \begin{pmatrix} C & b_n \\ b_n^\top & c_n \end{pmatrix}$$

satisfies simultaneously

$$\|C - \alpha I_m\|_{\text{op}} \leq e_T, \quad |c_n - \beta| \leq e_R, \quad \|b_n\|_2 \leq q_n. \tag{38}$$

Consequently, for

$$\varepsilon_n := \max \left\{ \frac{e_T}{\alpha}, \frac{e_R}{\beta} \right\} + \frac{q_n}{\sqrt{\alpha\beta}}, \tag{39}$$

one has the two-sided relative Löwner bound

$$(1 - \varepsilon_n)H_h(r, u) \preceq \widehat{H}_{h,n}(r, u) \preceq (1 + \varepsilon_n)H_h(r, u). \tag{40}$$

If $G := \alpha - \beta - e_T - e_R > 0$, the empirical matrix has a unique eigenvalue $\hat{\beta}$ below the tangential block spectrum, and its unit eigenvector $\hat{u}$ obeys

$$\tan \angle(\hat{u}, u) \leq \frac{q_n}{G}, \quad |\hat{\beta} - \beta| \leq e_R + \frac{q_n^2}{G}. \quad (41)$$

Every other ordered eigenvalue differs from α by at most $\max\{e_T, e_R\} + q_n$.

Proof. Write

$$C = \frac{1}{n} \sum_i W_i Y_i Y_i^\top, \quad b_n = \frac{1}{n} \sum_i W_i Z_i Y_i, \quad c_n = \frac{1}{n} \sum_i W_i Z_i^2.$$

Scalar Bernstein controls $\bar{W} - \alpha$ and, one-sidedly, $\sum_i W_i^2 \leq S_0^+$. Conditional on the weights, for a fixed unit v ,

$$v^\top \sum_i W_i (Y_i Y_i^\top - I_m) v = \sum_i W_i (G_i^2 - 1), \quad G_i \stackrel{\text{iid}}{\sim} N(0, 1).$$

The exact Gaussian moment-generating function gives

$$\left| \sum_i W_i (G_i^2 - 1) \right| \leq 2 \sqrt{u_0 \sum_i W_i^2} + 2H u_0$$

outside conditional probability $2e^{-u_0}$. A $1/4$ -net of S^{m-1} has at most 9^m points and $\|A\|_{\text{op}} \leq 2 \max_{v \text{ in net}} |v^\top A v|$, proving the first bound in (38). Bernstein applied to $W_i Z_i^2 \in [0, K_2]$ proves the radial bound.

Conditionally on the Z_i ,

$$b_n \sim N\left(0, \frac{\sum_i W_i^2 Z_i^2}{n^2} I_m\right).$$

A second one-sided Bernstein bound gives $\sum_i W_i^2 Z_i^2 \leq S_1^+$, and the Gaussian norm tail proves the cross bound. The listed failures total less than $9e^{-t} < \delta$.

Conjugating $\hat{H}_{h,n} - H_h$ by $H_h^{-1/2} = \text{diag}(\alpha^{-1/2} I_m, \beta^{-1/2})$ shows that its norm is at most (39); this is equivalent to (40). If $G > 0$, then $c_n < \lambda_{\min}(C)$. Interlacing leaves one eigenvalue below C . The eigenvector equation and the Schur complement give respectively $\tan \angle(\hat{u}, u) \leq q_n/G$ and $0 \leq c_n - \hat{\beta} \leq q_n^2/G$. Weyl's inequality gives the remaining assertion. $\square$

For $h_p = \sigma'^p$, all quantities in Theorem 8.1 are one-dimensional and computable. The elementary bound $\sigma'(s) \leq e^{-|s|}$ gives, for every $p > 0$,

$$H = 4^{-p}, \quad K_1 \leq \frac{e^{-2}}{p^2 r^2}, \quad K_2 \leq \frac{4e^{-2}}{p^2 r^2}, \quad (42)$$

and, with $\varphi(0) = (2\pi)^{-1/2}$,

$$\gamma_0 \leq \frac{\varphi(0)}{pr}, \quad \gamma_1 \leq \frac{\varphi(0)}{2p^3 r^3}, \quad \gamma_2 \leq \frac{3\varphi(0)}{2p^5 r^5}. \quad (43)$$

Indeed, after $s = rZ$, use $h_p(s)^2 \leq e^{-2p|s|}$ and $\int_{\mathbb{R}} |s|^j e^{-2p|s|} ds = 2j!/(2p)^{j+1}$ for $j = 0, 2, 4$.

Corollary 8.2 (effective sample size for every logistic power). *Let*

$$R_p^2 = \max \left\{ \frac{3m_2(p)}{m_0(p)}, \frac{m_4(p)}{m_2(p)} \right\}.$$

For $r \geq R_p$, $0 < \varepsilon \leq 1/2$ and $0 < \delta < 1$, the explicit condition $\varepsilon_n \leq \varepsilon$ in (39), with (42)–(43) substituted, guarantees (40). In particular, for each fixed $p > 0$ there are finite computable constants C_p, C'_p such that

$$n \geq C_p \frac{r [d + \log(12/\delta)]}{\varepsilon^2} \quad (44)$$

is sufficient and, after increasing C_p if necessary, $e_T + e_R \leq (\alpha - \beta)/2$. Hence $G \geq (\alpha - \beta)/2$ and, on the same event,

$$\sin \angle(\hat{u}, u) \leq C'_p \left(\sqrt{\frac{d + \log(1/\delta)}{nr}} + \frac{\sqrt{[d + \log(1/\delta)] \log(1/\delta)}}{n} \right).$$

Thus the effective number of transition-layer observations is n/r : relative resolution of a population eigenvalue of order r^{-3} does not itself require r^3 samples.

Proof. The brackets in Theorem 4.1 and the definition of R_p give

$$\alpha \geq \frac{\varphi(0)m_0(p)}{2r}, \quad \beta \geq \frac{\varphi(0)m_2(p)}{2r^3}, \quad \alpha - \beta \geq \frac{\varphi(0)m_0(p)}{2r}.$$

Insert these and (42)–(43) in (36)–(39). Each relative block term is bounded by a p -dependent constant times $\sqrt{r[d + \log(12/\delta)]/n} + r[d + \log(12/\delta)]/n$. The cross term has the same first order. Increasing a finite constant C_p makes their sum at most ε and also makes $e_T + e_R \leq (\alpha - \beta)/2$. Thus $G \geq (\alpha - \beta)/2 \geq \varphi(0)m_0(p)/(4r)$, and substitution of q_n into (41) gives the displayed bound with a finite computable C'_p . $\square$

For the two statistical powers, the conservative majorant used in this proof gives fully numerical choices:

p	R_p	C_p in (44)	C'_p in the angle bound
1	3.717183	22929	22.479249
2	2.153017	294162	52.368963

These are sufficient, not optimized, constants. The explicit-constants replay prints every intermediate coefficient and checks that the resulting quadratic majorant is at most ε for $0 < \varepsilon \leq 1/2$.

We next prove that the product scale in (44) is intrinsic for the stated bilateral loss. Choose any finite $a_p > 0$ satisfying

$$\int_{|s| > a_p} s^2 h_p(s) ds \leq \frac{m_2(p)}{16}.$$

For $r \geq \max\{\sqrt{m_4(p)/m_2(p)}, 4\varphi(0)a_p\}$, the empty-transition-slab event and Markov's inequality give

$$\Pr\left\{\lambda_{\min}(\hat{H}_{p,n}) < \frac{\beta_p(r)}{2}\right\} \geq \frac{1}{2} \exp\left(-\frac{4\varphi(0)a_p n}{r}\right). \quad (45)$$

Hence failure probability at most $\delta < 1/2$ requires

$$n \geq \frac{r}{4\varphi(0)a_p} \log \frac{1}{2\delta}. \quad (46)$$

To verify (45), condition on no $|Z_i| \leq a_p/r$. This event has probability at least the exponential shown; the conditional mean radial entry is at most $\beta_p/4$, so with conditional probability at least $1/2$ that entry, and therefore the smallest eigenvalue, is below $\beta_p/2$.

Theorem 8.3 (matching lower complexity for bilateral relative Loewner loss). *Fix $p > 0$. There are finite constants $\tilde{R}_p, c_p, \delta_p > 0$ such that, for $r \geq \tilde{R}_p$, $d \geq 3$, $0 < \varepsilon \leq 1/2$ and $0 < \delta \leq \delta_p$, the implication*

$$\Pr\{(1 - \varepsilon)H_p \preceq \hat{H}_{p,n} \preceq (1 + \varepsilon)H_p\} \geq 1 - \delta \implies n \geq c_p \frac{r[d + \log(1/\delta)]}{\varepsilon^2} \quad (47)$$

holds. At constant confidence the dimension term has the explicit sufficient obstruction

$$n < \frac{m_0(2p)}{16\varphi(0)m_0(p)^2} \frac{r(d-2)}{\varepsilon^2} \implies \Pr\{\text{relative Loewner failure}\} \geq \frac{1}{1152}, \quad (48)$$

provided $n\gamma_0 \geq H^2$ and the finite-radius moment brackets hold. The displayed coefficient equals $1/[96\varphi(0)]$ for $p = 1$ and $9/[560\varphi(0)]$ for $p = 2$.

Proof. Restrict the relative inequality to the $m = d - 1$ dimensional tangential block $C = n^{-1} \sum_i W_i Y_i Y_i^\top$. Put $\gamma_0 = \mathbb{E}W^2$ and $S_0 = \sum_i W_i^2$. Since $W \leq H$, Paley–Zygmund gives

$$\Pr\{S_0 \geq n\gamma_0/2\} \geq 1/8 \quad (n\gamma_0 \geq H^2).$$

Conditional on the weights, $T_0 = \sum_i W_i^2 Y_{i1}^2$ has second moment at most $3S_0^2$, so $\Pr\{T_0 \geq S_0/2 \mid W\} \geq 1/12$. Conditional on (W, Y_1) , the remaining $d - 2$ entries of the first off-diagonal column of C are Gaussian with common variance T_0/n^2 . A final Paley–Zygmund bound gives $\Pr\{\chi_{d-2}^2 \geq (d-2)/2\} \geq 1/12$. Hence, with probability at least $1/1152$,

$$\|C/\alpha - I_m\|_{\text{op}} \geq \sqrt{\frac{(d-2)\gamma_0}{8n\alpha^2}}.$$

The scaled integrals give

$$\frac{\gamma_0}{\alpha^2} \geq \frac{m_0(2p)}{2\varphi(0)m_0(p)^2} r,$$

which proves (48). If $n\gamma_0 < H^2$, then $n = O_p(r)$ and the slab bound (45) supplies a fixed failure probability after decreasing δ_p .

For the confidence term, normalize $A = W/\alpha$ and put $T = n^{-1} \sum_i A_i^2 Y_{i1}^2$. Its mean is $\rho = \mathbb{E}W^2/\alpha^2 \asymp_p r$, and

$$\frac{\mathbb{E}T^2}{\rho^2} \leq 1 + \frac{3\mathbb{E}W^4}{n(\mathbb{E}W^2)^2} \leq 1 + C_p \frac{r}{n}$$

by the same finite-radius moment brackets. Thus, when $n \geq a_p r$, $\Pr\{T \geq \rho/2\} \geq q_p > 0$ by Paley–Zygmund; when $n < a_p r$, the slab event supplies the required fixed failure probability after fixing δ_p . Conditional on T , one off-diagonal entry is $N(0, T/n)$. The elementary lower Gaussian tail $2\Phi(-x) \geq c_0 e^{-x^2}$ yields a failure probability at least $c'_p \exp(-2n\varepsilon^2/\rho)$. Therefore failure at most δ requires $n \gtrsim_p r\varepsilon^{-2} \log(1/\delta)$. Combining this with the dimension bound by $\max\{a, b\} \geq (a+b)/2$ proves (47). $\square$

The theorem is about the fixed-oracle weighted Gram matrix, not a minimax estimator. Nor should it be exported unchanged to weaker targets. Direct scalar Bernstein gives an $O_p(r/\varepsilon^2)$ radial-entry upper bound at constant confidence; the empty-slab argument gives only the fixed-accuracy $\Omega_p(r)$ obstruction recorded above. We do not prove a matching lower threshold for the ordinary bottom eigenspace; the sufficient angle event above remains $O_p(rd)$ at constant accuracy.

8.1 Uniform all-power sensitivity on a full dyadic shell

The pointwise theorem admits a proof-complete full-dyadic-shell extension at the price of an elementary covering logarithm. For $G_p(w) = \mathbb{E}[h_p(w^\top X)XX^\top]$ and its empirical analogue $\hat{G}_{p,n}(w)$, set

$$B_{n,\delta} = \sqrt{d} + \sqrt{2\log(2n/\delta)}, \quad \mu_{3,d} = [d(d+2)]^{3/4}, \quad L_p = p4^{-p},$$

$$b_p = \frac{\varphi(0)m_2(p)}{16}, \quad A_{p,n,\delta} = L_p(B_{n,\delta}^3 + \mu_{3,d}), \quad \eta = \min \left\{ R, \frac{\varepsilon b_p}{32A_{p,n,\delta}R^3} \right\}, \quad N = \left(1 + \frac{4R}{\eta} \right)^d.$$

Theorem 8.4 (uniform relative sensitivity on a full dyadic shell). *Fix $p > 0$, $R \geq R_p$, $0 < \varepsilon \leq 1/2$, and $0 < \delta < 1/2$. If*

$$n \geq 128C_p \frac{R}{\varepsilon^2} \left[d + \log \left(\frac{24N}{\delta} \right) \right], \quad (49)$$

then, with probability at least $1 - \delta$, simultaneously for every w with $R \leq \|w\| \leq 2R$,

$$(1 - \varepsilon)G_p(w) \preceq \hat{G}_{p,n}(w) \preceq (1 + \varepsilon)G_p(w). \quad (50)$$

The computable condition is implicit in n only through $B_{n,\delta}$; no optimality is claimed for its covering logarithm.

Proof. Since $\|h'_p\|_\infty \leq L_p$,

$$\|h_p(w^\top x)xx^\top - h_p(v^\top x)xx^\top\|_{\text{op}} \leq L_p\|w - v\|\|x\|^3.$$

With probability at least $1 - \delta/2$, all sampled Gaussian norms are at most $B_{n,\delta}$. On this event $D_n = \hat{G}_{p,n} - G_p$ is $A_{p,n,\delta}$ -Lipschitz in operator norm. Cover the shell by an η -net of size at most N and apply Corollary 8.2 at every net point with error $\varepsilon/8$ and failure $\delta/(2N)$; (49) is exactly a conservative sufficient condition. On the shell, $\lambda_{\min}G_p(w) \geq b_pR^{-3}$. Moving from a net point to w therefore changes both D_n and G_p by at most $(\varepsilon/32)\lambda_{\min}G_p(w)$. For every x ,

$$|x^\top D_n(w)x| \leq \left[\frac{\varepsilon}{8}(1 + \varepsilon/32) + \frac{\varepsilon}{32} \right] x^\top G_p(w)x \leq \varepsilon x^\top G_p(w)x,$$

which is (50). The two failure probabilities total at most δ . $\square$

For $p = 1$, $\hat{G}_{1,n}(w)$ is exactly the logistic-likelihood Hessian and is label-free. For $p \neq 1$ it is a sensitivity Gram, not that likelihood Hessian. Earlier work proves uniform local constant-factor sandwiches and an optimal-scale one-sided $p = 1$ lower bound [17, 16]; the narrower addition here is a full dyadic shell, all $p > 0$, and tunable bilateral precision. It is not by itself an estimator theorem.




Figure 2: Fixed-seed diagnostic at $d = 12, r = 6$, with 320 repetitions per sample size. Solid curves are medians and dashed curves are 90th percentiles. The first three panels display tangential relative error, ordinary smallest-eigenvalue relative error, and radial eigenspace angle. The fourth compares the exact empty-slab probability with its simulated frequency; exact values below 10^{-8} are clipped only in the plot and remain untruncated in the JSON certificate. The normalization $n/(rd)$ is the sharp scale for full bilateral relative Loewner loss; the ordinary eigenvalue and angle panels are diagnostics and make no claim that this is their optimal threshold.

9 What the law costs

9.1 Stationary fixed-step gradient descent

For either logistic profile, assume $d \geq 2$ and that one constant scalar step size η is reused at every iteration. At the teacher, the Jacobian of one gradient-descent step is $I - \eta H_h$. Its radial eigenvalue is $1 - \eta\beta$ and its tangential eigenvalue is $1 - \eta\alpha$. The best scalar step for this local quadratic model is

$$\eta_\star = \frac{2}{\alpha + \beta}, \quad \rho_\star = \frac{\alpha - \beta}{\alpha + \beta} = \frac{\kappa - 1}{\kappa + 1}. \quad (51)$$

Consequently $-\log \rho_\star \sim 2/\kappa$. For worst-case local error with components allowed in both eigenspaces, the e-folding iteration count $T_e := (-\log \rho_\star)^{-1}$ therefore has the exact leading scaling

$$T_{e,\text{sq}} \sim \frac{3}{2(\pi^2 - 6)} r^2, \quad T_{e,\text{B}} \sim \frac{3}{2\pi^2} r^2 \quad (52)$$

for the optimally tuned stationary scalar-step linearization. This is not a global iteration bound or an algorithm-independent obstruction: away from the teacher the Hessian contains residual terms, and nonstationary polynomial methods can exploit the two-point spectrum.

Indeed, the scalar schedule $\eta_1 = 1/\alpha$, $\eta_2 = 1/\beta$ makes $(I - \eta_2 H)(I - \eta_1 H) = 0$ in this exact local quadratic. Natural-gradient or matrix-preconditioned updates can also remove the local condition number by acting differently on the two eigenspaces.

9.2 Estimation

For n independent Gaussian-output observations with known variance τ^2 , the Fisher information is nH_{h_2}/τ^2 . The Cramér–Rao inequality gives, along any unit tangential vector $v \perp u$ and the radial vector u ,

$$\text{Var}(v^\top \hat{w}) \geq \frac{\tau^2}{n\alpha_2(r)}, \quad \text{Var}(u^\top \hat{w}) \geq \frac{\tau^2}{n\beta_2(r)}. \quad (53)$$

The ratio of the displayed radial lower bound to the tangential lower bound is $\kappa_2(r)$. Hence tangential variance is bounded below at order r/n , while radial variance is bounded below at order r^3/n . The analogous Bernoulli bounds replace h_2 by h_1 . The same quadratic ratio therefore appears in these coordinatewise information lower bounds and in stationary fixed-step descent, although the latter can be removed by a nonstationary schedule.

10 Bias and universality

A conventional neuron includes a bias b . At $b = 0$ the augmented parameter (w, b) has Fisher/Hessian block

$$\mathbb{E} \left[h(rZ) \begin{pmatrix} XX^\top & X \\ X^\top & 1 \end{pmatrix} \right].$$

For even h , the radial–bias cross term $\mathbb{E}[Zh(rZ)]$ vanishes. The bias eigenvalue equals $\alpha_h(r)$, so the spectrum contains one radial eigenvalue β_h and d tangential/bias eigenvalues α_h . A nonzero bias preserves the $(d - 1)$ tangential eigenspace but replaces the radial and bias coordinates by the explicit 2×2 block

$$\begin{pmatrix} \mathbb{E}[Z^2 h(rZ + b)] & \mathbb{E}[Zh(rZ + b)] \\ \mathbb{E}[Zh(rZ + b)] & \mathbb{E}[h(rZ + b)] \end{pmatrix}. \quad (54)$$

The centered theorem is therefore not an artifact of omitting the bias; it is the diagonal point of a fully explicit two-coordinate extension.

The exponent two is broader than the logistic curve. Theorem 4.1 shows that any localized sensitivity profile with finite positive m_0 and m_2 has the same r^2 anisotropy, with only the constant m_0/m_2 depending on the activation and observation model. Profiles with heavy sensitivity tails or vanishing second moment fall outside this universality class and can have different laws.

11 Relation to prior work

Table 1 states the claim boundary before the detailed discussion.

Prior result	Result here	Exact difference
Fixed-order Löwner/Schwarzian criterion and serial closure [3, 4]	closed logistic determinant and certified 2×2 witness	explicit specialization and certificate; not a new criterion or closure law
One-unit Gaussian Fisher split for ϕ'^2 [21]	spherical local bridge and Gaussian all- p laws	new coefficient cross-identification and extensions; not a new one-unit decomposition
Bernoulli exponents and finite-sample logistic Hessian bounds [14, 17, 16, 18]	pointwise and shell-uniform relative Gram bounds, plus matching pointwise complexity	tunable precision and all powers; not the first bilateral, uniform, or one-sided Hessian bound
Classical spherical representation [7] and regular-design reference exponents [16]	exact Q_R constants and inverse-radius phase transition	exact-constant and regularly-varying-boundary refinement; not a new spherical representation or first non-Gaussian exponent
Generalized-logistic transforms and cumulants [23, 22]	Gamma/polygamma moments inside neuronal anisotropy	transform is classical; the unified anisotropy, endpoint laws and monotonicity are the contribution
Lam’s historical v2 Fisher–Schwarzian construction [24]	Löwner–anisotropy coefficient identity	different objects and coefficient; no claim to the first general Fisher–Schwarzian connection

Table 1: Result-by-result priority boundary.

Löwner matrices characterize fixed-order matrix monotonicity; modern treatments include Hiai–Sano and Heinävaara [1, 2]. Kozlovski and Sands prove the relevant classical equivalence between $Sf \geq 0$ and order-two matrix monotonicity and record the Schwarzian composition law [3]. Hence the qualitative logistic obstruction and its serial-chain extension are classical consequences; negative Schwarzian also implies the all-pair determinant sign through strict convexity of $(f')^{-1/2}$. Cook, Hammerlindl and Tucker develop the algebraically equivalent nowhere-coexpanding inequality, prove serial closure, and explicitly discuss sigmoid activations and one-dimensional neural compositions [4]. Neural papers have proposed analytic quantum activations and trainable matrix activations [5, 6], but their objectives are circuit implementation or matrix-valued parameterization, not preservation of PSD order by the spectral logistic function. In the primary sources we located, we did not find the closed logistic identity (3), the explicit certified witness, or the coefficient identity (27) linking the Löwner defect to Gaussian radial–tangential anisotropy. This is a bounded novelty statement, not a claim that no earlier source can exist.

Learning one neuron is already a nontrivial optimization problem. Yehudai and Shamir study when gradient methods learn a single neuron under broad activations and input distributions, including positive and negative results [8]. Diakonikolas et al. treat monotone neurons with adversarial label noise and obtain polynomial learners for logistic activations under log-concave inputs [9]. Wu extends learnability results to nonmonotone activations [10], and Vardi et al. show that adding a bias qualitatively changes the ReLU landscape [11]. Recent single-index work analyzes SGD under anisotropic Gaussian inputs [12]. Under Gaussian logistic regression, Hsu and Mazumdar study the signal-dependent difficulty of direction and temperature estimation [13]. Chen and Mazumdar identify the Bernoulli radial and tangential Hessian functions, prove their r^{-3} and r^{-1} orders and show tangential curvature exceeds radial curvature for every nonzero signal [14]; they subsequently prove a minimax norm-estimation rate of order $\sqrt{r^3/n}$ [15]. Ostrovskii and Bach prove a uniform bilateral constant-factor empirical-Hessian sandwich for Gaussian logistic regression [17, Eq. (92)], while Fisher et

al. give a generic fixed-point standardized-Hessian concentration framework [18]. Chardon, Lerasle and Mourtada prove, for $p = 1$, a uniform one-sided empirical Hessian lower bound under the sufficient condition $n \gtrsim r(d + t)$ [16, Theorem 6]. We therefore do not claim the Bernoulli exponents, the $p = 1$ ordering, the effective transition-slab scale, the first bilateral Hessian concentration, or the first finite-sample lower bound. Our empirical contribution is the explicit pointwise $(1 \pm \varepsilon)$ approximation (40) for all $p > 0$, its matching bilateral lower complexity, and the full-dyadic-shell extension (50). The latter pays a conservative covering logarithm and is not asserted to improve the optimal local one-sided rate. Recent work already constructs minimax labeled estimators, including a split-sample direction/norm scheme with bias correction [15]; our oracle Gram theorem is not an estimator.

Fisher geometry and natural gradient are classical [19]. Karakida et al. study broad Fisher spectral statistics for random deep networks [20]; Amari, Karakida and Oizumi also give the direct one-unit Gaussian ϕ'^2 decomposition and bias coupling used here [21]. Our radial contribution is therefore not the two-eigenspace reduction. It is the profile-generic moment theorem, nonasymptotic brackets, closed all-power logistic constants, strict global monotonicity, and their connection through (27) to the Löwner order defect. The generalized-logistic Gamma transform and its polygamma cumulants in Lemma 7.1 are classical [23, 22]; the claim is their neural-geometric use, ratio monotonicity and endpoint package, not discovery of the distributional identity. Lam’s v2 preprint connected Schwarzian curvature with Fisher–Rao geometry on manifolds of densities [24]. The current v3 explicitly corrects and supersedes v1–v2 and removes that architecture [25]. Neither version contains the one-neuron radial–tangential coefficient or its equality with the two-point Löwner defect. We cite v2 only as a dated historical adjacency; its L^p parameter is unrelated to the sensitivity exponent in $h_p = \sigma'^p$.

12 Reproducibility and hostile checks

The supplied script `repro/verify_saturation_law.py` performs the following independent checks:

1. it spot-checks (3) at separated scales and constructs Proposition 2.2 with 256-bit directed Arb balls;
2. it checks the spherical fourth-moment coefficient for Gaussian and fixed-radius spherical inputs;
3. it compares numerical quadrature of m_0, m_2, m_4 with the Gamma/polygamma formula for several noninteger and integer p ;
4. it checks the small- r jet through r^6 and the refined large- r expansion;
5. it evaluates α, β directly over a logarithmic radius grid and checks positivity and grid monotonicity for $p = 1, 2$;
6. it checks (29)–(31) at declared finite scales;
7. for three declared fixed random seeds it evaluates the empirical radial, tangential and cross blocks and verifies that those realized errors lie inside (38);
8. it checks both sides of (19)–(20) at declared finite radii;
9. a second replay evaluates the explicit C_1, C_2 majorants and lower-bound constants; and
10. a fixed-seed replay regenerates Figure 2 and its JSON certificate, while the main replay regenerates Figure 1 and its certificate.

The numerical monotonicity and finite-radius checks are performed on declared grids; the universal statements are proved analytically above. The replay uses Python, python-flint, NumPy, SciPy, and Matplotlib. It is diagnostic rather than a substitute for the analytic proofs, which reduce to the Löwner determinant identity, the Schwarzian chain rule, Gaussian orthogonal decomposition, Bernstein/net bounds, an elementary exponential bound, and closed Fourier moments. The stable record, versioned source archives, audit reports and manifests are public at <https://arr-research.github.io/papers/ARR-2026-53CTRKDSP685PT51/>.

13 Limitations and next theorem

The Löwner theorem concerns spectral matrix functions, not entrywise vector activations; it does not say that a standard feedforward neuron reverses coordinatewise order. The empirical results assume Gaussian input. The local Schwarzian coefficient needs fourth moments, whereas the global spherical law needs transition-layer regularity and $\mathbb{E}R^{-1} < \infty$; covariance isotropy alone is insufficient by Proposition 5.2. Elliptical laws can be whitened, but radial/tangential directions then use the covariance metric.

The sharp complexity statement is pointwise for a fixed teacher and bilateral relative Loewner loss. The dyadic-shell theorem is uniform but pays a covering logarithm. Neither is an estimator guarantee using labels or a minimax parameter lower bound. A weaker target can have a different threshold: the radial scalar needs no dimension factor in its upper bound, while no matching lower threshold is proved here for the bottom eigenspace. Proving such an eigenspace lower bound, or removing the shell covering logarithm, is the strongest proof-complete successor direction. The optimization result remains local, and the Cramér–Rao statement concerns regular unbiased estimation.

AI assistance disclosure

AI tools assisted with public-literature and repository triage, proof stress testing, numerical replay, drafting, and PDF production. The author remains responsible for the mathematical claims and any remaining errors.

References

- [1] F. Hiai and T. Sano, *Löwner matrices of matrix convex and monotone functions*, Journal of the Mathematical Society of Japan 64(2) (2012), 343–364. <https://arxiv.org/abs/1007.2478>
- [2] O. Heinävaara, *Characterizing matrix monotonicity of fixed order on general sets*, preprint (2019). <https://arxiv.org/abs/1906.06155>
- [3] O. Kozlovski and D. Sands, *Higher order Schwarzian derivatives in interval dynamics*, Fundamenta Mathematicae 206 (2009), 217–239. <https://arxiv.org/abs/0812.2646>
- [4] A. Cook, A. Hammerlindl and W. Tucker, *Nowhere coexpanding functions*, Chaos 33 (2023), 123105. <https://arxiv.org/abs/2303.12814>
- [5] M. Maronese, C. Destri and E. Prati, *Quantum activation functions for quantum neural networks*, Quantum Information Processing 21 (2022), 128. <https://arxiv.org/abs/2201.03700>
- [6] Z. Liu, S. Cao, Y. Li and L. Zikatanov, *Neural networks with trainable matrix activation functions*, Journal of Machine Learning for Modeling and Computing 6(2) (2025), 1–11. <https://arxiv.org/abs/2109.09948>
- [7] S. Cambanis, S. Huang and G. Simons, *On the theory of elliptically contoured distributions*, Journal of Multivariate Analysis 11(3) (1981), 368–385. [https://doi.org/10.1016/0047-259X\(81\)90082-8](https://doi.org/10.1016/0047-259X(81)90082-8)

- [8] G. Yehudai and O. Shamir, *Learning a Single Neuron with Gradient Methods*, Proceedings of Machine Learning Research 125 (2020), 3756–3786. <https://proceedings.mlr.press/v125/yehudai20a.html>
- [9] I. Diakonikolas, V. Kontonis, C. Tzamos and N. Zarifis, *Learning a Single Neuron with Adversarial Label Noise via Gradient Descent*, Proceedings of Machine Learning Research 178 (2022), 4313–4361. <https://proceedings.mlr.press/v178/diakonikolas22c.html>
- [10] L. Wu, *Learning a Single Neuron for Non-monotonic Activation Functions*, Proceedings of Machine Learning Research 151 (2022), 4178–4197. <https://proceedings.mlr.press/v151/wu22c.html>
- [11] G. Vardi, G. Yehudai and O. Shamir, *Learning a Single Neuron with Bias Using Gradient Descent*, Advances in Neural Information Processing Systems 34 (2021), 28690–28700. <https://arxiv.org/abs/2106.01101>
- [12] G. Braun, M. H. Quang and M. Imaizumi, *Learning a Single Index Model from Anisotropic Data with Vanilla Stochastic Gradient Descent*, Proceedings of Machine Learning Research 258 (2025), 1216–1224. <https://proceedings.mlr.press/v258/braun25a.html>
- [13] D. Hsu and A. Mazumdar, *On the Sample Complexity of Parameter Estimation in Logistic Regression with Normal Design*, Proceedings of Machine Learning Research 247 (2024), 2418–2437. <https://proceedings.mlr.press/v247/hsu24a.html>
- [14] J. Chen and A. Mazumdar, *Finite-Sample Performance of Gradient Descent in Logistic Regression with Gaussian Design*, arXiv:2606.21683 (2026). <https://arxiv.org/abs/2606.21683>
- [15] J. Chen and A. Mazumdar, *Minimax Optimal Estimator and Improved Error Rate for the MLE in Logistic Regression with Gaussian Design*, arXiv:2608.17260 (2026). <https://arxiv.org/abs/2608.17260>
- [16] H. Chardon, M. Lerasle and J. Mourtada, *Finite-sample performance of the maximum likelihood estimator in logistic regression*, arXiv:2411.02137v3 (first posted 2024; v3 2026). <https://arxiv.org/abs/2411.02137v3>
- [17] D. M. Ostrovskii and F. Bach, *Finite-sample analysis of M-estimators using self-concordance*, Electronic Journal of Statistics 15(1) (2021), 326–391. <https://doi.org/10.1214/20-EJS1780>
- [18] J. Fisher, L. Liu, K. Pillutla, Y. Choi and Z. Harchaoui, *Influence Diagnostics under Self-concordance*, Proceedings of Machine Learning Research 206 (2023), 10028–10076. <https://proceedings.mlr.press/v206/fisher23a.html>
- [19] S.-i. Amari, *Natural Gradient Works Efficiently in Learning*, Neural Computation 10 (1998), 251–276. <https://doi.org/10.1162/089976698300017746>
- [20] R. Karakida, S. Akaho and S.-i. Amari, *Universal Statistics of Fisher Information in Deep Neural Networks: Mean Field Approach*, Proceedings of Machine Learning Research 89 (2019), 1032–1041. <https://proceedings.mlr.press/v89/karakida19a.html>
- [21] S.-i. Amari, R. Karakida and M. Oizumi, *Fisher Information and Natural Gradient Learning in Random Deep Networks*, Proceedings of Machine Learning Research 89 (2019), 694–702. <https://proceedings.mlr.press/v89/amari19a.html>
- [22] C. J. Lee, A. Zito, H. Sang and D. B. Dunson, *Logistic-beta processes for dependent random probabilities with beta marginals*, Bayesian Analysis 20(4) (2025), 1345–1369. <https://doi.org/10.1214/25-BA1541>
- [23] M. O. Ojo and A. K. Olapade, *On a Six-Parameter Generalized Logistic Distribution*, Kragujevac Journal of Mathematics 26 (2004), 31–38. https://imi.pmf.kg.ac.rs/kjm/pub/12616736649184_5.pdf
- [24] H. P. G. Lam, *Real Bers embedding on the line: Fisher–Rao linearization, Schwarzian curvature, and scattering coordinates*, arXiv:2602.07373v2 (2026). <https://arxiv.org/abs/2602.07373v2>
- [25] H. P. G. Lam, *Zero-energy scattering and the real Bers image on the line*, arXiv:2602.07373v3 (2026), correcting and superseding v1–v2. <https://arxiv.org/abs/2602.07373v3>