Lawfulness, Residuals, and Evidence

A conceptual companion to sequential algebraic coherence in curve-decoding

Jonathan R. Landers

July 2026

Abstract

This note isolates the conceptual lens behind a connection between proximity gaps, a flight-path mixture model, and a statistical extension of curve-decodability. The common question is not merely whether observations are close, but what repeated closeness says about an underlying law. In curve-decoding, sufficiently many local codeword matches cannot remain unrelated: their row span must contain a common low-degree polynomial explanation. In the flight-path model, a proposed single mechanism makes a sharp distributional prediction; its failure forces multiple regimes, after which persistent local residuals identify one regime for an individual path. The algebraic extension acts on the same matrix used in the curve-decoding formalism, but replaces a structural yes-or-no statement with a residual and a null rarity rate. The resulting principle is: identify the simplest lawful subspace, remove what it explains, and measure how surprising the remaining coherence would be under an accidental mechanism.

1 The original structural move: closeness forces lawfulness

Key conceptual points

- A single close point may be accidental; many close points along a structured curve are not free to be arbitrary.
- The row span is the object that remembers whether the nearby codewords share one polynomial law.
- Their proximity result is local-to-global: enough local agreement forces global algebraic coherence.

Let $C \subseteq \mathbb{F}_q^n$ be a linear code and let

$$u(\alpha) = \sum_{j=0}^{\ell} u_j \alpha^j$$

be a degree-ℓ curve in the ambient word space. Suppose that for many parameters $\alpha \in A$, there is a codeword $f(\alpha) \in C$ satisfying

$$\Delta\big(u(\alpha), f(\alpha)\big) \leq \delta.$$

The structural question is whether these local decoding successes can be governed by unrelated explanations. Curve-decodability says that, once there are enough of them, many must lie on one codeword-valued polynomial curve,

$$f(\alpha) = c(\alpha), \qquad c(\alpha) = \sum_{j=0}^{\ell} c_j \alpha^j, \quad c_j \in C.$$

Place the nearby codewords into the columns of

$$M = [f(\alpha_1) \ \cdots \ f(\alpha_m)], \qquad U = \mathrm{RowSpan}(M).$$

For a subset $B \subseteq A$, the columns indexed by B obey one degree-ℓ law exactly when

$$\pi_B(U) \subseteq \mathrm{RS}_\ell(B), \qquad \mathrm{RS}_\ell(B) = \{(p(\alpha))_{\alpha \in B} : \deg p \leq \ell\}.$$

Their key matrix move is therefore to preserve U, because coordinatewise closeness alone does not remember how the candidate codewords relate across the whole collection [1].

Their arc

many local proximity events $\implies$ row-span containment $\implies$ one shared polynomial law.

The intuitive statement is: the points are *too consistently close along a structured object* for every match to come from a different underlying lawfulness.

2 The flight-path move: a law predicts a distribution

Key conceptual points

- A proposed mechanism is useful only if it predicts observable features that can fail.
- In the flight model, broad random scale should wash out ordinary bridge-shape non-Gaussianity.
- The observed distribution is too different from that prediction, so one universal mechanism is not enough.

The flight-path note begins in the opposite direction. Rather than starting from many local matches and asking what law they force, it starts from one proposed law and asks what distribution that law must produce [2]. The RMS departure from an endpoint geodesic is written

$$D = AQ,$$

where A is route scale and Q is the dimensionless shape of an endpoint-pinned bridge. With $U = \log(A/d_0)$ and $V = \log Q$,

$$\log(D/d_0) = U + V.$$

If U is Gaussian and independent of V, broad variation in U suppresses the standardized higher cumulants contributed by the bridge shape. Writing

$$\rho = \frac{\mathrm{Var}(V)}{\mathrm{Var}(U) + \mathrm{Var}(V)},$$

one obtains

$$\mathrm{Skew}(\log D) = \mathrm{Skew}(\log Q)\rho^{3/2}, \qquad \mathrm{ExKurt}(\log D) = \mathrm{ExKurt}(\log Q)\rho^{2}.$$

Thus one mechanism predicts an almost Gaussian law in log space. Under the Brownian-bridge benchmark, the fitted scale predicts approximately

$$\mathrm{Skew}(\log D) \approx 0.003, \qquad \mathrm{ExKurt}(\log D) \approx -0.002,$$

whereas the observed cohort has

$$\mathrm{Skew}(\log D) = -2.40, \qquad \mathrm{ExKurt}(\log D) = 8.11.$$

The discrepancy is not a minor defect in the center of the histogram. It says that at least one defining assumption of the single-regime explanation is false. The visible near-geodesic boundary then motivates a mixture,

$$f_D(d) = \pi_0 f_0(d) + \sum_{c=1}^{C} \pi_c f_D(d \mid c).$$

The cohort-level flight arc

one proposed mechanism $\Longrightarrow$ a sharp distributional signature $\Longrightarrow$ observed mismatch $\Longrightarrow$ more than one mechanism.

This is not the opposite of curve-decodability so much as its conceptual complement. Their theorem says that repeated compatibility forces common structure. The flight argument says that a global structure must earn its status by surviving a distributional test.

3 Persistence turns local evidence into mechanism assignment

Key conceptual points

- A short straight-looking prefix is compatible with either flight regime; persistence is what becomes diagnostic.
- Evidence accumulates when the null mechanism must repeatedly spend probability mass to imitate the alternative.
- This sequential logic is the bridge from the flight problem back to curve-decoding.

After the mixture is forced, the flight problem reverses direction. Let $C = 0$ denote the near-geodesic boundary regime and $C = 1$ the ordinary bridge regime. For the first n observations, define the residual energy

$$R_n^2 = \frac{1}{n} \sum_{j=1}^{n} \eta(t_j)^2.$$

If persistent small residuals are exponentially rare under the ordinary population,

$$\Pr_1(R_n \le \tau) \le e^{-nI(\tau)},$$

then every additional small-residual observation spends more null probability. Bayes' rule turns that expenditure into increasing evidence for the boundary mechanism:

$$\Pr(C = 0 \mid R_n \le \tau) \gtrsim \frac{\pi_0}{\pi_0 + (1 - \pi_0)e^{-nI(\tau)}}.$$

The global cohort rejects one law; the local sequence then identifies which law is active for one trajectory.

The full flight arc

wrong global prediction $\Longrightarrow$ multiple regimes $\Longrightarrow$ persistent local residual $\Longrightarrow$ one regime assignment.

The key transferable idea is not "use a mixture model." It is the more general rule

$$\boxed{\text{local compatibility becomes global evidence only after the null law has paid for its persistence.}}$$

4 The new matrix lens: from row span to residual evidence

Key conceptual points

- Their row span identifies what a common algebraic law looks like.
- The new rank residual measures what remains after that lawful component is removed.
- The entropy bound measures how unlikely exact coherence would be under an accidental decoder law.

The statistical extension acts on the same matrix M, but asks a different question. Let $H_{\ell,m}$ be a parity-check matrix for the degree-ℓ evaluation space on $A_m = \{\alpha_1, \dots, \alpha_m\}$. Define

$$r_m = \operatorname{rank}\left(M_m H_{\ell,m}^{\mathsf{T}}\right).$$

Because $H_{\ell,m}$ annihilates every degree-ℓ evaluation vector,

$$r_m = 0 \quad \Longleftrightarrow \quad \operatorname{RowSpan}(M_m) \subseteq \operatorname{RS}_\ell(A_m).$$

Thus r_m is an algebraic residual: it counts independent row-space directions that the proposed polynomial law cannot absorb. It is the finite-field analogue of residual bridge energy.

There is also an interpolation barrier. Any $\ell + 1$ values can be fit by a degree-ℓ polynomial, so

$$m \le \ell + 1 \quad \Longrightarrow \quad r_m = 0 \quad \text{under both coherent and accidental explanations.}$$

Evidence begins only at observation $\ell + 2$. If the next decoder output has conditional point mass at most q^{-h_i} after all previous outputs and proximity information are known, then accidental exact coherence satisfies

$$\Pr_{\text{acc}}(r_m = 0) \le q^{-\sum_{i=\ell+2}^{m} h_i}.$$

This is the direct algebraic analogue of the small-ball rate $e^{-nI(\tau)}$. The row-span theorem tells us what lawfulness is; the entropy bound tells us how much evidential weight to assign when that lawfulness appears.

The combined matrix arc

row span identifies the lawful subspace $\Longrightarrow$ syndrome rank measures the unexplained part $\Longrightarrow$ conditional entropy prices accidental coherence.

At the same conceptual level, the division of labor is

Structural statement:	enough local proximity forces a coherent subset;
Inferential statement:	the observed coherence is expensive under an accidental law.

The first says that a law exists. The second says how strongly the data support that law over a competing explanation.

5 A reusable conceptual template

Key conceptual points

- "Distance" enters three times: distance to a model, residual outside the lawful subspace, and distance between competing probability laws.
- A residual alone is not evidence; it becomes evidence only when paired with a null rarity rate.
- The general method applies whenever a simple mechanism defines a low-complexity set of lawful observations.

The common lens can be separated into four objects:

$$\boxed{\begin{array}{c}\text{candidate lawful class } \mathcal{L} \\ \Downarrow \\ \text{residual } \mathcal{R}(X;\mathcal{L}) \\ \Downarrow \\ \text{null rarity } \Pr_0\{\mathcal{R} \le \tau\} \\ \Downarrow \\ \text{global mechanism statement.}\end{array}}$$

For flight paths, $\mathcal{L}$ is an ordinary bridge regime, $\mathcal{R}$ is prefix RMS residual energy, and the null rarity is a small-ball probability. For curve-decoding, $\mathcal{L}$ is the Reed–Solomon evaluation subspace, $\mathcal{R}$ is syndrome rank, and the null rarity is controlled by conditional min-entropy.

This clarifies three distinct meanings of closeness:

$$\begin{aligned}\text{geometric closeness} &\quad:\quad \Delta(u(\alpha), C) \le \delta, \\ \text{algebraic closeness} &\quad:\quad r_m \text{ is zero or small}, \\ \text{distributional separation} &\quad:\quad W_p(P_{\text{law}}, P_{\text{acc}}) \text{ is large.}\end{aligned}$$

They should not be conflated. Geometric closeness supplies candidate matches. Algebraic residual asks whether the matches share one low-complexity explanation. Distributional separation asks whether that explanation can be distinguished from an accidental mechanism at finite sample size.

The broad research direction is therefore not merely to prove more proximity gaps. It is to build a statistical geometry of lawfulness: quantify approximate residuals, estimate entropy rates for actual decoders, compare finite-length residual laws across code ensembles, and construct sequential tests whose false-alarm guarantees survive optional stopping. The same lens may apply wherever repeated local compatibility is suspected to reveal a shared latent rule.

Conceptual endpoint

A structural theorem says that the observations must admit a common law.
A residual says what that law fails to explain.
A probability bound says whether the law is genuine evidence or an inexpensive accident.

References

[1] R. Goyal, V. Guruswami, Y. Sun, and M. Wootters, "Locality of Curve-Decoding and Improved Proximity Gaps," *arXiv preprint arXiv:2607.08516*, 2026. `https://arxiv.org/abs/2607.08516`.

[2] J. R. Landers, "When Does a Random-Scale Flight Path Look Lognormal? A Spherical Bridge Note from OpenSky Trajectories," GitHub repository, 2026. `https://github.com/jonland82/aerostatsome`.

[3] J. R. Landers, "Sequential Detection of Algebraic Coherence in Curve-Decoding," unpublished technical companion, 2026.

[4] C. Villani, *Optimal Transport: Old and New*. Springer, 2009.