From Syntax to Semantics: Geometric Stability as the Missing Axis of Perturbation Biology

Prashant C. Raju
rajuprashant@gmail.com

Abstract

The capacity to precisely edit genomes has outpaced our ability to predict the consequences. A cell can be genetically perfect and therapeutically useless: edited exactly as intended, yet unstable, drifting toward unintended fates, or selected for properties that compromise safety. This paradox reflects a deeper gap in how we evaluate biological intervention. Current frameworks excel at measuring what was done to a cell but remain blind to what the cell has become. We argue that this blindness stems from treating cells as collections of independent variables rather than as dynamical systems occupying positions on high-dimensional state manifolds. Drawing on Waddington's epigenetic landscape, we propose geometric stability as a missing axis of evaluation: the directional coherence of cellular responses to perturbation. This metric distinguishes interventions that guide cells coherently toward stable states from those that scatter them across the state manifold. Validation across diverse perturbation datasets reveals that geometric stability captures regulatory architecture invisible to conventional metrics, discriminating pleiotropic master regulators from lineage-specific factors without prior biological annotation. As precision medicine increasingly relies on cellular reprogramming, the question shifts from "did the intervention occur?" to "is the resulting state stable?" Geometric stability provides a framework for answering.

The Syntax-Semantics Gap

Biology is navigating a transition analogous to the shift from assembly code to high-level programming: we have mastered the ability to write characters (DNA), but we struggle to compile stable programs (cell states). For half a century, the central dogma provided a reductionist map: DNA to RNA to Protein. In the era of observation, this map was sufficient. We sequenced, catalogued, and mapped. In the era of active authorship, where we edit genomes to cure disease and engineer cell therapies (Boeke et al. 2016), it is failing.

The crisis manifests as a paradox: a cell can be genetically perfect and therapeutically useless. We treat genes as independent variables in a linear equation, assuming that a precise edit at a specific locus will produce a contained, predictable effect. But cells are complex adaptive systems existing on high-dimensional manifolds (Moon et al. 2019; Wagner, Regev, and Yosef 2016). In this nonlinear regime, identical genetic perturbations can trigger divergent phenotypic trajectories depending on the cell's initial position in state space (Jensen et al. 2017) or the local curvature of the regulatory landscape.

Current evaluation metrics, designed for the linear paradigm, measure the "syntax" of the edit: indel rates, off-target cleavage, and sequence fidelity (Brinkman et al. 2014; Tsai et al. 2014). These metrics answer the engineer's question: "Did I change the code?" They fail to answer the biologist's question: "Did the system stabilize?" By focusing on sequence fidelity rather than state stability, we miss the structural failures that define the success or failure of a perturbation.

Three such failures illustrate what the linear paradigm cannot see. First, selection artifacts. In human pluripotent stem cells, Cas9-induced breaks trigger p53-mediated toxicity even when targeting non-essential genes. The cells that survive are often those that have spontaneously acquired p53 mutations: a "successful" edit that has inadvertently selected for oncogenic potential (Haapaniemi et al. 2018; Ihry et al. 2018). Second, regulatory disruption. We now know that even "clean" integrations can disrupt 3D genome organization or enhancer dynamics, pushing cells toward malignancy without

mutating coding sequences (Cavazza, Moiani, and Mavilio 2013; Lupiáñez et al. 2015). On-target edits can trigger large deletions or chromothripsis invisible to standard sequencing (Kosicki, Tomberg, and Bradley 2018; Leibowitz et al. 2021). Third, phenotypic heterogeneity. It is increasingly clear that two cells carrying the exact same edit often behave differently: one differentiating, one remaining stem-like. This is not noise. It reflects the initial position of each cell on the state manifold and the local geometry of that landscape (Replogle et al. 2022; Weinreb et al. 2020).

These failures share a common feature: they occur in cell state space, not sequence space. We have mastered the syntax of the genome (Doudna and Charpentier 2014; Jiang and Doudna 2017; Jinek et al. 2012). We remain blind to its semantics. To move forward, we must recognize that stability is not a property of the sequence. It is a geometric property of the state.

The Manifold Reality

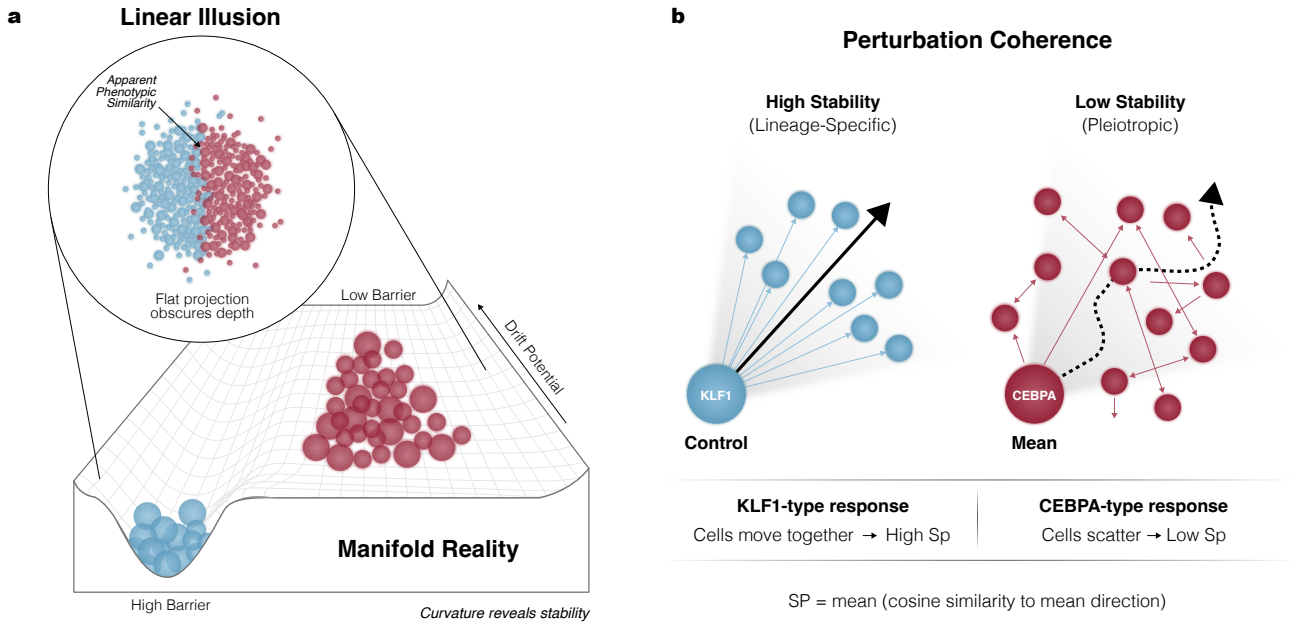

Figure 1: **The Geometric Tax: linear metrics obscure biological stability. a.** Standard dimensionality reduction projects high-dimensional cell states onto a flat plane (Linear Illusion, inset), where two populations (blue, red) appear to overlap, suggesting similar phenotypes. Mapping these populations onto the underlying biological manifold (Manifold Reality) reveals distinct stability properties invisible to linear projections. The blue population occupies a deep valley (high barrier), representing a robust cell state resistant to perturbation. The red population sits on a shallow ridge (low barrier), representing an unstable state prone to drift. This stability difference constitutes the Geometric Tax of engineering cells into non-native configurations. **b.** Geometric stability quantified through perturbation coherence. High-stability perturbations (left, e.g., KLF1) produce shift vectors that align coherently, indicating cells move together along a shared trajectory toward the mean direction (solid arrow). Low-stability perturbations (right, e.g., CEBPA) scatter cells in divergent directions despite similar magnitude shifts, with the mean direction (dashed arc) representing dispersed cellular responses. The Shesha stability score (Sp) captures this distinction as the mean cosine similarity between individual shift vectors and the population mean. Together, panels a and b demonstrate how manifold curvature, invisible to linear projections, determines whether perturbations produce stable or fragile cellular states.

To resolve the syntax-semantics gap, we must pivot from a sequence-centric view of biology to a geometric one. This is not a radical invention but a return to the foundational insight of developmental mechanics: the epigenetic landscape. When Conrad Waddington depicted cell development as a ball rolling down an undulating surface (Slack 2002; Waddington 1957), he was not merely offering an illustration (Fard et al. 2016). He was describing the topology of a dynamical system (Ferrell 2012). In this view, "valleys" are not mere metaphors; they are attractor basins, stable regions of state space where regulatory networks minimize the system's quasi-potential energy (Enver et al. 2009; Huang 2009; J. Wang et al. 2011).

Modern single-cell genomics has transformed this topology from a concept into a measurable reality (Rand et al. 2021). We

now understand that a human cell, expressing approximately 20,000 genes, occupies a point in a high-dimensional phase space. Gene regulatory networks function as the vector field governing this space (A. Raju and Siggia 2023), defining the flow that pulls cells toward attractors (stable cell types) and separates them via repellers (unstable intermediates). The fundamental law of this physics is clear: stability is geometry. A deep, steep-walled basin represents a robust state where perturbations are rapidly dampened by restoring forces, a property historically termed canalization (Siegal and Bergman 2002). A shallow, flat plateau represents fragility, where transcriptional noise dominates and small fluctuations drive lineage drift (Elowitz et al. 2002; Mojtahedi et al. 2016). While early maps of this terrain were theoretical, new computational frameworks are finally allowing us to reconstruct these landscapes directly from data (Cislo et al. 2025; Han et al. 2025).

Yet our analytical toolkit remains stubbornly Newtonian in an Einsteinian world. Principal Component Analysis (PCA), the workhorse of perturbation biology (Tsuyuzaki et al. 2020), commits the "Euclidean Error": it forces the complex, curved geometry of the manifold onto a flat linear hyperplane (Moon et al. 2019; Zhou and Sharpee 2021). By measuring distances along straight lines rather than geodesic curves, we distort the very signal we seek to measure. Two cells might appear effectively identical in a 2D projection, yet be separated by a high energetic barrier that prevents state transition. In flattening the manifold, we erase the topography that defines stability. We measure the shadow of the mountain, not the mountain itself.

Recovering the "semantics" of a perturbation therefore requires measuring the curvature of the manifold (Moon et al. 2019). A cell pushed into a deep attractor moves coherently: regulatory constraints channel the perturbation vector along a defined trajectory (Huang 2009). A cell pushed onto a flat plain scatters: lacking strong restoring forces, the population diffuses stochastically (Elowitz et al. 2002; Mojtahedi et al. 2016). This geometric divergence, coherence versus scatter, is the signature of biological meaning. It distinguishes a specific signal from entropic noise.

Quantifying Geometric Stability

If stability is geometry, we need metrics that measure geometry rather than magnitude. Standard perturbation analysis asks: how far did cells move in expression space? This captures effect size but misses a critical dimension (Dixit et al. 2016; Replogle et al. 2022). The question that predicts functional outcomes is different: did cells move *together*? A perturbation that shifts cells coherently along a shared trajectory has engaged a robust regulatory program. A perturbation that scatters cells in divergent directions has pushed them onto an unstable region of the manifold where regulatory constraints are weak (Scheffer et al. 2009). The distinction is invisible to magnitude-based metrics but determines whether an edit produces a stable therapeutic product or a heterogeneous population primed for failure.

The Shesha stability score operationalizes this intuition through directional coherence (P. C. Raju 2026a). Consider the hydrodynamics of a river: laminar flow moves water molecules in parallel layers, preserving local topology even as mass shifts downstream. Turbulent flow generates chaotic eddies where adjacent molecules scatter in orthogonal directions. Both regimes may transport mass the same distance, but only one preserves the structural integrity of the stream. Shesha measures whether a genetic perturbation induces laminar flow (coherent state transition) or turbulence (entropic scatter).

A clarification is necessary. We argued above that PCA distorts global manifold geometry, and this remains true. However, what we seek to measure is not global distance but *local directional coherence*: whether cells receiving the same perturbation move in similar directions relative to controls. Within the dominant principal components that capture the majority of biological variance, local angular relationships are preserved well enough to extract this signal. PCA serves as a tractable first approximation, computationally efficient and compatible with standard single-cell workflows (Tsuyuzaki et al. 2020). The path to refinement is clear: diffusion maps, UMAP with appropriate parameterization (McInnes et al. 2018), or learned embeddings from foundation models (Theodoris et al. 2023) offer progressively better manifold representations. But even in linear PCA space, directional coherence distinguishes biologically meaningful patterns, as validation demonstrates.

Box 1 | Measuring Geometric Stability: The Shesha Framework

We formalize stability as directional coherence in reduced-dimensional space. Let $\mathbf{c}$ denote the centroid of unperturbed control cells, establishing the baseline state:

$$\mathbf{c} = \frac{1}{n_{\text{ctrl}}} \sum_i \mathbf{x}_i^{\text{ctrl}} \tag{1}$$

For each cell j receiving perturbation p, the shift vector $\mathbf{v}_j = \mathbf{x}_j^p - \mathbf{c}$ captures its displacement from baseline. The population's mean shift direction $\bar{\mathbf{v}}$ and its norm define effect magnitude:

$$\bar{\mathbf{v}} = \frac{1}{n_p} \sum_j \mathbf{v}_j, \quad \text{Magnitude} = \|\bar{\mathbf{v}}\| \tag{2}$$

Stability then follows as the average alignment between individual trajectories and the collective response:

$$S_p = \frac{1}{|V|} \sum_{j \in V} \frac{\mathbf{v}_j \cdot \bar{\mathbf{v}}}{\|\mathbf{v}_j\| \|\bar{\mathbf{v}}\|} \tag{3}$$

where $V = \{j : \|\mathbf{v}_j\| > 10^{-6}\}$ excludes cells with negligible displacement. The score ranges from -1 to 1: values approaching 1 indicate *laminar* response (cells moving in unison), while values near 0 indicate *turbulent* scatter (cells diffusing stochastically). Perturbations with fewer than 10 cells or mean magnitude below 10^{-6} are excluded to ensure statistical robustness.

Implementation. Shesha is available as a pip-installable Python package compatible with AnnData objects (Virshup, Bredikhin, et al. 2023; Virshup, Rybakov, et al. 2024): `pip install shesha-geometry` (P. C. Raju 2026b)

Validation across four independent single-cell CRISPR datasets confirms that stability captures consistent biological signal. Analyzing 422 perturbations across 212,865 cells from CRISPRa (Norman et al. 2019), CRISPRi (Adamson et al. 2016; Dixit et al. 2016), and pooled screens (Papalexi et al. 2021), we observe strong magnitude-stability correlations ranging from $\rho = 0.746$ to $\rho = 0.985$, with a pooled correlation of $\rho = 0.833$ (Table 1). This relationship is robust across methodological choices: Euclidean distance, Mahalanobis-whitened coordinates, and local k-nearest-neighbor centroids all yield correlations above 0.74. The consistency across datasets, modalities, and cell types suggests a universal geometric relationship between effect size and directional coherence. Yet the correlation is not perfect: cases where magnitude and stability diverge prove biologically informative, revealing regulatory architecture invisible to magnitude alone.

Table 1: Magnitude-stability Spearman correlations across CRISPR datasets with 95% bootstrap confidence intervals.

Dataset	Perturbations	Cells	ρ_{Mag}	95% CI	p
Norman 2019	236	99,420	0.953	[0.934, 0.965]	$< 10^{-100}$
Adamson 2016	8	5,752	0.929	[0.407, 1.000]	$< 10^{-3}$
Dixit 2016	153	89,350	0.746	[0.641, 0.827]	$< 10^{-100}$
Papalexi 2021	25	18,343	0.985	[0.939, 0.997]	$< 10^{-18}$
Pooled	422	212,865	0.833	—	$< 10^{-100}$

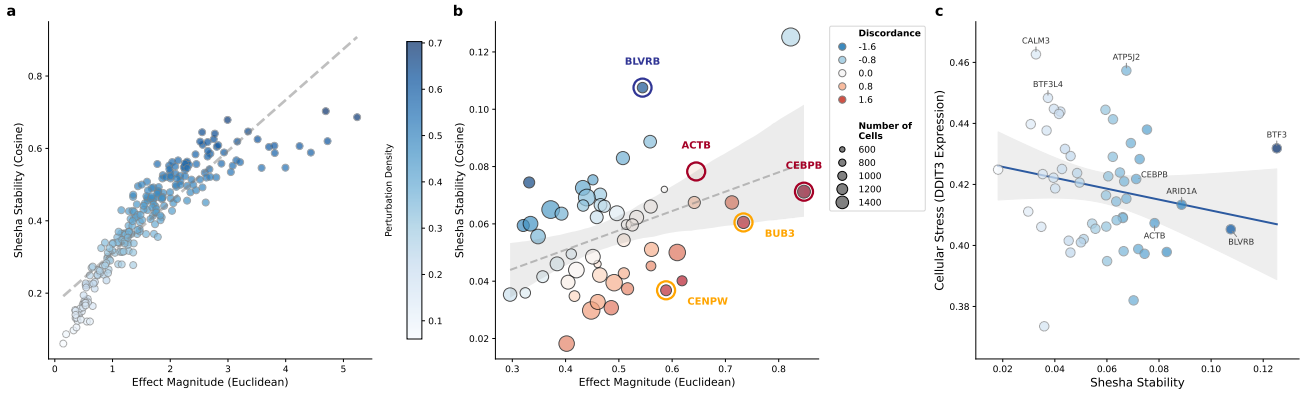

Figure 2: **Geometric stability validated across CRISPR datasets and linked to cellular stress. a.** Magnitude-stability relationship in Norman et al. CRISPRa dataset (Norman et al. 2019) ($n = 236$ perturbations). Shesha stability score correlates strongly with effect magnitude (Spearman $\rho = 0.953$, $p < 10^{-100}$). Color indicates local perturbation density. Dashed line shows linear fit. **b.** Independent validation in the Replogle et al. genome-scale CRISPRi screen (Replogle et al. 2022) (K562 cells). Point color indicates discordance (deviation from expected stability given magnitude); point size indicates cell count per perturbation. Labeled genes illustrate biological interpretation: BLVRB (biliverdin reductase, metabolic) shows high stability relative to magnitude, consistent with pathway-specific effects. CEBPB (C/EBP family transcription factor) and CENPW (centromere protein) show low stability relative to magnitude, consistent with pleiotropic or cell-division-wide effects. BUB3 (spindle checkpoint) demonstrates that low stability is not merely a proxy for cell cycle arrest. Grey shading indicates 95% confidence interval. **c.** Functional consequence of geometric instability. Stability negatively correlates with DDIT3 expression (Spearman $\rho = -0.28$, $p = 0.041$), a canonical marker of cellular stress (Oyadomari and Mori 2003). Perturbations producing incoherent cellular responses (low stability) induce elevated stress signatures. Point shading encodes stability (darker = higher). Linear fit shown with 95% confidence interval. Notably, no perturbations occupy the high-stability/high-stress quadrant, suggesting that geometric coherence is a prerequisite for cellular homeostasis.

Critically, however, the magnitude-stability trade-off is not an artifact of linear projection. Validation using embeddings from scGPT, a foundation model trained on 33 million cells that learns nonlinear representations of cell state (Cui, C. Wang, et al. 2024), confirms that this geometric relationship persists robustly in nonlinear latent space ($\rho = 0.935$, 95% CI $[0.911, 0.951]$, $p < 10^{-107}$ vs. $\rho = 0.953$ in PCA for Norman et al. 2019 (Norman et al. 2019)). While diffusion maps (Coifman and Lafon 2006), PHATE (Moon et al. 2019), and other manifold-aware embeddings may reveal additional structure, this indicates that the tax is a property of biological state space itself, independent of dimensionality reduction technique.

Discordance Reveals Regulatory Architecture

The strong correlation between perturbation magnitude and geometric stability might suggest these metrics are redundant. They are not. The divergent cases, where magnitude and stability decouple, expose a fundamental distinction in how genes shape the regulatory landscape. This divergence captures the difference between regulators that coordinate cellular programs and those that merely perturb them.

Master regulators pay a Geometric Tax. CEBPA, the canonical driver of myeloid differentiation, activates over 24 downstream pathways spanning immune response, cell cycle control, metabolism, and lineage commitment (Friedman 2007, 2015; Norman et al. 2019). When CEBPA is perturbed, cells move far from controls in transcriptomic space, but they scatter incoherently, each cell activating a different subset of CEBPA's broad target repertoire. This is not a failure of the perturbation; it is a consequence of regulatory architecture. A single node attempting to coordinate thousands of competing variables cannot maintain geometric coherence. The Tax is the cost of pleiotropy: large effects, diffuse outcomes. Despite ranking among the highest-magnitude perturbations in the Norman CRISPRa dataset, CEBPA combinations cluster below the regression line, their stability lower than effect size would predict.

Lineage-specific factors operate under different dynamical regime. KLF1, the erythroid-specific transcription factor essential for globin expression and terminal red blood cell maturation (Miller and J J Bieker 1993; Siatecka and James J. Bieker 2011; Tallack and Perkins 2010; Tallack, Whitington, et al. 2010), regulates a narrower but tightly coordinated gene

program. KLF1 perturbations produce moderate magnitude shifts but exceptional coherence: cells move together along a shared trajectory corresponding to erythroid differentiation arrest (Pilon et al. 2008). KLF1 pays minimal Geometric Tax because its targets are functionally aligned, part of a single regulatory module rather than distributed across the genome. The stability score captures this alignment without any prior knowledge of KLF1's biology.

This distinction, pleiotropic master regulators versus lineage-specific factors, generalizes across datasets and screening modalities. In an independent genome-scale CRISPRi screen with over 2,000 perturbations in K562 cells (Replogle et al. 2022), CEBPB, a member of the same transcription factor family as CEBPA (Jakobsen et al. 2013), shows the identical signature: high magnitude, low stability. CENPW, encoding a centromere protein required for chromosome segregation (Hori et al. 2008; Liu et al. 2024), shows even more pronounced discordance. Disrupting the cell division machinery produces large but geometrically incoherent responses as cells scatter into diverse failure states. Conversely, BLVRB shows high stability despite moderate magnitude, consistent with its role encoding biliverdin reductase, an enzyme in a specific metabolic pathway (O'Brien et al. 2015; Paukovich et al. 2018). The pattern is not dataset-specific. It reflects how regulatory architecture shapes the topology of state space.

The statistical signature confirms what the landscape framework predicts. Perturbations in the high-discordance quartile (high magnitude, low stability) exhibit nearly three-fold higher variance in stability scores ($SD = 0.202$) compared to low-discordance perturbations ($SD = 0.074$). Shallow valleys produce scatter; deep valleys produce coherence. The variance itself encodes the geometry.

The Geometric Tax has functional consequences beyond phenotypic heterogeneity. Across essential gene perturbations in the Replogle dataset, stability negatively correlates with DDIT3 expression (Spearman $\rho = -0.28$, $p = 0.041$), a canonical marker of the integrated stress response (Oyadomari and Mori 2003; Zinszner et al. 1998). Perturbations that scatter cells incoherently across transcriptomic space induce elevated stress signatures. Geometric instability is not merely a statistical abstraction; it manifests as cellular dysfunction.

One might argue that stability simply measures proliferation arrest: cells that stop dividing would appear artificially coherent. The data refute this. CENPW (Zhang et al. 2024) and BUB3 (Logarinho et al. 2008), both cell cycle regulators whose perturbation induces mitotic arrest, show low stability despite halting division. BLVRB (O'Brien et al. 2015) shows high stability while cells continue cycling. The metric captures regulatory coherence, not cessation of proliferation.

The implications extend beyond biological understanding. For CRISPR screens, geometric stability can prioritize hits likely to replicate, distinguishing perturbations that engage specific programs from those causing nonspecific toxicity (Tycko et al. 2019). For cell therapy manufacturing, stability can flag products at risk of lineage drift or functional heterogeneity before clinical deployment (Kliegman et al. 2024; Lipsitz, Timmins, and Zandstra 2016). Indel rates answer whether the edit occurred (Sentmanat et al. 2018). Stability answers whether the resulting state will persist.

Stability as a Design Principle

The preceding sections establish geometric stability as a measurable property of perturbations. But stability is more than a metric. It is a design principle that evolution discovered long before we learned to measure it.

Natural selection operates on phenotypes, not genotypes (Mayr 1963). A mutation that produces large transcriptional effects but scatters cells into incoherent states will be rapidly purged: such cells cannot coordinate the programs required for survival, proliferation, or function (Fisher 1930; Kirschner and Gerhart 1998). Conversely, mutations that engage robust regulatory modules, moving cells coherently along trajectories that preserve functional capacity, can propagate (Kirschner and Gerhart 1998; A. Raju, Xue, and Leibler 2023). Over evolutionary time, this selection pressure has carved the deep valleys of Waddington's landscape (Waddington 1957; J. Wang et al. 2011). The attractor basins we observe in single-cell data are not arbitrary; they are the stable configurations that have survived billions of years of geometric selection (Huang 2009; A. Raju and Siggia 2023; Wagner, Regev, and Yosef 2016). While early theoretical work showed that

such canalization emerges naturally in evolving networks (Siegal and Bergman 2002), recent studies have mathematically formalized Waddington's landscape, demonstrating how these geometric constraints direct evolutionary trajectories (A. Raju, Xue, and Leibler 2023). Thus, gene regulatory networks are optimized not merely for specific expression patterns but for the stability of those patterns under perturbation.

This framing recasts the challenge of biological engineering. When we edit a genome, we are not simply changing a sequence; we are perturbing a system that evolution has optimized for stability (Kitano 2004). The Geometric Tax is the cost of fighting this optimization. Edits that work with the grain of regulatory architecture, engaging existing modules and respecting attractor boundaries, pay minimal tax. Edits that work against it, forcing cells into states that evolution never selected for, pay heavily in the currency of heterogeneity, stress (Zinszner et al. 1998), and drift (Weinreb et al. 2020).

This is biology's alignment problem. In artificial intelligence, alignment refers to the challenge of ensuring that optimized systems pursue intended objectives rather than exploiting loopholes in their reward functions (Amodei et al. 2016; Hubinger et al. 2019; Krakovna et al. 2020; Russell 2019). In biological engineering, the analogous challenge is ensuring that edited cells maintain intended phenotypes rather than drifting toward unintended attractors. A CAR-T cell optimized for tumor killing may instead find a local minimum in exhaustion (Fraietta et al. 2018; Philip et al. 2017; Sen et al. 2016). An iPSC-derived beta cell may appear differentiated by marker expression yet occupy a shallow basin that permits dedifferentiation under metabolic stress (Lipsitz, Timmins, and Zandstra 2016; Nair et al. 2019; Talchai et al. 2012). Current evaluation frameworks, focused on sequence fidelity and marker expression, cannot detect misalignment until it manifests clinically (Bravery et al. 2013; Capelli et al. 2023; Fraietta et al. 2018; Salmikangas et al. 2015). Geometric stability offers an objective function for alignment: perturbations that produce high stability are those where the intended state coincides with a deep attractor.

The practical implications follow from this principle. In screening, stability prioritizes hits that engage robust regulatory programs over those that produce large but incoherent effects, predicting which perturbations will replicate across contexts (Tycko et al. 2019). In manufacturing, stability flags cell products balanced on flat manifold regions before clinical deployment, complementing marker-based QC with a measure of phenotypic robustness (Fraietta et al. 2018; Levine et al. 2017; Lipsitz, Timmins, and Zandstra 2016). In regulatory review, stability addresses a failure mode that current genotoxicity and tumorigenicity assays cannot access: the edited cell that is genetically clean but phenotypically fragile (U.S. Food and Drug Administration 2013, 2024).

This imperative extends beyond the bench to a transformation underway in computational biology. Foundation models trained on transcriptomic data, including scGPT (Cui, C. Wang, et al. 2024), Geneformer (Theodoris et al. 2023), and Universal Cell Embeddings (Rosen et al. 2023), learn representations where distances correspond to biological relationships. These models implicitly encode the geometry of cell state space and inherit the same fragility problem. Geometric stability operationalizes a question that matters equally for evaluating any learned biological embedding: do perturbations move cells coherently through latent space, or do they scatter them? A foundation model that places biologically stable states in deep basins and unstable intermediates on flat regions has learned something true about cellular dynamics. A model that scrambles this geometry has learned a representation that will fail when deployed. Geometric stability thus serves as a necessary criterion for representation quality.

The deeper point transcends any particular implementation. Evolution has spent billions of years optimizing gene regulatory networks for stability. We are only now developing the tools to measure what it optimized for. As biological engineering matures from sequence writing to state programming, geometric stability offers a bridge: between the intuitions of machine learning and the constraints of cellular dynamics, between what we intend and what cells become.

Finally, this framework resolves a critical challenge for the emerging era of generative biology. As foundation models gain the capability for *in silico* perturbation (Cui, Tejada-Lapuerta, et al. 2025), we face a new validation gap: distinguishing viable biological attractors from hallucinatory intermediates (Rathkopf 2025). Just as Shesha identifies unstable cells in experimental screens, it provides the necessary physical constraint for generative models, ensuring that computational

predictions remain faithful to biological dynamics. Geometric stability becomes the physics engine that keeps our digital explorations grounded in cellular reality.

What the Thunder Said

We have argued that the central challenge of biological engineering is not writing the genome but compiling stable programs from it. The syntax-semantics gap is not a limitation of current tools; it is a consequence of applying linear intuitions to nonlinear systems. Cells are not collections of independent variables. They are dynamical systems occupying positions on high-dimensional manifolds, governed by vector fields that evolution has spent billions of years optimizing. When we edit a genome, we perturb this system. The outcome depends not on the edit alone, but on the geometry of the landscape at the point of perturbation.

This geometry imposes a tax. Master regulators like CEBPA, coordinating dozens of downstream pathways, pay heavily: their perturbations produce large effects but scatter cells into incoherent states. Lineage-specific factors like KLF1, engaging tightly coordinated modules, pay minimally: their perturbations move cells together along defined trajectories. The Geometric Tax is not a metaphor. It is the measurable cost of working against regulatory architecture, quantifiable as the divergence between effect magnitude and directional coherence. Current evaluation frameworks, designed for the linear paradigm, cannot see this cost. They measure the shadow of the manifold, not the manifold itself.

Geometric stability fills this gap. It asks not just whether the edit occurred, but whether the resulting state is robust. It distinguishes laminar response from turbulent scatter, deep attractors from shallow plateaus, aligned perturbations from misaligned ones. Validation across CRISPR modalities demonstrates that this measure captures regulatory architecture invisible to sequence-level metrics. The CEBPA-KLF1 distinction emerges from the data without supervision, as does the correlation between instability and cellular stress. The geometry is real. It can be measured. And it predicts functional outcomes that magnitude alone cannot.

As biological engineering matures from sequence editing to state programming, evaluation must follow. Efficiency, specificity, and stability should become the three axes of assessment: efficiency to confirm the edit occurred; specificity to confirm it occurred only where intended; and stability to confirm the resulting state will persist. The first two axes are well established. The third is the geometric imperative.

We have learned to spell the genome. What remains is to learn its grammar. In the wasteland of unstable phenotypes, where raw edits bring only noise, the thunder finally speaks (Eliot 1922). Its command is simple: Control the geometry.

Code Availability

The full code necessary to reproduce all experiments, benchmarks, and analysis described in this paper is publicly available at `https://github.com/prashantcraju/geometric-stability-crispr`.

References

[1] Britt Adamson, Thomas M. Norman, Marco Jost, Min Y. Cho, James K. Nuñez, Yuwen Chen, Jacqueline E. Villalta, Luke A. Gilbert, Max A. Horlbeck, Marco Y. Hein, Ryan A. Pak, Andrew N. Gray, Carol A. Gross, Atray Dixit, Oren Parnas, Aviv Regev, and Jonathan S. Weissman. "A Multiplexed Single-Cell CRISPR Screening Platform Enables Systematic Dissection of the Unfolded Protein Response". In: *Cell* 167.7 (2016), 1867–1882.e21. ISSN: 0092-8674. DOI: 10.1016/j.cell.2016.11.048.

[2] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Francis Christiano, John Schulman, and Dandelion Mané. "Concrete Problems in AI Safety". In: *arXiv preprint arxiv:1606.06565* (2016).

[3] Jef D. Boeke, George Church, Andrew Hessel, Nancy J. Kelley, Adam Arkin, Yizhi Cai, Rob Carlson, Aravinda Chakravarti, Virginia W. Cornish, Liam Holt, Farren J. Isaacs, Todd Kuiken, Marc Lajoie, Tracy Lessor, Jeantine Lunshof, Matthew T. Maurano, Leslie A. Mitchell, Jasper Rine, Susan Rosser, Neville E. Sanjana, Pamela A. Silver, David Valle, Harris Wang, Jeffrey C. Way, and Luhan Yang. "The Genome Project-Write". In: *Science* 353.6295 (July 2016), pp. 126–127. ISSN: 1095-9203. DOI: 10.1126/science.aaf6850. URL: http://dx.doi.org/10.1126/science.aaf6850.

[4] Christopher A. Bravery, Jessica Carmen, Timothy Fong, Wanda Oprea, Karin H. Hoogendoorn, Juliana Woda, Scott R. Burger, Jon A. Rowley, Mark L. Bonyhadi, and Wouter Van't Hof. "Potency assay development for cellular therapy products: an ISCT review of the requirements and experiences in the industry". In: *Cytotherapy* 15.1 (Jan. 2013), 9–19.e9. ISSN: 1465-3249. DOI: 10.1016/j.jcyt.2012.10.008. URL: http://dx.doi.org/10.1016/j.jcyt.2012.10.008.

[5] Eva K. Brinkman, Tao Chen, Mario Amendola, and Bas van Steensel. "Easy quantitative assessment of genome editing by sequence trace decomposition". In: *Nucleic Acids Research* 42.22 (Oct. 2014), e168–e168. ISSN: 0305-1048. DOI: 10.1093/nar/gku936. URL: http://dx.doi.org/10.1093/nar/gku936.

[6] Chiara Capelli, Carolina Cuofano, Chiara Pavoni, Simona Frigerio, Daniela Lisini, Sara Nava, Michele Quaroni, Valentina Colombo, Francesco Galli, Svetlana Bezukladova, Paola Panina-Bordignon, Giuseppe Gaipa, Patrizia Comoli, Giulio Cossu, Gianvito Martino, Andrea Biondi, Martino Introna, and Josée Golay. "Potency assays and biomarkers for cell-based advanced therapy medicinal products". In: *Frontiers in Immunology* 14 (June 2023). ISSN: 1664-3224. DOI: 10.3389/fimmu.2023.1186224. URL: http://dx.doi.org/10.3389/fimmu.2023.1186224.

[7] Alessia Cavazza, Arianna Moiani, and Fulvio Mavilio. "Mechanisms of Retroviral Integration and Mutagenesis". In: *Human Gene Therapy* 24.2 (Feb. 2013), pp. 119–131. ISSN: 1557-7422. DOI: 10.1089/hum.2012.203. URL: http://dx.doi.org/10.1089/hum.2012.203.

[8] Dillon J. Cislo, M. Joaquina Delás, James Briscoe, and Eric D. Siggia. "Reconstructing Waddington's landscape from data". In: *Proceedings of the National Academy of Sciences* 122.49 (Dec. 2025). ISSN: 1091-6490. DOI: 10.1073/pnas.2521762122. URL: http://dx.doi.org/10.1073/pnas.2521762122.

[9] Ronald R. Coifman and Stéphane Lafon. "Diffusion maps". In: *Applied and Computational Harmonic Analysis* 21.1 (July 2006), pp. 5–30. ISSN: 1063-5203. DOI: 10.1016/j.acha.2006.04.006. URL: http://dx.doi.org/10.1016/j.acha.2006.04.006.

[10] Haotian Cui, Alejandro Tejada-Lapuerta, Maria Brbić, Julio Saez-Rodriguez, Simona Cristea, Hani Goodarzi, Mohammad Lotfollahi, Fabian J. Theis, and Bo Wang. "Towards multimodal foundation models in molecular cell biology". In: *Nature* 640.8059 (Apr. 2025), pp. 623–633. ISSN: 1476-4687. DOI: 10.1038/s41586-025-08710-y. URL: http://dx.doi.org/10.1038/s41586-025-08710-y.

[11] Haotian Cui, Chloe Wang, Hassaan Maan, Kuan Pang, Fengning Luo, Nan Duan, and Bo Wang. "scGPT: toward building a foundation model for single-cell multi-omics using generative AI". In: *Nature Methods* 21.8 (Feb. 2024), pp. 1470–1480. ISSN: 1548-7105. DOI: 10.1038/s41592-024-02201-0.

[12] Atray Dixit, Oren Parnas, Biyu Li, Jenny Chen, Charles P. Fulco, Livnat Jerby-Arnon, Nemanja D. Marjanovic, Danielle Dionne, Tyler Burks, Raktima Raychowdhury, Britt Adamson, Thomas M. Norman, Eric S. Lander, Jonathan S. Weissman, Nir Friedman, and Aviv Regev. "Perturb-Seq: dissecting molecular circuits with scalable single-cell RNA profiling of pooled genetic screens". In: *Cell* 167.7 (2016), 1853–1866.e17. ISSN: 0092-8674. DOI: 10.1016/j.cell.2016.11.038.

[13] Jennifer A. Doudna and Emmanuelle Charpentier. "The new frontier of genome engineering with CRISPR-Cas9". In: *Science* 346.6213 (Nov. 2014). ISSN: 1095-9203. DOI: 10.1126/science.1258096. URL: http://dx.doi.org/10.1126/science.1258096.

[14] T. S. Eliot. *The Waste Land*. New York: Boni and Liveright, 1922.

[15] Michael B. Elowitz, Arnold J. Levine, Eric D. Siggia, and Peter S. Swain. "Stochastic Gene Expression in a Single Cell". In: *Science* 297.5584 (Aug. 2002), pp. 1183–1186. ISSN: 1095-9203. DOI: 10.1126/science.1070919. URL: http://dx.doi.org/10.1126/science.1070919.

[16] Tariq Enver, Martin Pera, Carsten Peterson, and Peter W. Andrews. "Stem Cell States, Fates, and the Rules of Attraction". In: *Cell Stem Cell* 4.5 (May 2009), pp. 387–397. ISSN: 1934-5909. DOI: 10.1016/j.stem.2009.04.011. URL: http://dx.doi.org/10.1016/j.stem.2009.04.011.

[17] Atefeh Taherian Fard, Sriganesh Srihari, Jessica C Mar, and Mark A Ragan. "Not just a colourful metaphor: modelling the landscape of cellular development using Hopfield networks". In: *npj Systems Biology and Applications* 2.1 (Feb. 2016). ISSN: 2056-7189. DOI: 10.1038/npjsba.2016.1. URL: http://dx.doi.org/10.1038/npjsba.2016.1.

[18] James E. Ferrell. "Bistability, Bifurcations, and Waddington's Epigenetic Landscape". In: *Current Biology* 22.11 (June 2012), R458–R466. ISSN: 0960-9822. DOI: 10.1016/j.cub.2012.03.045. URL: http://dx.doi.org/10.1016/j.cub.2012.03.045.

[19] Ronald Aylmer Fisher. *The genetical theory of natural selection.* Clarendon Press, 1930. DOI: 10.5962/bhl.title.27468. URL: http://dx.doi.org/10.5962/bhl.title.27468.

[20] Joseph A. Fraietta, Simon F. Lacey, Elena J. Orlando, Iulian Pruteanu-Malinici, Mercy Gohil, Stefan Lundh, Alina C. Boesteanu, Yan Wang, Roddy S. O'Connor, Wei-Ting Hwang, Edward Pequignot, David E. Ambrose, Changfeng Zhang, Nicholas Wilcox, Felipe Bedoya, Corin Dorfmeier, Fang Chen, Lifeng Tian, Harit Parakandi, Minnal Gupta, Regina M. Young, F. Brad Johnson, Irina Kulikovskaya, Li Liu, Jun Xu, Sadik H. Kassim, Megan M. Davis, Bruce L. Levine, Noelle V. Frey, Donald L. Siegel, Alexander C. Huang, E. John Wherry, Hans Bitter, Jennifer L. Brogdon, David L. Porter, Carl H. June, and J. Joseph Melenhorst. "Determinants of response and resistance to CD19 chimeric antigen receptor (CAR) T cell therapy of chronic lymphocytic leukemia". In: *Nature Medicine* 24.5 (Apr. 2018), pp. 563–571. ISSN: 1546-170X. DOI: 10.1038/s41591-018-0010-1. URL: http://dx.doi.org/10.1038/s41591-018-0010-1.

[21] Alan D. Friedman. "C/ebpα induces pu.1 and interacts with ap-1 and nf-κb to regulate myeloid development". In: *Blood Cells, Molecules, and Diseases* 39.3 (2007), pp. 340–343. ISSN: 1079-9796. DOI: 10.1016/j.bcmd.2007.06.010.

[22] Alan D. Friedman. "C/EBPα in normal and malignant myelopoiesis". In: *International Journal of Hematology* 101.4 (2015), pp. 330–341. ISSN: 1865-3774. DOI: 10.1007/s12185-015-1764-6.

[23] Emma Haapaniemi, Sandeep Botla, Jenna Persson, Bernhard Schmierer, and Jussi Taipale. "CRISPR-Cas9 genome editing induces a p53-mediated DNA damage response". In: *Nature Medicine* 24.7 (June 2018), pp. 927–930. ISSN: 1546-170X. DOI: 10.1038/s41591-018-0049-z. URL: http://dx.doi.org/10.1038/s41591-018-0049-z.

[24] Yourui Han, Bolin Chen, Zhongwen Bi, Jianjun Zhang, Youpeng Hu, Jun Bian, Ruiming Kang, and Xuequn Shang. "Reconstructing Waddington Landscape from Cell Migration and Proliferation". In: *Interdisciplinary Sciences: Computational Life Sciences* 17.3 (Jan. 2025), pp. 541–554. ISSN: 1867-1462. DOI: 10.1007/s12539-024-00686-z. URL: http://dx.doi.org/10.1007/s12539-024-00686-z.

[25] Tetsuya Hori, Miho Amano, Aussie Suzuki, Chelsea B. Backer, Julie P. Welburn, Yimin Dong, Bruce F. McEwen, Wei-Hao Shang, Emiko Suzuki, Katsuya Okawa, Iain M. Cheeseman, and Tatsuo Fukagawa. "CCAN Makes Multiple Contacts with Centromeric DNA to Provide Distinct Pathways to the Outer Kinetochore". In: *Cell* 135.6 (Dec. 2008), pp. 1039–1052. ISSN: 0092-8674. DOI: 10.1016/j.cell.2008.10.019. URL: http://dx.doi.org/10.1016/j.cell.2008.10.019.

[26] Sui Huang. "Reprogramming cell fates: reconciling rarity with robustness". In: *BioEssays* 31.5 (Apr. 2009), pp. 546–560. ISSN: 1521-1878. DOI: 10.1002/bies.200800189. URL: http://dx.doi.org/10.1002/bies.200800189.

[27] Evan Hubinger, Chris van Merwijk, Vladimir Mikulik, Joar Skalse, and Scott Garrabrant. "Risks from Learned Optimization in Advanced Machine Learning Systems". In: *arXiv preprint arxiv:1906.01820* (2019).

[28] Robert J. Ihry, Kathleen A. Worringer, Max R. Salick, Elizabeth Frias, Daniel Ho, Kraig Theriault, Sravya Kommineni, Julie Chen, Marie Sondey, Chaoyang Ye, Ranjit Randhawa, Tripti Kulkarni, Zinger Yang, Gregory McAllister, Carsten Russ, John Reece-Hoyes, William Forrester, Gregory R. Hoffman, Ricardo Dolmetsch, and Ajamete Kaykas. "p53 inhibits CRISPR-Cas9 engineering in human pluripotent stem cells". In: *Nature Medicine* 24.7 (June 2018), pp. 939–946. ISSN: 1546-170X. DOI: 10.1038/s41591-018-0050-6. URL: http://dx.doi.org/10.1038/s41591-018-0050-6.

[29] Janus Schou Jakobsen, Johannes Waage, Nicolas Rapin, Hanne Cathrine Bisgaard, Fin Stolze Larsen, and Bo Torben Porse. "Temporal mapping of CEBPA and CEBPB binding during liver regeneration reveals dynamic occupancy and specific regulatory codes for homeostatic and cell cycle gene batteries". In: *Genome Research* 23.4 (Feb. 2013), pp. 592–603. ISSN: 1088-9051. DOI: 10.1101/gr.146399.112. URL: http://dx.doi.org/10.1101/gr.146399.112.

[30] Kristopher Torp Jensen, Lasse Fløe, Trine Skov Petersen, Jinrong Huang, Fengping Xu, Lars Bolund, Yonglun Luo, and Lin Lin. "Chromatin accessibility and guide sequence secondary structure affect CRISPR-Cas9 gene editing efficiency". In: *FEBS Letters* 591.13 (June 2017), pp. 1892–1901. ISSN: 1873-3468. DOI: 10.1002/1873-3468.12707. URL: http://dx.doi.org/10.1002/1873-3468.12707.

[31] Fuguo Jiang and Jennifer A. Doudna. "CRISPR-Cas9 Structures and Mechanisms". In: *Annual Review of Biophysics* 46.1 (May 2017), pp. 505–529. ISSN: 1936-1238. DOI: 10.1146/annurev-biophys-062215-010822. URL: http://dx.doi.org/10.1146/annurev-biophys-062215-010822.

[32] Martin Jinek, Krzysztof Chylinski, Ines Fonfara, Michael Hauer, Jennifer A. Doudna, and Emmanuelle Charpentier. "A Programmable Dual-RNA–Guided DNA Endonuclease in Adaptive Bacterial Immunity". In: *Science* 337.6096

(Aug. 2012), pp. 816–821. ISSN: 1095-9203. DOI: 10.1126/science.1225829. URL: http://dx.doi.org/10.1126/science.1225829.

[33] Marc Kirschner and John Gerhart. "Evolvability". In: *Proceedings of the National Academy of Sciences* 95.15 (July 1998), pp. 8420–8427. ISSN: 1091-6490. DOI: 10.1073/pnas.95.15.8420. URL: http://dx.doi.org/10.1073/pnas.95.15.8420.

[34] Hiroaki Kitano. "Biological robustness". In: *Nature Reviews Genetics* 5.11 (2004), pp. 826–837. ISSN: 1471-0064. DOI: 10.1038/nrg1471.

[35] Melinda Kliegman, Manar Zaghlula, Susan Abrahamson, Jonathan H. Esensten, Ross C. Wilson, Fyodor D. Urnov, and Jennifer A. Doudna. "A roadmap for affordable genetic medicines". In: *Nature* 634.8033 (July 2024), pp. 307–314. ISSN: 1476-4687. DOI: 10.1038/s41586-024-07800-7. URL: http://dx.doi.org/10.1038/s41586-024-07800-7.

[36] Michael Kosicki, Kärt Tomberg, and Allan Bradley. "Repair of double-strand breaks induced by CRISPR-Cas9 leads to large deletions and complex rearrangements". In: *Nature Biotechnology* 36.8 (July 2018), pp. 765–771. ISSN: 1546-1696. DOI: 10.1038/nbt.4192. URL: http://dx.doi.org/10.1038/nbt.4192.

[37] Victoria Krakovna, Jonathan Uesato, Vladimir Mikulik, Matthew Rahtz, Tom Everitt, Ramana Kumar, Zachary Kenton, Jan Leike, and Shane Legg. *Specification gaming: the flip side of AI ingenuity*. DeepMind Blog. Accessed: 2026-01-23. 2020. URL: https://deepmind.google/discover/blog/specification-gaming-the-flip-side-of-ai-ingenuity/.

[38] Mitchell L. Leibowitz, Stamatis Papathanasiou, Phillip A. Doerfler, Logan J. Blaine, Lili Sun, Yu Yao, Cheng-Zhong Zhang, Mitchell J. Weiss, and David Pellman. "Chromothripsis as an on-target consequence of CRISPR-Cas9 genome editing". In: *Nature Genetics* 53.6 (Apr. 2021), pp. 895–905. ISSN: 1546-1718. DOI: 10.1038/s41588-021-00838-7. URL: http://dx.doi.org/10.1038/s41588-021-00838-7.

[39] Bruce L. Levine, James Miskin, Keith Wonnacott, and Christopher Keir. "Global Manufacturing of CAR T Cell Therapy". In: *Molecular Therapy - Methods & Clinical Development* 4 (Mar. 2017), pp. 92–101. ISSN: 2329-0501. DOI: 10.1016/j.omtm.2016.12.006. URL: http://dx.doi.org/10.1016/j.omtm.2016.12.006.

[40] Yonatan Y Lipsitz, Nicholas E Timmins, and Peter W Zandstra. "Quality cell therapy manufacturing by design". In: *Nature Biotechnology* 34.4 (Apr. 2016), pp. 393–400. ISSN: 1546-1696. DOI: 10.1038/nbt.3525. URL: http://dx.doi.org/10.1038/nbt.3525.

[41] Wei Liu, Zhen Dou, Chunyue Wang, Gangyin Zhao, Fengge Wu, Chunli Wang, Felix Aikhionbare, Mingliang Ye, Divine Mensah Sedzro, Zhenye Yang, Chuanhai Fu, Zhikai Wang, Xinjiao Gao, Xuebiao Yao, Xiaoyu Song, and Xing Liu. "Aurora B promotes the CENP-T–CENP-W interaction to guide accurate chromosome segregation in mitosis". In: *Journal of Molecular Cell Biology* 16.2 (Jan. 2024). Ed. by Zhiyuan Shen. ISSN: 1759-4685. DOI: 10.1093/jmcb/mjae001. URL: http://dx.doi.org/10.1093/jmcb/mjae001.

[42] Elsa Logarinho, Tatiana Resende, Cláudia Torres, and Hassan Bousbaa. "The Human Spindle Assembly Checkpoint Protein Bub3 Is Required for the Establishment of Efficient Kinetochore–Microtubule Attachments". In: *Molecular Biology of the Cell* 19.4 (Apr. 2008). Ed. by Stephen Doxsey, pp. 1798–1813. ISSN: 1939-4586. DOI: 10.1091/mbc.e07-07-0633. URL: http://dx.doi.org/10.1091/mbc.E07-07-0633.

[43] Darío G. Lupiáñez, Katerina Kraft, Verena Heinrich, Peter Krawitz, Francesco Brancati, Eva Klopocki, Denise Horn, Hülya Kayserili, John M. Opitz, Renata Laxova, Fernando Santos-Simarro, Brigitte Gilbert-Dussardier, Lars Wittler, Marina Borschiwer, Stefan A. Haas, Marco Osterwalder, Martin Franke, Bernd Timmermann, Jochen Hecht, Malte Spielmann, Axel Visel, and Stefan Mundlos. "Disruptions of Topological Chromatin Domains Cause Pathogenic Rewiring of Gene-Enhancer Interactions". In: *Cell* 161.5 (May 2015), pp. 1012–1025. ISSN: 0092-8674. DOI: 10.1016/j.cell.2015.04.004. URL: http://dx.doi.org/10.1016/j.cell.2015.04.004.

[44] Ernst Mayr. *Animal Species and Evolution*. London, England: Harvard University Press, Dec. 1963.

[45] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger. "UMAP: Uniform Manifold Approximation and Projection". In: *Journal of Open Source Software* 3.29 (Sept. 2018), p. 861. ISSN: 2475-9066. DOI: 10.21105/joss.00861. URL: http://dx.doi.org/10.21105/joss.00861.

[46] I J Miller and J J Bieker. "A novel, erythroid cell-specific murine transcription factor that binds to the CACCC element and is related to the Krüppel family of nuclear proteins." In: *Molecular and Cellular Biology* 13.5 (1993), pp. 2776–2786. ISSN: 1098-5549. DOI: 10.1128/mcb.13.5.2776.

[47] Mitra Mojtahedi, Alexander Skupin, Joseph Zhou, Ivan G. Castaño, Rebecca Y. Y. Leong-Quong, Hannah Chang, Kalliopi Trachana, Alessandro Giuliani, and Sui Huang. "Cell Fate Decision as High-Dimensional Critical State Transition". In: *PLOS Biology* 14.12 (Dec. 2016), e2000640. ISSN: 1545-7885. DOI: 10.1371/journal.pbio.2000640. URL: http://dx.doi.org/10.1371/journal.pbio.2000640.

[48] Kevin R. Moon, David van Dijk, Zheng Wang, Scott Gigante, Daniel B. Burkhardt, William S. Chen, Kristina Yim, Antonia van den Elzen, Matthew J. Hirn, Ronald R. Coifman, Natalia B. Ivanova, Guy Wolf, and Smita Krishnaswamy. "Visualizing structure and transitions in high-dimensional biological data". In: *Nature Biotechnology* 37.12 (Dec. 2019), pp. 1482–1492. ISSN: 1546-1696. DOI: 10.1038/s41587-019-0336-3. URL: http://dx.doi.org/10.1038/s41587-019-0336-3.

[49] Gopika G. Nair, Jennifer S. Liu, Holger A. Russ, Stella Tran, Michael S. Saxton, Richard Chen, Charity Juang, Meilan Li, Vinh Q. Nguyen, Simone Giacometti, Sapna Puri, Yuan Xing, Yong Wang, Gregory L. Szot, Jose Oberholzer, Anil Bhushan, and Matthias Hebrok. "Recapitulating endocrine cell clustering in culture promotes maturation of human stem-cell-derived β cells". In: *Nature Cell Biology* 21.2 (Feb. 2019), pp. 263–274. ISSN: 1476-4679. DOI: 10.1038/s41556-018-0271-4. URL: http://dx.doi.org/10.1038/s41556-018-0271-4.

[50] Thomas M Norman, Max A Horlbeck, Joseph M Replogle, Y Alexander Ge, Alex Xu, Marco Jost, Luke A Gilbert, and Jonathan S Weissman. "Exploring genetic interaction manifolds constructed from rich single-cell phenotypes". In: *Science* 365.6455 (2019), pp. 786–793. ISSN: 1095-9203. DOI: 10.1126/science.aax4438.

[51] Luke O'Brien, Peter A. Hosick, Kezia John, David E. Stec, and Terry D. Hinds. "Biliverdin reductase isozymes in metabolism". In: *Trends in Endocrinology & Metabolism* 26.4 (Apr. 2015), pp. 212–220. ISSN: 1043-2760. DOI: 10.1016/j.tem.2015.02.001. URL: http://dx.doi.org/10.1016/j.tem.2015.02.001.

[52] S Oyadomari and M Mori. "Roles of CHOP/GADD153 in endoplasmic reticulum stress". In: *Cell Death & Differentiation* 11.4 (Dec. 2003), pp. 381–389. ISSN: 1476-5403. DOI: 10.1038/sj.cdd.4401373. URL: http://dx.doi.org/10.1038/sj.cdd.4401373.

[53] Efthymia Papalexi, Eleni P. Mimitou, Andrew W. Butler, Samantha Foster, Bernadette Bracken, William M. Mauck, Hans-Hermann Wessels, Yuhan Hao, Bertrand Z. Yeung, Peter Smibert, and Rahul Satija. "Characterizing the molecular regulation of inhibitory immune checkpoints with multimodal single-cell screens". In: *Nature Genetics* 53.3 (2021), pp. 322–331. ISSN: 1546-1718. DOI: 10.1038/s41588-021-00778-2.

[54] Natasia Paukovich, Mengjun Xue, James R. Elder, Jasmina S. Redzic, Ashley Blue, Hamish Pike, Brian G. Miller, Todd M. Pitts, David D. Pollock, Kirk Hansen, Angelo D'Alessandro, and Elan Zohar Eisenmesser. "Biliverdin Reductase B Dynamics Are Coupled to Coenzyme Binding". In: *Journal of Molecular Biology* 430.18 (Sept. 2018), pp. 3234–3250. ISSN: 0022-2836. DOI: 10.1016/j.jmb.2018.06.015. URL: http://dx.doi.org/10.1016/j.jmb.2018.06.015.

[55] Mary Philip, Lauren Fairchild, Liping Sun, Ellen L. Horste, Steven Camara, Mojdeh Shakiba, Andrew C. Scott, Agnes Viale, Peter Lauer, Taha Merghoub, Matthew D. Hellmann, Jedd D. Wolchok, Christina S. Leslie, and Andrea Schietinger. "Chromatin states define tumour-specific T cell dysfunction and reprogramming". In: *Nature* 545.7655 (May 2017), pp. 452–456. ISSN: 1476-4687. DOI: 10.1038/nature22367. URL: http://dx.doi.org/10.1038/nature22367.

[56] Andre M. Pilon, Murat O. Arcasoy, Holly K. Dressman, Serena E. Vayda, Yelena D. Maksimova, Jose I. Sangerman, Patrick G. Gallagher, and David M. Bodine. "Failure of Terminal Erythroid Differentiation in EKLF-Deficient Mice Is Associated with Cell Cycle Perturbation and Reduced Expression of E2F2". In: *Molecular and Cellular Biology* 28.24 (2008), pp. 7394–7401. ISSN: 1098-5549. DOI: 10.1128/mcb.01087-08.

[57] Archishman Raju and Eric D. Siggia. "A geometrical perspective on development". In: *Development, Growth & Differentiation* 65.5 (May 2023), pp. 245–254. ISSN: 1440-169X. DOI: 10.1111/dgd.12855. URL: http://dx.doi.org/10.1111/dgd.12855.

[58] Archishman Raju, BingKan Xue, and Stanislas Leibler. "A theoretical perspective on Waddington's genetic assimilation experiments". In: *Proceedings of the National Academy of Sciences* 120.51 (Dec. 2023). ISSN: 1091-6490. DOI: 10.1073/pnas.2309760120. URL: http://dx.doi.org/10.1073/pnas.2309760120.

[59] Prashant C. Raju. "Geometric Stability: The Missing Axis of Representations". In: *arXiv preprint arXiv:2601.09173* (2026).

[60] Prashant C. Raju. *Shesha: Self-Consistency Metrics for Representational Stability*. Zenodo. 2026. DOI: 10.5281/zenodo.18227453. URL: https://doi.org/10.5281/zenodo.18227453.

[61] David A. Rand, Archishman Raju, Meritxell Sáez, Francis Corson, and Eric D. Siggia. "Geometry of gene regulatory dynamics". In: *Proceedings of the National Academy of Sciences* 118.38 (Sept. 2021). ISSN: 1091-6490. DOI: 10.1073/pnas.2109729118. URL: http://dx.doi.org/10.1073/pnas.2109729118.

[62] Charles Rathkopf. "Hallucination, reliability, and the role of generative AI in science". In: *arXiv preprint arxiv:2504.08526* (2025).

[63] Joseph M. Replogle, Reuben A. Saunders, Angela N. Pogson, Jeffrey A. Hussmann, Alexander Lenail, Alina Guna, Lauren Mascibroda, Eric J. Wagner, Karen Adelman, Gila Lithwick-Yanai, Nika Iremadze, Florian Oberstrass, Doron Lipson, Jessica L. Bonnar, Marco Jost, Thomas M. Norman, and Jonathan S. Weissman. "Mapping information-rich

genotype-phenotype landscapes with genome-scale Perturb-seq". In: *Cell* 185.14 (July 2022), 2559–2575.e28. ISSN: 0092-8674. DOI: 10.1016/j.cell.2022.05.013.

[64] Yanay Rosen, Yusuf Roohani, Ayush Agrawal, Leon Samotorcan, Tabula Sapiens Consortium, Stephen R. Quake, and Jure Leskovec. "Universal Cell Embeddings: A Foundation Model for Cell Biology". In: *bioRxiv* (2023). DOI: 10.1101/2023.11.28.568918.

[65] Stuart Russell. *Human Compatible: Artificial Intelligence and the Problem of Control*. Viking, Oct. 2019. ISBN: 978-0-525-55861-3.

[66] Paula Salmikangas, Margarida Menezes-Ferreira, Ilona Reischl, Asterios Tsiftsoglou, Jan Kyselovic, John Joseph Borg, Sol Ruiz, Egbert Flory, Jean-Hugues Trouvin, Patrick Celis, Janis Ancans, Marcos Timon, Guido Pante, Dariusz Sladowski, Metoda Lipnik-Stangelj, and Christian K Schneider. "Manufacturing, Characterization and Control of Cell-Based Medicinal Products: Challenging Paradigms Toward Commercial Use". In: *Regenerative Medicine* 10.1 (Jan. 2015), pp. 65–78. ISSN: 1746-076X. DOI: 10.2217/rme.14.65. URL: http://dx.doi.org/10.2217/rme.14.65.

[67] Marten Scheffer, Jordi Bascompte, William A. Brock, Victor Brovkin, Stephen R. Carpenter, Vasilis Dakos, Hermann Held, Egbert H. van Nes, Max Rietkerk, and George Sugihara. "Early-warning signals for critical transitions". In: *Nature* 461.7260 (Sept. 2009), pp. 53–59. ISSN: 1476-4687. DOI: 10.1038/nature08227. URL: http://dx.doi.org/10.1038/nature08227.

[68] Debattama R. Sen, James Kaminski, R. Anthony Barnitz, Makoto Kurachi, Ulrike Gerdemann, Kathleen B. Yates, Hsiao-Wei Tsao, Jernej Godec, Martin W. LaFleur, Flavian D. Brown, Pierre Tonnerre, Raymond T. Chung, Damien C. Tully, Todd M. Allen, Nicole Frahm, Georg M. Lauer, E. John Wherry, Nir Yosef, and W. Nicholas Haining. "The epigenetic landscape of T cell exhaustion". In: *Science* 354.6316 (Dec. 2016), pp. 1165–1169. ISSN: 1095-9203. DOI: 10.1126/science.aae0491. URL: http://dx.doi.org/10.1126/science.aae0491.

[69] Monica F. Sentmanat, Samuel T. Peters, Colin P. Florian, Jon P. Connelly, and Shondra M. Pruett-Miller. "A Survey of Validation Strategies for CRISPR-Cas9 Editing". In: *Scientific Reports* 8.1 (Jan. 2018). ISSN: 2045-2322. DOI: 10.1038/s41598-018-19441-8. URL: http://dx.doi.org/10.1038/s41598-018-19441-8.

[70] Miroslawa Siatecka and James J. Bieker. "The multifunctional role of EKLF/KLF1 during erythropoiesis". In: *Blood* 118.8 (2011), pp. 2044–2054. ISSN: 1528-0020. DOI: 10.1182/blood-2011-03-331371.

[71] Mark L. Siegal and Aviv Bergman. "Waddington's canalization revisited: Developmental stability and evolution". In: *Proceedings of the National Academy of Sciences* 99.16 (June 2002), pp. 10528–10532. ISSN: 1091-6490. DOI: 10.1073/pnas.102303999. URL: http://dx.doi.org/10.1073/pnas.102303999.

[72] Jonathan M. W. Slack. "Conrad Hal Waddington: the last Renaissance biologist?" In: *Nature Reviews Genetics* 3.11 (Nov. 2002), pp. 889–895. ISSN: 1471-0064. DOI: 10.1038/nrg933. URL: http://dx.doi.org/10.1038/nrg933.

[73] Chutima Talchai, Shouhong Xuan, Hua V. Lin, Lori Sussel, and Domenico Accili. "Pancreatic β Cell Dedifferentiation as a Mechanism of Diabetic β Cell Failure". In: *Cell* 150.6 (Sept. 2012), pp. 1223–1234. ISSN: 0092-8674. DOI: 10.1016/j.cell.2012.07.029. URL: http://dx.doi.org/10.1016/j.cell.2012.07.029.

[74] Michael R. Tallack and Andrew C. Perkins. "KLF1 directly coordinates almost all aspects of terminal erythroid differentiation". In: *IUBMB Life* 62.12 (2010), pp. 886–890. ISSN: 1521-6551. DOI: 10.1002/iub.404.

[75] Michael R. Tallack, Tom Whitington, Wai Shan Yuen, Elanor N. Wainwright, Janelle R. Keys, Brooke B. Gardiner, Ehsan Nourbakhsh, Nicole Cloonan, Sean M. Grimmond, Timothy L. Bailey, and Andrew C. Perkins. "A global role for KLF1 in erythropoiesis revealed by ChIP-seq in primary erythroid cells". In: *Genome Research* 20.8 (2010), pp. 1052–1063. ISSN: 1088-9051. DOI: 10.1101/gr.106575.110.

[76] Christina V. Theodoris, Ling Xiao, Anant Chopra, Mark D. Chaffin, Zeina R. Al Sayed, Matthew C. Hill, Helene Mantineo, Elizabeth M. Brydon, Zexian Zeng, X. Shirley Liu, and Patrick T. Ellinor. "Transfer learning enables predictions in network biology". In: *Nature* 618.7965 (May 2023), pp. 616–624. ISSN: 1476-4687. DOI: 10.1038/s41586-023-06139-9.

[77] Shengdar Q Tsai, Zongli Zheng, Nhu T Nguyen, Matthew Liebers, Ved V Topkar, Vishal Thapar, Nicolas Wyvekens, Cyd Khayter, A John Iafrate, Long P Le, Martin J Aryee, and J Keith Joung. "GUIDE-seq enables genome-wide profiling of off-target cleavage by CRISPR-Cas9 nucleases". In: *Nature Biotechnology* 33.2 (Dec. 2014), pp. 187–197. ISSN: 1546-1696. DOI: 10.1038/nbt.3117. URL: http://dx.doi.org/10.1038/nbt.3117.

[78] Koki Tsuyuzaki, Hiroyuki Sato, Kenta Sato, and Itoshi Nikaido. "Benchmarking principal component analysis for large-scale single-cell RNA-sequencing". In: *Genome Biology* 21.1 (Jan. 2020). ISSN: 1474-760X. DOI: 10.1186/s13059-019-1900-3. URL: http://dx.doi.org/10.1186/s13059-019-1900-3.

[79] Josh Tycko, Michael Wainberg, Georgi K. Marinov, Oana Ursu, Gaelen T. Hess, Braeden K. Ego, Aradhana, Amy Li, Alisa Truong, Alexandro E. Trevino, Kaitlyn Spees, David Yao, Irene M. Kaplow, Peyton G. Greenside, David W. Morgens, Douglas H. Phanstiel, Michael P. Snyder, Lacramioara Bintu, William J. Greenleaf, Anshul Kundaje, and

Michael C. Bassik. “Mitigation of off-target toxicity in CRISPR-Cas9 screens for essential non-coding elements”. In: *Nature Communications* 10.1 (Sept. 2019). ISSN: 2041-1723. DOI: 10.1038/s41467-019-11955-7. URL: http://dx.doi.org/10.1038/s41467-019-11955-7.

[80] U.S. Food and Drug Administration. *Guidance for Industry: Preclinical Assessment of Investigational Cellular and Gene Therapy Products*. 2013. URL: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/preclinical-assessment-investigational-cellular-and-gene-therapy-products.

[81] U.S. Food and Drug Administration. *Human Gene Therapy Products Incorporating Human Genome Editing; Draft Guidance for Industry*. 2024. URL: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/human-gene-therapy-products-incorporating-human-genome-editing.

[82] Isaac Virshup, Danila Bredikhin, Lukas Heumos, Giovanni Palla, Gregor Sturm, Adam Gayoso, Ilia Kats, Mikaela Koutrouli, Philipp Angerer, Volker Bergen, Pierre Boyeau, Maren Büttner, Gokcen Eraslan, David Fischer, Max Frank, Justin Hong, Michal Klein, Marius Lange, Romain Lopez, Mohammad Lotfollahi, Malte D. Luecken, Fidel Ramirez, Jeffrey Regier, Sergei Rybakov, Anna C. Schaar, Valeh Valiollah Pour Amiri, Philipp Weiler, Galen Xing, Bonnie Berger, Dana Pe’er, Aviv Regev, Sarah A. Teichmann, Francesca Finotello, F. Alexander Wolf, Nir Yosef, Oliver Stegle, and Fabian J. Theis. “The scverse project provides a computational ecosystem for single-cell omics data analysis”. In: *Nature Biotechnology* 41.5 (Apr. 2023), pp. 604–606. ISSN: 1546-1696. DOI: 10.1038/s41587-023-01733-8. URL: http://dx.doi.org/10.1038/s41587-023-01733-8.

[83] Isaac Virshup, Sergei Rybakov, Fabian J. Theis, Philipp Angerer, and F. Alexander Wolf. “anndata: Access and store annotated data matrices”. In: *Journal of Open Source Software* 9.101 (Sept. 2024), p. 4371. ISSN: 2475-9066. DOI: 10.21105/joss.04371. URL: http://dx.doi.org/10.21105/joss.04371.

[84] Conrad Hal Waddington. *The Strategy of the Genes: A Discussion of Some Aspects of Theoretical Biology*. London: George Allen & Unwin, 1957.

[85] Allon Wagner, Aviv Regev, and Nir Yosef. “Revealing the vectors of cellular identity with single-cell genomics”. In: *Nature Biotechnology* 34.11 (Nov. 2016), pp. 1145–1160. ISSN: 1546-1696. DOI: 10.1038/nbt.3711. URL: http://dx.doi.org/10.1038/nbt.3711.

[86] Jin Wang, Kun Zhang, Li Xu, and Erkang Wang. “Quantifying the Waddington landscape and biological paths for development and differentiation”. In: *Proceedings of the National Academy of Sciences* 108.20 (May 2011), pp. 8257–8262. ISSN: 1091-6490. DOI: 10.1073/pnas.1017017108. URL: http://dx.doi.org/10.1073/pnas.1017017108.

[87] Caleb Weinreb, Alejo Rodriguez-Fraticelli, Fernando D. Camargo, and Allon M. Klein. “Lineage tracing on transcriptional landscapes links state to fate during differentiation”. In: *Science* 367.6479 (Feb. 2020). ISSN: 1095-9203. DOI: 10.1126/science.aaw3381. URL: http://dx.doi.org/10.1126/science.aaw3381.

[88] Peng Zhang, Qian Yang, Xulong Chen, Xiaolong Chen, Qing Wang, Kun Chen, Yu An, Kehua Jiang, and Fa Sun. “CENPW knockdown inhibits progression of bladder cancer through inducing cell cycle arrest and apoptosis”. In: *Journal of Cancer* 15.3 (2024), pp. 858–870. ISSN: 1837-9664. DOI: 10.7150/jca.90449. URL: http://dx.doi.org/10.7150/jca.90449.

[89] Yuansheng Zhou and Tatyana O. Sharpee. “Hyperbolic geometry of gene expression”. In: *iScience* 24.3 (Mar. 2021), p. 102225. ISSN: 2589-0042. DOI: 10.1016/j.isci.2021.102225. URL: http://dx.doi.org/10.1016/j.isci.2021.102225.

[90] H. Zinszner, M. Kuroda, X. Wang, N. Batchvarova, R. T. Lightfoot, H. Remotti, J. L. Stevens, and D. Ron. “CHOP is implicated in programmed cell death in response to impaired function of the endoplasmic reticulum”. In: *Genes & Development* 12.7 (Apr. 1998), pp. 982–995. ISSN: 0890-9369. DOI: 10.1101/gad.12.7.982. URL: http://dx.doi.org/10.1101/gad.12.7.982.