Attractors are Less Stable than their Basins: Canalization Creates a Coherence Gap in Gene Regulatory Networks

Venkata Sai Narayana Bavisetty [*] and Matthew Wheeler [†]
Department of Medicine, University of Florida, Gainesville, Florida 32610, USA

Claus Kadelka [‡]
Department of Mathematics, Iowa State University, Ames, Iowa 50011, USA

(Received 29 November 2025; accepted 20 April 2026; published 26 May 2026)

Waddington's epigenetic landscape has served as biology's central metaphor for cellular differentiation for over half a century, depicting mature cell types as balls resting in stable valley floors. Boolean networks—introduced by Kauffman in 1969 to model gene regulatory dynamics—provide a mathematical formalization of this landscape, where attractors represent phenotypes, and basins of attraction correspond to developmental valleys. Traditional stability measures quantify robustness by perturbing arbitrary states, yet biological systems typically reside at attractors rather than in transient states. Here we formalize and systematically analyze attractor coherence, a stability measure Kauffman originally envisioned but never rigorously developed, which quantifies how likely it is that a perturbation of an attractor state causes phenotype switching. Analyzing 95 expert-curated biological Boolean models, we reveal a striking paradox: attractors representing mature cell types are consistently less stable than the developmental trajectories approaching them. Large-scale simulations of random networks demonstrate that this coherence gap arises from canalization—a hallmark of biological regulation where individual genes can override other genes. While canalization increases overall network stability, it disproportionately stabilizes transient states, positioning attractors near basin boundaries. The gap's magnitude is almost perfectly predicted by network bias (Spearman's $\rho = -0.997$), itself modulated by canalization. These findings revise Waddington's landscape: canalization carves deep protective valleys ensuring developmental robustness, yet simultaneously flattens ridges near valley floors, facilitating phenotypic plasticity when multiple fates coexist. This explains how biological systems achieve both reliable development and plasticity, with implications for understanding development, disease-related transitions, and designing robust yet controllable synthetic gene circuits.

DOI: 10.1103/v7kd-n9xh

I. INTRODUCTION

Waddington's epigenetic landscape metaphor has profoundly shaped how biologists conceptualize cellular differentiation and stability [1–3]. In this metaphor, a ball rolling down a landscape represents a cell progressing through development, with valleys corresponding to different cell fates or phenotypes. The depth of these valleys represents stability: once a cell reaches the valley floor—a mature, differentiated state—it should be difficult to perturb into a different phenotype. This intuitive picture suggests that mature cell types occupy the most stable positions in the regulatory landscape, protected by steep barriers from phenotype-switching perturbations.

Boolean networks (BNs) provide a natural mathematical formalization of Waddington's landscape [4]. In this framework, genes and proteins are represented by binary nodes, regulatory interactions by edges, and the dynamics evolve according to discrete update rules. The phenotypes or cell types correspond to *attractors*: fixed points or limit cycles to which the system eventually converges. The basin of attraction surrounding each attractor comprises all states that eventually transition to it, directly analogous to valleys in Waddington's landscape. This formalization has made Boolean networks a powerful tool for studying gene regulatory dynamics across diverse biological processes, from developmental programs to cancer signaling pathways [5].

The stability of Boolean network attractors has been a central question since Kauffman's foundational work on metabolic stability and epigenesis [4]. Early metrics like the Derrida value [6] characterized dynamical regimes—ordered, chaotic, or critical—by measuring sensitivity to perturbations after a single update. These analyses led to the criticality hypothesis: that biological systems operate near the critical regime to optimize the tradeoff between robustness and adaptability [7–10]. However, the Derrida value fails to capture the influence of network topology and feedback loops

[*]Contact author: sbavisetty@ucla.edu
[†]Contact author: matthew.wheeler@medicine.ufl.edu
[‡]Contact author: ckadelka@iastate.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

[11], depending only on local sensitivities rather than global structure [12,13].

To address these limitations, *coherence* was introduced as a measure of phenotypic robustness [14–17]. Coherence quantifies the probability that a single-node perturbation does not cause a transition to a different basin of attraction. Importantly, coherence is computed as an average over all states in the state space—both transient states along developmental trajectories and attractor states representing mature phenotypes. Yet biological systems predominantly reside at attractors, not in transient states: A cell spends the vast majority of its lifetime in a differentiated state, not in the process of differentiating. This raises a fundamental question: Are attractors themselves as stable as the basins surrounding them?

Kauffman himself recognized this distinction in his original 1969 paper, informally introducing the concept of attractor-specific stability [4]. However, this concept was never formalized or systematically analyzed. Here, we formalize *attractor coherence*, which is the stability of attractor states to perturbations. Attractor coherence is similar to *basin coherence* [14], which measures the average stability of all states leading to an attractor. By comparing these two measures across 95 expert-curated biological Boolean network models and large ensembles of random networks, we reveal a striking and counterintuitive pattern: attractors are often less stable than the developmental trajectories that lead to them.

Large-scale simulations of random networks demonstrate that this coherence gap is not incidental but arises from fundamental design principles of biological regulation. Canalization—the property whereby individual regulatory inputs can override others—is a hallmark of biological gene regulatory networks [8,10,18,19]. We show that while canalization increases overall network stability, it increases basin coherence more strongly than attractor coherence, creating a stability gradient within basins. Attractors become positioned closer to basin boundaries, making them more vulnerable to phenotype-switching perturbations. The magnitude of this coherence gap is almost perfectly predicted by network bias, which is itself strongly modulated by canalization [12,20].

These findings revise Waddington's landscape metaphor (Fig. 1). While canalization creates deep valleys that robustly guide cells along developmental trajectories—explaining why development is highly reproducible—it also creates a stability gradient within valleys. Attractors reside in regions with relatively flatter boundaries compared to mid-trajectory states, facilitating phenotype switching when multiple fates are accessible. This architectural feature is most biologically relevant when multiple phenotypically meaningful attractors coexist. It helps explain how biological systems achieve both developmental robustness (high absolute coherence protects trajectories) and phenotypic plasticity (coherence gap enables regulated transitions), with important implications for understanding processes like stem cell reprogramming [21], wound healing, and disease-related transitions such as cancer differentiation [2,22,23].

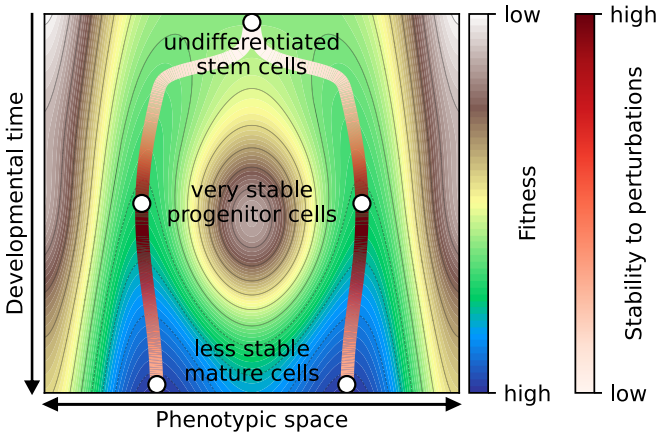


FIG. 1. Canalization creates a coherence gap in Waddington's epigenetic landscape. While canalization increases both basin and attractor coherence in absolute terms, it increases basin coherence more, creating a stability gradient within each valley. Attractors (representing mature cell types or phenotypes) are positioned closer to basin boundaries relative to transient states in their developmental trajectories. The ridges between valleys become relatively flatter near attractors compared to mid-trajectory, making phenotype switching more likely from mature states when multiple phenotypes are accessible. This coherence gap is almost perfectly predicted by network bias, which is itself modulated by canalization. Here, "fitness" is schematic and follows the classical developmental biology metaphor rather than representing a formally defined potential function. Stability to perturbations is quantified independently through coherence.

II. METHODS

A. Boolean networks

A Boolean network (BN) is a finite dynamical system, given by a function

$$F : \{0,1\}^N \longrightarrow \{0,1\}^N, \tag{1}$$

$$x \longmapsto (f_1(x), \ldots, f_n(x)). \tag{2}$$

It consists of N nodes, where each node can take on only two values, 0 (e.g., OFF, unexpressed, absent) and 1 (e.g., ON, expressed, present). The state of the entire BN at time t is represented by the vector $\mathbf{x}(t) = (x_1(t), \ldots, x_N(t)) \in \{0,1\}^N$. The *state space* $\{0,1\}^N$ is the set of all possible states and coincides with the vertices of the N-dimensional Boolean hypercube. Every node is updated in discrete time steps. The state of node i at time $t+1$ is determined by its *update function*

$$f_i : \{0,1\}^N \to \{0,1\}, \tag{3}$$

$$\mathbf{x}(t) \mapsto x_i(t+1). \tag{4}$$

In a given BN, f_i may depend only on a subset of nodes. The *in-degree* of node i is the number of essential inputs of f_i.

In this paper, we focus on synchronously updated BNs, where all nodes are updated simultaneously, yielding

deterministic dynamics:

$$\mathbf{x}(t+1) = F(\mathbf{x}(t)) = (f_1(\mathbf{x}(t)), \ldots, f_N(\mathbf{x}(t))). \quad (5)$$

The dynamics are fully described by the *state transition graph*, whose vertices are the 2^N possible states $\mathbf{x} = (x_1, \ldots, x_N) \in \{0, 1\}^N$ and there is an edge from state $\mathbf{x}$ to state $\mathbf{y}$ if $F(\mathbf{x}) = \mathbf{y}$. Because BNs have finite size and deterministic dynamics, they eventually exhibit periodic behavior. That is, they settle into *attractors*, which can be fixed points (also known as steady states) or limit cycles. In biological contexts, attractors represent phenotypes or cell types.

Given a BN F, let $\mathcal{A}(F)$ denote the set of all its attractors. The *basin of attraction* $B_F(A)$ of a given attractor $A \in \mathcal{A}(F)$ consists of all states that eventually transition to A, including the states within A itself. Consider the function

$$D_F : \{0, 1\}^N \longrightarrow \mathcal{A}(F), \quad (6)$$

which assigns each state to its attractor. With this notation, we have

$$B_F(A) = \{\mathbf{x} \in \{0, 1\}^N \mid D_F(\mathbf{x}) = A\}, \quad (7)$$

where the sum of all basin sizes $|B_F(A)|$ is

$$\sum_{A \in \mathcal{A}(F)} |B_F(A)| = 2^N. \quad (8)$$

The relative basin sizes are simply the basin sizes divided by the number of total states, 2^N.

B. Coherence

Coherence quantifies the probability that a single-node perturbation does not result in a change of phenotype [14]. Given a state $\mathbf{x}$ of a BN F, let $\mathbf{x} \oplus e_i$ denote the state where the ith bit in $\mathbf{x}$ is flipped. The *coherence* of state $\mathbf{x}$, denoted $\psi_{\mathbf{x}}$, is the fraction of neighboring states $\mathbf{x} \oplus e_1, \ldots, \mathbf{x} \oplus e_N$ that are in the same basin of attraction as $\mathbf{x}$. That is,

$$\psi_{\mathbf{x}} = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}[D_F(\mathbf{x}) = D_F(\mathbf{x} \oplus e_i)], \quad (9)$$

where $\mathbb{1}[\cdot]$ is the indicator function. The *basin coherence* ψ_B is defined as the average coherence of the states within a basin of attraction B, and is given by

$$\psi_B = \frac{1}{|B|} \sum_{\mathbf{x} \in B} \psi_{\mathbf{x}}. \quad (10)$$

Similarly, the network coherence ψ_F is given by

$$\psi_F = \frac{1}{2^N} \sum_{\mathbf{x} \in \{0,1\}^N} \psi_{\mathbf{x}}, \quad (11)$$

which is the average coherence of all states in the network. Network coherence can alternatively be obtained by simply summing all basin coherences, weighted by the relative basin sizes.

Typically, larger basins tend to have higher basin coherence. In particular, in [24] we established the following bounds on the expected basin coherence for a given basin size:

$$\frac{|B|}{2^N} \leqslant \mathbb{E}[\psi_B] \leqslant \frac{\log_2(|B|)}{N} + O\left(\frac{1}{N}\right). \quad (12)$$

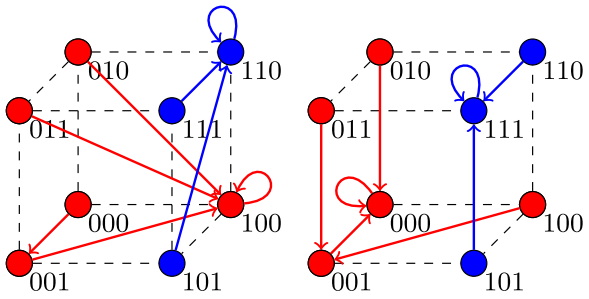


FIG. 2. State transition graph of two three-node Boolean networks. Despite identical basin sizes and basin coherence, both attractors in the left network F have coherence 1/3, whereas the attractors 000 and 111 in the right network G have coherence 1 and 2/3, respectively. This demonstrates that attractor coherence is not determined solely by basin properties.

The lower bound shows that basin coherence grows at least proportionally to the relative basin size, while the upper bound shows that it grows at most logarithmically in the basin size (as a fraction of N). The upper bound is asymptotic and the error term is typically negligible for $N > 12$.

C. Attractor coherence

Conventional stability metrics for BNs (e.g., the coherence, fragility, or Derrida values) describe the effect of perturbing a random network state (see, e.g., [16]). Since most BNs possess many more transient states than attractor states, these stability measures primarily reflect the resilience of transient states to perturbations rather than the resilience of the attractors. However, an unperturbed biological system rests at an attractor state. Therefore, to accurately describe the resilience of real-world systems, it is necessary to use a measure that quantifies the stability of attractors themselves.

Motivated by this, we study *attractor coherence*, a measure closely related to the previously defined basin coherence. It was already informally introduced by Kauffman in 1969 [4] (also see [25]). Attractor coherence describes the average coherence of states within an attractor, measuring the probability that a single-node perturbation of an attractor state does not push the system into another attractor. Analogous to basin coherence, the coherence of an attractor $A \in \mathcal{A}(F)$ can be defined by

$$\psi_A = \frac{1}{|A|} \sum_{\mathbf{x} \in A} \psi_{\mathbf{x}}, \quad (13)$$

where $|A|$ denotes the periodicity of the attractor.

As an example, consider the three-node BNs

$$F(x_1, x_2, x_3) = (x_1 \vee x_2 \vee x_3, x_1 \wedge (x_2 \vee x_3), \bar{x}_1 \wedge \bar{x}_2 \wedge \bar{x}_3)$$

and

$$G(x_1, x_2, x_3) = (x_1 \wedge (x_2 \vee x_3), x_1 \wedge (x_2 \vee x_3), x_1 \vee (x_2 \wedge x_3)),$$

whose state transition graphs are shown in Fig. 2. Both networks have two steady-state attractors with same basins of sizes 5 and 3. The basin coherences are $2/3$ and $4/9$, respectively, which represent the maximum possible coherence for basins of these sizes [24]. Thus, the total network coherence

for both F and G is $5/8 \times 2/3 + 3/8 \times 4/9 = 7/12$. However, the attractor coherences between the two networks differ strongly. F has steady-state attractors 100 and 110, which both have coherence 1/3 because only one of each of their three neighboring states is in their basin of attraction. On the contrary, the steady-state attractors 000 and 111 of G have much higher attractor coherence of 1 and 2/3, respectively.

Our analysis in this paper shows that most biological networks behave like F: the attractor coherence tends to be lower than the basin coherence, which has important implications for our understanding of network robustness.

To derive an expected value for attractor coherence and understand its theoretical baseline, we consider the random map model [26]. In this model, we consider a Boolean network defined by a random map,

$$F_{\text{rand}} : \{0,1\}^N \to \{0,1\}^N, \tag{14}$$

where the successor state $\mathbf{x}'$ of any state $\mathbf{x}$ is assigned uniformly at random. In this case, as demonstrated in [14], the coherence of a state depends solely on the basin size and follows a scaled binomial distribution $B(\frac{|b|-1}{2^N-1}, N)/N$. This implies that the expected basin coherence is

$$\mathbb{E}[\psi_B] = \frac{|B|-1}{2^N-1} \sim \frac{|B|}{2^N}. \tag{15}$$

Furthermore, they show that the coherence of a state in a basin follows a scaled binomial distribution $\text{Binom}(\frac{|b|-1}{2^N-1}, N)/N$. Similarly, for an attractor A of length l, we get that attractor coherence follows a scaled binomial $\text{Binom}(\frac{|b|-1}{2^N-1}, Nl)/Nl$. Crucially, in this random map model, attractor coherence and basin coherence have the same expected value: both equal $\frac{|b|-1}{2^N-1} \sim \frac{|b|}{2^N}$, which is simply the relative basin size and which does not depend on the length of the attractor.

However, the basin and attractor coherence of biological networks deviate substantially from this theoretical baseline because successor states are *not* assigned uniformly at random.

D. Coherence gap

Our main goal is to understand how the average basin coherence ψ_B, the attractor coherence ψ_A, and the difference $\psi_B - \psi_A$ vary between network ensembles. The expected values depend on the distribution $\phi(s)$ of relative basin sizes. Specifically, we have

$$\begin{aligned}\mathbb{E}[\psi_B] &= \int_0^1 \mathbb{E}[\psi_B \mid |B| = s]\,\phi(s)\,ds,\\ \mathbb{E}[\psi_A] &= \int_0^1 \mathbb{E}[\psi_A \mid |B| = s]\,\phi(s)\,ds,\end{aligned} \tag{16}$$

which shows that the overall expectations are weighted by the frequency of basin sizes, which tends to be skewed towards both very small and very large basins.

Different ensembles of BNs may possess different basin size distributions. For example, it is known that the number of attractors increases as the sensitivity of a network to perturbations increases [27]. To enable a comparison of ψ_B, ψ_A, and $\psi_B - \psi_A$ across network ensembles, we need to remove the confounding effect of basin size distribution. We achieve this by computing a standardized measure that considers all basin sizes equally likely in Eq. (16):

$$\begin{aligned}\text{AUC:BC} &= \int_0^1 \mathbb{E}[\psi_B \mid |B| = s]\,ds,\\ \text{AUC:AC} &= \int_0^1 \mathbb{E}[\psi_A \mid |B| = s]\,ds.\end{aligned} \tag{17}$$

AUC:BC and AUC:AC denote the area under the curves of the functions $\mathbb{E}[\psi_B \mid |B| = b]$ and $\mathbb{E}[\psi_A \mid |B(A)| = b]$, respectively. To compare the basin and the attractor coherence, we consider the *coherence gap*

$$\begin{aligned}\Delta\text{AUC} &= \int_0^1 \mathbb{E}[\psi_B - \psi_A \mid |B| = b]\,db\\ &= \text{AUC:BC} - \text{AUC:AC}.\end{aligned} \tag{18}$$

In the random map model, introduced in [26] (see Sec. II C),

$$\text{AUC:AC} = \text{AUC:BC} \approx \int_0^1 s\,ds = 0.5. \tag{19}$$

We thus use this value as a baseline for comparison.

In practice, we approximate these integrals using empirical conditional means computed over 1000 equal-sized sub-intervals of [0,1]. When an ensemble's basin size distribution exhibits gaps $\geqslant$0.02, a non-negligible proportion of the interval [0,1] lacks empirical data for computing $\mathbb{E}[\psi_B \mid |B| = s]$ and $\mathbb{E}[\psi_A \mid |B| = s]$. For such ensembles, the integral approximations underlying AUC:BC, AUC:AC, and ΔAUC are unreliable, and we therefore excluded them from our analysis. Specifically, this criterion meant we did not compute AUC:BC, AUC:AC, and ΔAUC for ensembles of BNs governed by highly biased NCFs with layer structures [5], [6], [4,2], [7], [5,2], [4,1,2], [4,3], and [3,1,3].

E. Bias, sensitivity, and dynamical regime

Sensitivity and bias are two key properties that shape the dynamical regime of Boolean networks. The *average sensitivity* $S(f)$ of a Boolean function $f(x_1, \ldots, x_n)$ quantifies how sensitive its output is to a change in a single random input [28]. That is,

$$S(f) = \frac{1}{2^n} \sum_{x \in \{0,1\}^n} \sum_{i=1}^{n} \mathbb{1}[f(x) \neq f(x \oplus e_i)]. \tag{20}$$

The mean of the average sensitivities of all update rules of a Boolean network is a proxy for its *dynamical regime* [6,12,29]. Networks with average sensitivity below 1 lie in the *ordered dynamical regime* where perturbations vanish over time. Conversely, networks with sensitivity above 1 lie in the *chaotic dynamical regime* where perturbations amplify. Between these two extremes lies the *critical dynamical regime*, associated with an average sensitivity value near 1.

In this paper, the *bias* p of a Boolean function is defined as the proportion of 1's in a function's truth table [30]. It quantifies the level of output asymmetry. Functions with $p = 0.5$ are *unbiased*, containing an equal number of 0s and 1s. We report bias using two related metrics: *absolute bias* $a(p) = |2p - 1|$ (ranging from 0 for unbiased to 1 for maximally biased functions) and *standardized bias* $\zeta(p) = p(1-p)$ (ranging from

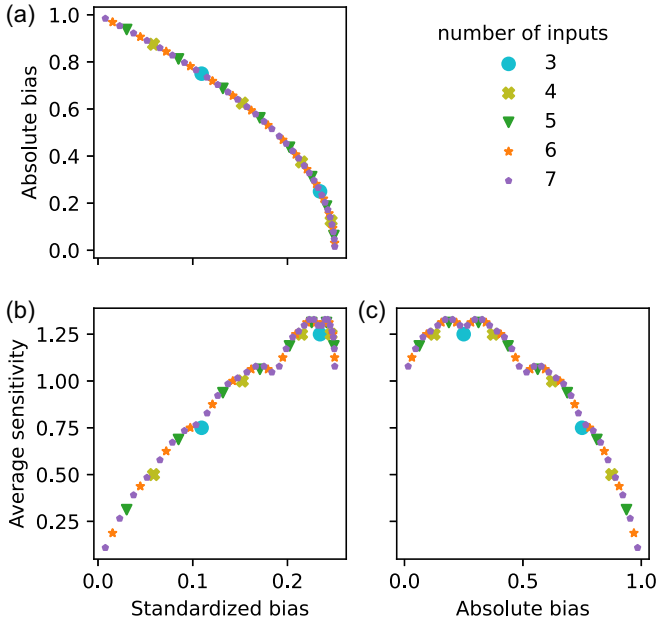


FIG. 3. Relationship between standardized bias, absolute bias, and average sensitivity for all nested canalizing functions (NCFs) with three to seven inputs. Each dot corresponds to an NCF of defined degree (color) and specific layer structure.

0 for maximally biased to 0.25 for unbiased functions). The standardized bias and absolute bias are negatively related [Fig. 3(a)]:

$$4\zeta(p) = 1 - a^2(p). \tag{21}$$

Importantly, bias and sensitivity are not independent [Figs. 3(b) and 3(c)]: higher bias generally leads to lower sensitivity, though this relationship is nonmonotonic [12]. Functions with extreme bias (p near 0 or 1) typically exhibit low sensitivity because their outputs are dominated by the preferred value. However, moderately biased functions exhibit varying sensitivity, depending on their degree of canalization.

F. Canalization

Canalization was introduced in developmental biology to describe the regulative behavior of developmental pathways, i.e., their tendency to return to a normal trajectory despite alterations in their state [1]. Consider a Boolean function $f(x_1, \ldots, x_n)$ of degree n. The variable x_i, $1 \leqslant i \leqslant n$, is a *canalizing variable* if there exists a fixed value $a \in \{0, 1\}$ such that when $x_i = a$, the output of f is fully determined, independently of the other inputs [31]. The function f is *canalizing* if it has a canalizing variable. This definition formalizes the idea of canalization by describing how specific variables can enforce stability in developmental pathways by effectively diminishing the impact of other influences.

Canalizing Boolean functions can be further classified by their *(canalizing) layer structure*, which can be computed iteratively [32]. First, identify all k_1 canalizing variables and their corresponding canalizing values. These variables and their values form the first canalizing layer. To determine the next layer, substitute the noncanalizing values into the function. This results in a new function that depends only on the variables not included in the first layer. Then, identify the new function's canalizing variables and their corresponding canalizing values. These form the second layer and their number is denoted by k_2. Repeat this process until either a constant function or a noncanalizing function is reached after $r \geqslant 0$ steps. The vector $(k_1, \ldots, k_r)$ constitutes the canalizing layer structure of the original function f [12,33]. The *canalizing depth* of f is $k = \sum_{i=1}^{r} k_i$ and describes the number of variables that become canalizing eventually. Note that, a Boolean function is canalizing when the canalizing depth $k \geqslant 1$. *Nested canalizing functions* (NCFs) are Boolean functions whose canalizing depth equals their degree. In other words, all the inputs of NCFs become eventually canalizing. Most updates rules in biological BNs are NCFs [10].

To illustrate how layer structure affects dynamics, consider the three-input Boolean functions $f(x, y, z) = x \wedge y \wedge z$ and $g(x, y, z) = x \wedge (y \vee z)$. Both are nested canalizing, but they have different sensitivities. The function f has layer structure (3): all three variables are canalizing. Setting any variable to zero fixes the output to zero regardless of other inputs. The function has low sensitivity, i.e., it is robust to perturbations. In contrast, g has layer structure (1,2): only x is canalizing. When $x = 1$ (its noncanalizing value), the function reduces to $y \vee z$, revealing a second layer where both y and z become canalizing. Such a function has higher sensitivity because more variable combinations can alter the output.

NCFs with the same layer structure have the same average sensitivity and absolute bias [12]. For example, the function f has absolute bias $3/4$ because its truth table contains seven 1s. Similarly, the function g contains five 1s and thus has absolute bias $1/4$. For a given number of variables $n \geqslant 2$, there are 2^{n-2} NCFs with different layer structure, each appearing equally likely by chance. Table I illustrates this for $n = 5$, showing how layer structure determines both bias and sensitivity—a relationship we exploit in our results.

For a broader treatment and alternative definitions of canalization in BNs, we refer the reader to [34–37]. In these papers, the traditional definition of canalization by Kauffman is generalized to the *canalizing strength*, which quantifies the level of collective canalization in a Boolean function [36]. Canalization can also be interpreted as *input redundancy*, which is the probability that an input to a Boolean function is not needed to determine its output [35,37]. By contrast, the *effective degree* describes the average number of inputs needed to determine the output.

In this paper, we present evidence that the gap between attractor and basin coherence widens as the level of canalization in the network increases, irrespective of how canalization is assessed.

G. Analysis of biological network models

In this study, we considered two collections of BN models: synthetic random networks and published biological models. The latter is a collection of 122 distinct models that contain between 5 and 342 nodes [10]. For small BNs, the entire state transition graph, basin coherences, and basin sizes can be computed quickly. However, as the number of nodes increases, the exponential growth of the state space makes it infeasible to compute basin sizes and basin coherences

TABLE I. Stratification of all five-input NCFs by layer structure, similar to [12].

Layer structure	Example	Hamming weight	Absolute bias	Standardized bias	Average sensitivity
(5)	A ∨ B ∨ C ∨ D ∨ E	1 or 31	15/16	0.0303	0.3125
(3, 2)	A ∨ B ∨ C ∨ (D ∧ E)	3 or 29	13/16	0.0850	0.6875
(2, 1, 2)	A ∨ B ∨ [C ∧ (D ∨ E)]	5 or 27	11/16	0.1318	0.9375
(2, 3)	A ∨ B ∨ (C ∧ D ∧ E)	7 or 25	9/16	0.1709	1.0625
(1, 1, 3)	A ∨ [B ∧ (C ∨ D ∨ E)]	9 or 23	7/16	0.2021	1.1875
(1, 1, 1, 2)	A ∨ (B ∧ [C ∨ (D ∧ E)])	11 or 21	5/16	0.2256	1.3125
(1, 2, 2)	A ∨ [B ∧ C ∧ (D ∨ E)]	13 or 19	3/16	0.2412	1.3125
(1, 4)	A ∨ (B ∧ C ∧ D ∧ E)	15 or 17	1/16	0.2490	1.1875

exactly. Therefore, we conducted two types of biological model analyses: one for smaller networks ($N \leqslant 20$) and another for larger networks ($N \leqslant 64$).

For the first type of analysis, from the full collection of 122 biological BN models, we considered the 42 models that have $N \leqslant 20$ nodes. For these networks, we computed the entire state transition graph for up to 16 different fixed-source networks as described below. Consequently, we obtained exact values for the basin sizes, basin coherences ψ_B, and attractor coherences ψ_A.

For the second type of analysis, we considered the 95 models with $N \leqslant 64$. For these networks, we approximated the state transition graph using `BoolForge` [38]. We generated $q = 1000$ random initial states $\mathbf{x} \in \{0, 1\}^N$ and repeatedly updated $\mathbf{x}$ until it transitioned into a periodic orbit, i.e., to an attractor. The proportion of random initial states that transition to a specific attractor provides an unbiased estimate for the corresponding relative basin size. Similarly, the coherence of these states offers an unbiased estimate of basin coherence. This procedure identifies an attractor of relative basin size $b \in (0, 1]$ with probability $p(b, N) = 1 - (1 - b)^q$ [39]. Consequently, attractors with very small basins are often overlooked. However, biologically meaningful attractors typically have sufficiently large basin sizes and are discovered with high probability. For example, an attractor with relative basin size $b = 1\%$ is found with probability 99.996%. Irrespective of the network size, we computed the exact attractor coherence ψ_A for every identified attractor A. Thus, for this analysis, we obtained the exact value of ψ_A and unbiased estimates of ψ_B and basin sizes for all the identified attractors.

H. Treatment of source nodes

Some of the biological network models contain *source nodes*—variables with no regulatory inputs whose values remain constant throughout the dynamics [16]. For a network with M source nodes, the full state transition graph of size 2^N decomposes into 2^M independent subgraphs, each corresponding to one fixed configuration of source node states. Each subgraph evolves under identical regulatory logic for the remaining N-M nodes but with different constant inputs from the sources. This partitioning can artificially inflate the number of attractors and deflate basin sizes, confounding comparisons of coherence between networks that differ in the number of source nodes.

To remove this effect, we analyzed only *fixed-source networks* in the biological datasets, where the states of the source nodes were held constant during the dynamical simulations. For networks with multiple source nodes, we randomly analyzed up to 16 of the 2^M independent state transition graphs and average the resulting attractor statistics. This approach isolates the contribution of the regulatory logic from trivial fragmentation effects introduced by variable numbers of source nodes.

I. Analysis of random Boolean networks

Random models offer a useful baseline for interpreting how various features of biological networks shape their dynamics. We conducted two experiments on random ensembles of BNs to examine how canalization influences the relationship between attractor and basin coherence. In both experiments, each ensemble consisted of 10 000 randomly generated 12-node networks, created and exactly analyzed using the Python package `BoolForge` [38]. In the first experiment, all networks within an ensemble share update rules with the same degree and canalizing depth, while these parameters vary across different ensembles. We ensured that (i) the wiring diagrams of all networks were strongly connected, and (ii) the Boolean functions were nondegenerate (i.e., function outputs genuinely depend on all input variables indicated in the wiring diagram, not a subset). This excludes, for example, $f(x, y, z) = x \vee y$ as a three-input update function. In the second experiment, all networks within an ensemble were governed by the same type of NCF, specified by canalizing layer structure.

III. RESULTS

A. Basin and attractor coherence in biological networks

A recent meta-analysis of 122 published, expert-curated biological BN models revealed that they possess a number of remarkable structural and dynamical features [10]. They are sparsely connected, and the Boolean rules describing the regulatory logic are significantly more canalizing, redundant and biased than expected by chance. To assess how these features influence the robustness of biological BNs, we first considered the 42 models of size $N \leqslant 20$. These models are small enough to be simulated exhaustively, allowing us to compute their coherence exactly.

For each of the 42 biological models, we generated 100 corresponding random null models. These null models preserve the wiring diagram but used randomized, nondegenerate update rules. To examine the difference in coherence between

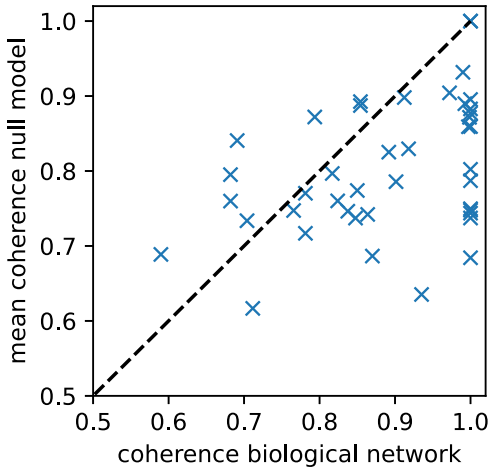


FIG. 4. Biological Boolean network models exhibit higher-than-expected coherence. For each of the 42 biological BN models from [10] with size $N \leqslant 20$, the network coherence is plotted against the mean coherence of 100 corresponding random null models. The null models have the same wiring diagram but use randomized, nondegenerate update rules. The dashed diagonal line indicates equality between biological and expected coherence.

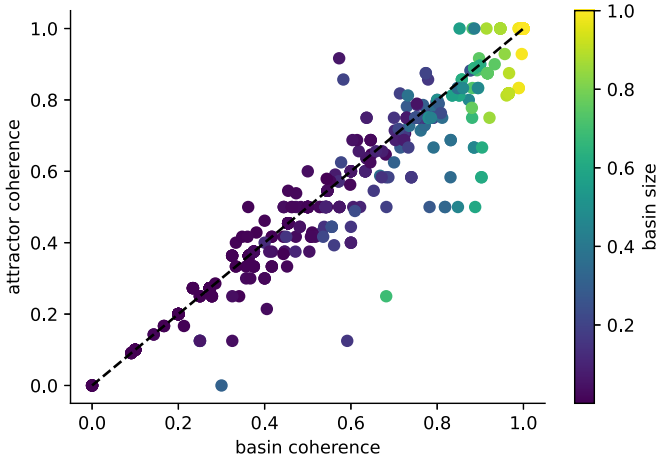


FIG. 5. Exact basin and attractor coherence of biological Boolean network models. For each attractor found in Boolean biological network models in [10] of size $N \leqslant 20$ (42 models), attractor coherence is plotted against basin coherence. Color indicates the relative basin size. A substantial number of attractors lie below the diagonal (basin coherence > attractor coherence; sign test $p = 2 \times 10^{-11}$), indicating that attractors are on average less resilient to perturbations than their basins.

the null models and the biological models, we computed the difference between the coherence of the biological networks and the mean coherence of the corresponding random models. The distribution of paired coherence differences did not deviate from normality (Shapiro-Wilk $p = 0.31$), and a one-sided paired t-test confirmed that biological networks exhibited significantly higher coherence than the random null models with the same topology ($p = 9 \times 10^{-6}$, $n = 42$; Fig. 4). This demonstrates that the high coherence in biological networks is not solely due to their network topology but likely arises from their specific regulatory logic.

To further investigate the mechanisms underlying this enhanced stability, we examined the coherence of basins and attractors in biological networks. For each network attractor, we computed the basin size, as well as the basin and attractor coherence. Across all attractors, larger basins are generally more robust (Fig. 5), and basin and attractor coherence are, unsurprisingly, highly correlated (Pearson's r = 0.97). It is known that BNs with low connectivity have high average basin coherence, while highly connected BNs possess basins that are almost randomly distributed throughout the state space [14]. Biological BNs have relatively low connectivity (mean = 2.56, median = 2.27) [10]. This partly explains why even attractors with relatively small basins were on average fairly coherent.

We observed a significant number of attractors with high basin coherence but relatively low attractor coherence. Because the distribution of differences is non-normal and asymmetric, we used the sign test to measure the statistical significance of our observations ($p = 2 \times 10^{-11}$, $n = 597$; Fig. 5). This suggests that some attractors lie near the "boundary" of their basin.

This pattern partially persists when sampling the dynamics of the 95 biological BN models of size $N \leqslant 64$ (Fig. 6). For attractors with relatively large, biologically relevant basins, the average basin coherence is higher than the average attractor coherence. In contrast, attractors with tiny basins exhibit on average the opposite trend. Altogether, the area under the basin coherence curve (AUC:BC = 0.853) is larger than the area under the attractor coherence curve (AUC:AC = 0.825). This is consistent with the landscape model illustrated in Fig. 1, where attractors occupy positions near relatively flatter boundaries of their basins. To disentangle what features of biological networks render their attractors less robust than their basins, we next investigated large ensembles of random BNs with a fixed in-degree.

B. Canalization increases both basin and attractor coherence but creates a coherence gap

One key feature of biological BN models is their high level of canalization [8,10,18,37]. To assess the effect of canalization on the basin and attractor coherence, we generated ensembles of 10 000 12-node BNs for different combinations of degree and canalizing depth (Figs. 7 and 8).

In the absence of canalization, increasing the degree led to randomly distributed basins in the state space. That is, both basin and attractor coherence converged toward the random baseline given by the relative basin size. For example at degree 3, the AUC:BC and AUC:AC values were 18.6% and 20% above the random baseline of 0.5, respectively [Fig. 7(a)]. These values dropped to 7.2% and 8.2% at degree 5 [Fig. 7(b)]. In contrast, canalizing and nested canalizing networks exhibited substantially higher basin and attractor coherence, regardless of degree. Specifically, the AUC:BC values were 27% and 27.6% higher in canalizing networks of degree 3 and 5, respectively [Figs. 7(c) and 7(d)]. In nested canalizing networks, these numbers rose to 50.8% and 50.6%, respectively [Figs. 7(e) and 7(f)]. These observations are

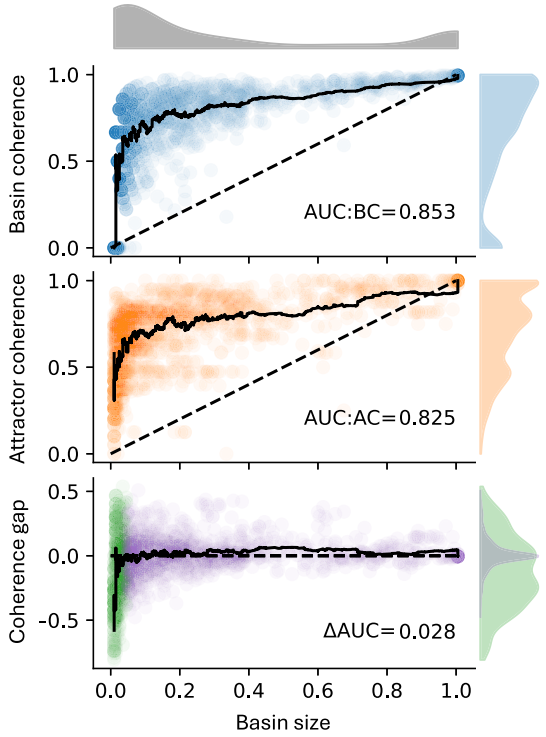


FIG. 6. Coherence of biological Boolean network attractors and their basins. The synchronous state transition graph (attractors, basin coherence, attractor coherence) was approximated for all Boolean biological network models in [10] of size $N \leqslant 64$ (95 models; see Methods for details). The approximated basin coherence (top; blue), the exact attractor coherence (middle; orange), and the difference per attractor (bottom) are plotted against the approximate basin size (x-axis) for each discovered attractor in the biological networks. Black lines depict an unweighted rolling-window mean with window size 50. Distributions of the different metrics are color-coded. The distribution of the difference in basin minus attractor coherence is shown for basin sizes $< 5\%$ (green) and $\geqslant 5\%$ (purple). AUC quantifies the area under the respective lines.

consistent with earlier findings that biological networks, typically highly canalizing, tend to be highly coherent (Fig. 4).

Canalization increases both basin and attractor coherence in absolute terms, but increases basin coherence more strongly, creating a coherence gap. This implies that transient states in canalizing and nested canalizing networks are generally more coherent than attractor states, while there exists hardly any difference in noncanalizing networks. For canalizing networks, the AUC:AC values were only 20.8% higher than the random baseline [Figs. 7(c) and 7(d)]. This represents a relative drop of $\Delta\text{AUC}/(\text{AUC:BC} - 0.5) = 22.9\%$ for the case $n = 3$ and 25.2% for the case $n = 5$. In nested canalizing networks, AUC:AC values were 38.4 to 38.8% higher than the baseline, corresponding to a similar relative drop of 24.4% ($n = 3$) and 23.3% ($n = 5$). Biological BNs are often dominated by NCFs [10], which may explain the observed coherence gap in biological systems.

There are several alternative measures of canalization in Boolean functions, including canalizing strength and input redundancy [36,37]. In this study, we chose canalizing depth as our primary metric because the standard monomial form of a Boolean function provides a straightforward method for generating networks with a specified canalizing depth [32]. Functions with higher canalizing depth tend to exhibit higher canalizing strength and input redundancy (Fig. 8). These high correlations suggest that the findings are insensitive to the specific choice of canalization metric.

C. Sensitivity drives basin coherence

Sensitivity and bias are two key network properties that shape a network's dynamical behavior. Both properties are closely linked to canalization. In particular, canalizing functions tend to be biased [20], and they typically exhibit low sensitivity. Furthermore, other well-established predictors of a network's dynamical regime, such as the metric $\langle K_e \rangle \langle p(1-p) \rangle$, which combines average effective degree and average standardized bias, are also highly correlated with canalization [10,40]. These relationships suggest that the dynamical regime of a network, shaped by canalization and bias, may play a key role in determining the coherence of its basins and attractors (Fig. 8).

To systematically dissect the roles of sensitivity and bias in shaping the observed coherence gap, we reanalyzed the networks governed by NCFs [Figs. 7(e) and 7(f)]. To achieve finer resolution, we further stratified NCFs according to their canalizing layer structure. Because NCFs with the same layer structure exhibit identical bias and average sensitivity [12], this approach allowed us to independently examine how sensitivity and bias influence basin coherence, attractor coherence, and the coherence gap. Specifically, for each five-input NCF with different layer structure, we generated an ensemble of 10 000 12-node BNs, yielding networks with varying levels of sensitivity and bias (Table I).

Basin coherence was high across all network ensembles, but showed clear sensitivity dependence (Fig. 9). Networks governed by NCFs with the highest average sensitivity of 1.3125—those with layer structures (1,1,1,2) and (1,2,2)—exhibited the lowest AUC:BC values [Table I, Figs. 9(f) and 9(g)]. In contrast, bias proved to be a weak predictor of basin coherence. For example, NCFs with layer structure (1,4), which are nearly unbiased (absolute bias $= 1/16$), and NCFs with layer structure (3,2), which are highly biased (absolute bias $= 13/16$), both exhibited similarly high AUC:BC values. These findings suggest that the sensitivity plays a more important role than bias in determining AUC:BC values.

Repeating this analysis for network ensembles with different connectivity (3-7 inputs per node) revealed a consistent pattern. AUC:BC values were consistently lowest in networks governed by the most sensitive NCFs [Fig. 10(a)]. However, the observed relationship between AUC:BC and the average sensitivity was far from monotonic (Spearman's $\rho = -0.629$): Critical networks (sensitivity ≈ 1) exhibited locally minimal basin coherence, while the most ordered networks (lowest sensitivity) did not possess maximal AUC:BC values. Similar trends persisted when comparing AUC:BC values with the standardized bias of the NCFs [Fig. 10(b); Spearman's $\rho = -0.312$], with one notable exception—networks governed by the least biased NCFs maintained relatively high basin coherence. This exception is in accordance with the underlying nonmonotonic relationship between bias and sensitivity

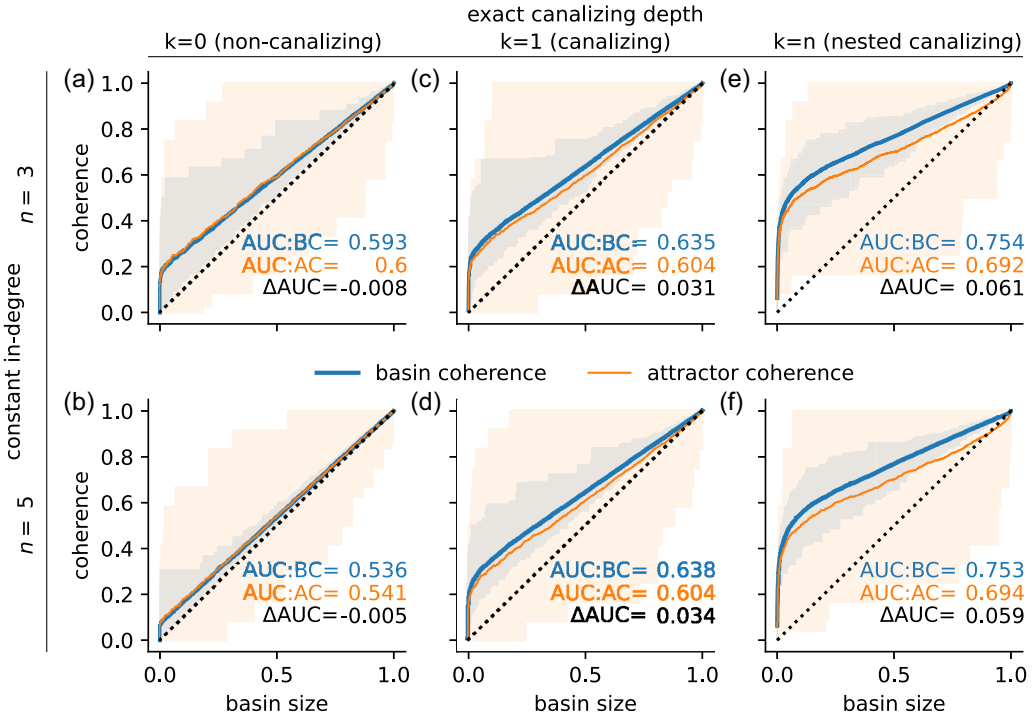


FIG. 7. Robustness of random Boolean network attractors and their basins. For 10 000 random 12-node BNs with random (a), (b), canalizing (c), (d), and nested canalizing (e), (f) update rules of in-degree 3 (a), (c), (e) and 5 (b), (d), (f), the coherence of each attractor (orange) and the coherence of its corresponding basin (blue) are stratified by basin size (x-axis). Shaded areas indicate the range of observed values. Lines depict rolling-window means with window size 500 and AUC quantifies the area under the respective lines.

[Figs. 3(b) and 3(c)]: NCFs with low absolute bias exhibited lower sensitivity values.

D. Bias drives attractor coherence

In contrast to basin coherence, attractor coherence varied substantially across network ensembles and depended more strongly on bias than sensitivity [Figs. 10(c) and 10(d)]. Generally, the lower the standardized bias of the NCFs (i.e., the higher the absolute bias), the lower was the observed attractor coherence (Spearman's $\rho = 0.983$). This suggests that, in networks governed by more biased NCFs, attractors with small and moderately sized basins are less coherent. Such attractors are, however, also less frequent in these biased networks (see Figs. 12, 14, and 16).

	Degree					
	n=3			n=5		
	Exact canalizing depth					
	k = 0	k = 1	k = 3	k = 0	k = 1	k = 5
AUC basin coherence (AUC:BC)	0.593	0.635	0.754	0.536	0.638	0.753
AUC attractor coherence (AUC:AC)	0.600	0.604	0.692	0.541	0.604	0.694
Coherence gap (ΔAUC)	-0.008	0.031	0.061	-0.005	0.034	0.059
Average canalizing strength	0.267	0.5	0.777	0.171	0.437	0.618
Average normalized input redundancy	0.271	0.375	0.531	0.308	0.534	0.685
Average effective degree, K_e	2.186	1.875	1.407	3.460	2.332	1.576
Average standardized bias, p(1-p)	0.239	0.188	0.171	0.242	0.184	0.166
Product, K_e p(1-p)	0.520	0.352	0.252	0.839	0.430	0.277
Average sensitivity	1.800	1.5	1	2.502	1.518	1

FIG. 8. Properties of noncanalizing, canalizing, and nested canalizing Boolean functions and networks governed by them. For each metric and a given degree, a separate linear color gradient is used.

Intriguingly, the attractor coherence of networks in the ordered dynamical regime, governed by NCFs with layer structure (3,2) and (2,1,2) with average sensitivity below 1 (i.e., with ordered dynamics), appeared to be nonmonotonic in basin size [Figs. 9(b) and 9(c)]. For example, attractors in networks governed by NCFs with layer structure (3,2) that attract roughly half the states were on average less coherent than attractors with a basin size of around 30%. More connected networks governed by six-input and seven-input NCFs exhibited the same trends, indicating robust results (see Figs. 13–16).

E. Bias perfectly explains the coherence gap

The coherence gap exhibited a remarkably strong linear relationship with standardized bias [Spearman's $\rho = -0.997$; Fig. 10(f)]. Networks governed by least biased NCFs (with standardized bias $p \approx 0.25$) possessed roughly equal attractor and basin coherence. As the standardized bias decreases (i.e., as the absolute bias increases), the coherence gap increased nearly linearly. This pattern held across all investigated in-degrees and suggests that bias, modulated through canalization [12,20], largely determines whether, and by how much, attractor and basin coherence differ.

We hypothesized that this arises because highly biased networks confine attractors to similar regions of state space. In these networks, most nodes become frozen in the same state across attractors, with dynamics confined to a small non-frozen core [41–43]. This clustering reduces the mean Hamming distance between attractors relative to random state pairs, creating relatively flatter basin boundaries near

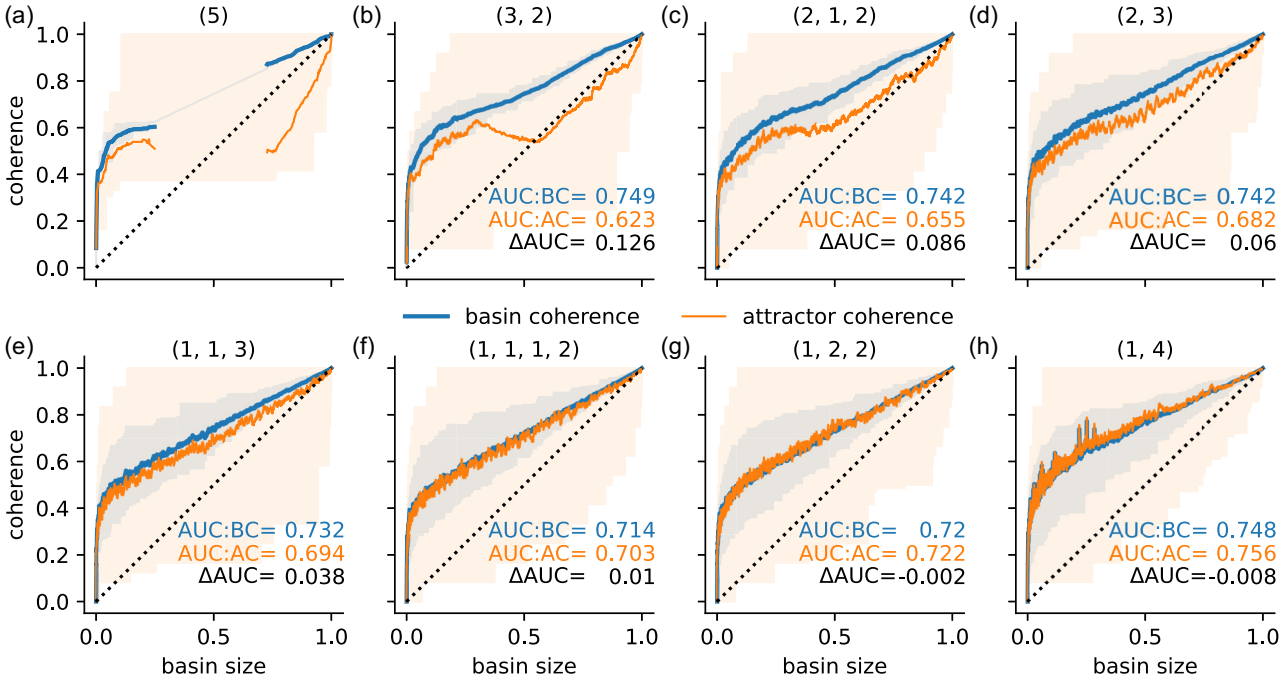


FIG. 9. Robustness of nested canalizing Boolean network attractors and their basins. For 10 000 random 12-node BNs governed by five-input NCFs of different layer structure (specified above each panel), the coherence of each attractor (orange) and the coherence of its corresponding basin (blue) are stratified by basin size (x-axis). Shaded areas indicate the range of observed values. Lines depict rolling-window means with window size 50, and AUC is the area under the respective lines. Figure 12 shows the basin size distribution for each network type.

attractors compared to mid-trajectory states, while maintaining high absolute coherence throughout the basin.

To test this hypothesis, we computed the average number of frozen nodes and the average distance between attractors in the ensembles of 12-node BNs governed by NCFs of different layer structure. We found strong monotonic trends (Fig. 11): The higher the absolute bias of a nested canalizing network (i.e., the lower the standardized bias), the more nodes were on average frozen and the lower was the average Hamming distance between attractors. These trends persisted when only considering networks in each ensemble with a fixed number of attractors.

These findings provide the mechanism underlying the within-basin stability gradient shown in Fig. 1: bias-driven compression of attractors into similar state-space regions creates the coherence gap. The biological relevance of this gradient depends on whether multiple attractors with substantial basins coexist (see Fig. 12).

IV. DISCUSSION

Our analyses demonstrate that biological Boolean networks exhibit markedly higher coherence than random ensembles of comparable size and connectivity. This enhanced resilience to perturbations can be traced back to characteristic design principles of biological systems, such as abundant canalization and high bias. Each of these properties has been observed consistently across curated biological network models and contributes to systematic dynamical differences between biological and synthetic networks [8,10,29,44].

A. Attractors are less stable than their basins

A central result of this study is the discovery that attractor coherence is lower than basin coherence in biological networks. This finding revises our understanding of phenotypic stability. Mature cell types and phenotypes still maintain high absolute coherence. However, rather than occupying the most stable positions in the regulatory landscape, they are positioned closer to basin boundaries than transient states. This within-basin stability gradient means that, within the same basin, perturbations of attractor states are more likely to induce phenotype switching than perturbations of transient states.

B. Canalization creates the coherence gap through bias

Our systematic analysis of random network ensembles reveals the mechanism underlying the coherence gap. Canalization—a pervasive feature of biological regulation where specific inputs can override others—simultaneously increases overall network stability and magnifies the coherence gap. This dual effect arises because canalization induces strong bias in Boolean functions [12,20].

The coherence gap is almost perfectly predicted by network bias (Spearman's $\rho = -0.997$), with highly biased networks showing the largest disparities between basin and attractor coherence. One possible explanation is that attractors in highly biased and canalized networks typically cluster in similar regions of the state space. This occurs because high bias is known to cause many nodes to become frozen (always off or always on), leaving dynamics confined to a small nonfrozen core [41–43]. Because attractors are concentrated

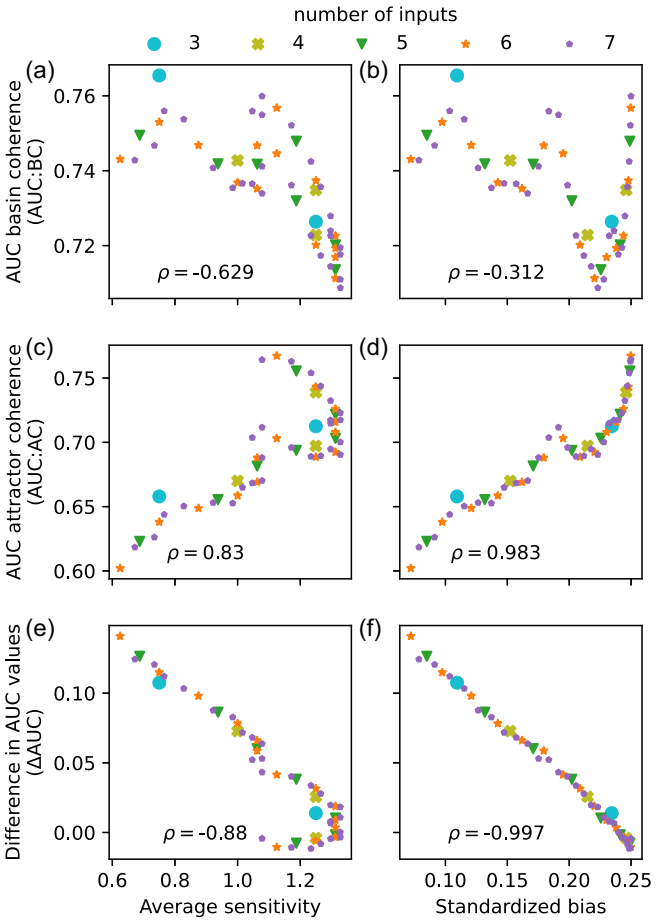


FIG. 10. Association between different metrics. For ensembles of random 12-node BNs governed by nested canalizing update rules in three to seven inputs (color) and of different layer structure, AUC:BC (top row), AUC:AC (middle row), and the coherence gap (bottom row) are computed as in Fig. 9, and plotted against the average sensitivity (left column) and the standardized bias $p(1-p)$ (right column). Each data point is derived from 10 000 networks. Across all data points, Spearman's ρ correlation coefficient is computed.

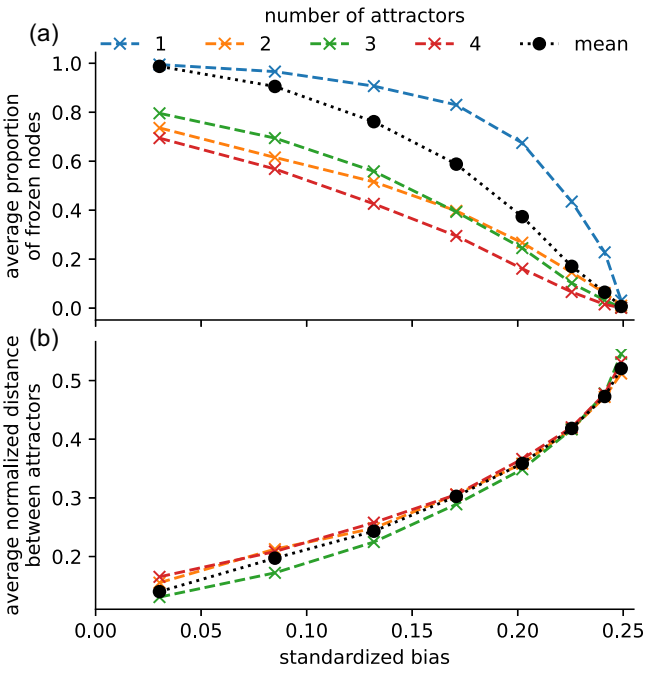


FIG. 11. Attractor similarity in nested canalizing Boolean networks. For 10 000 random 12-node BNs governed by five-input NCFs of different layer structure, the (a) average proportion of frozen nodes and the (b) average normalized distance between attractors is plotted against the standardized bias associated with all NCFs of specific layer structure. Results are stratified by the number of network attractors (dashed colored lines). In (b), networks with a single attractor do not contribute to the mean (dotted black line).

in this small core, they sit closer together in state space, reducing the mean Hamming distance between attractors and effectively flattening the ridges separating basins near their "valley floors." Small perturbations at these attractor locations can thus easily induce inter-attractor transitions.

Importantly, this relationship holds across multiple measures of canalization, including canalizing depth [32], canalizing strength [36], and input redundancy [37]. This indicates robustness of our findings to methodological choices. The dynamical regime (ordered, critical, or chaotic) plays a secondary role, primarily influencing basin coherence, while bias dominates attractor coherence.

C. Context-dependent relevance of the coherence gap

Consistent with the high robustness of biological systems, canalization increases both basin and attractor coherence in absolute terms. At the same time, canalization increases the coherence gap, suggesting the presence of a within-basin stability gradient (Fig. 1). The biological relevance of the coherence gap depends strongly on the basin size distribution.

In highly biased, canalized networks, the number of attractors is often small (see Fig. 12). These networks may be dominated by one or a few major attractors occupying most of state space. In such cases, the coherence gap has limited functional impact because alternative phenotypes have been effectively or completely eliminated.

However, the coherence gap becomes functionally important in biological contexts where multiple phenotypically meaningful attractors coexist. In these contexts, our finding that attractors are less stable relative to their basins predicts enhanced phenotypic plasticity despite overall high stability. The impact of canalization is thus context-dependent: protective when one fate dominates (high absolute coherence, coherence gap irrelevant), and permissive when choices exist (high absolute coherence, but the coherence gap facilitates phenotypic transitions). This dual nature may help resolve the apparent paradox of how biological systems can be simultaneously robust and adaptable.

D. Balancing robustness and adaptability: Insights into the criticality hypothesis

Our findings provide insights into the longstanding criticality hypothesis, which posits that biological systems balance between order and chaos to maintain both robustness and adaptability (see, e.g., [7,9,30,45]). Park *et al.* (2023)

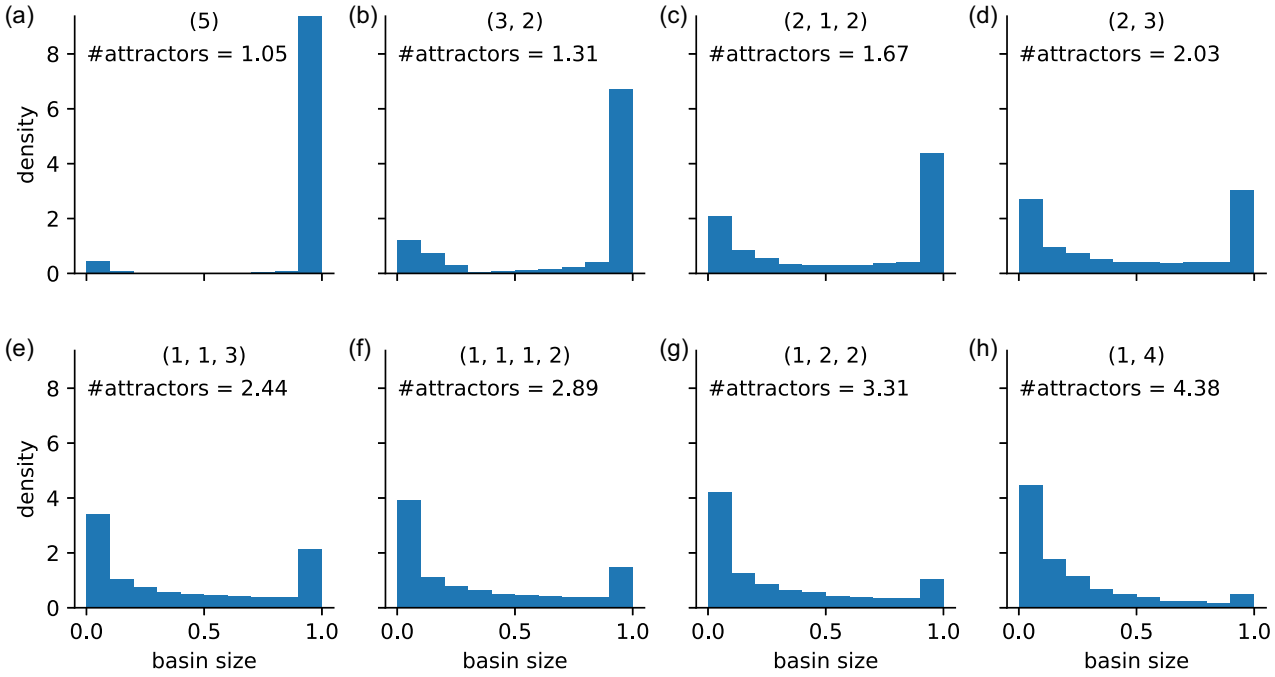


FIG. 12. Basin size distribution for different types of nested canalizing Boolean networks. For 10 000 random 12-node BNs governed by five-input nested canalizing update rules of different layer structure (specified above each panel), the distribution of the basin size is shown, in addition to the average number of network attractors.

recently challenged this hypothesis: biological networks are substantially more robust than previously appreciated when disregarding perturbations of source nodes [16]. However, in their analysis, they considered perturbations of any network state. Since biological networks rest at attractors, attractor coherence primarily determines the resilience to perturbations. Our direct examination of attractor stability reveals that biological attractors are, on average, less stable than the transient states approaching them. This indicates that biological systems operate closer to the critical edge than suggested by Park *et al*.

Canalization and bias create a dual-purpose architecture operating at two levels. First, at the level of absolute stability, both basin and attractor coherence are high, ensuring that random perturbations rarely cause phenotype switching. This explains developmental robustness—stochastic fluctuations do not derail differentiation, and mature cell types resist random noise. Second, within each basin, a stability gradient positions attractors closer to boundaries relative to transient states, enabling regulated phenotype transitions in response to appropriate signals.

This architecture may be evolutionarily advantageous. During normal development, high absolute coherence protects cells from stochastic perturbations as they follow canalized paths through state space. The basin's depth ensures reliable lineage commitment. However, when cells reach decision points where multiple fates have substantial basins, the coherence gap provides a mechanism for controlled plasticity: Cells can respond to developmental signals, injury, or stress, by transitioning to alternative fates without requiring dramatic rewiring of regulatory networks. The system is robust by default but plastic on demand.

E. Implications for development, disease, reprogramming, and design

The attractor-centric perspective introduced here has potentially broad implications across biological and engineering contexts.

Development and cell fate decisions. High basin coherence during differentiation ensures robust lineage commitment, while the within-basin stability gradient means attractors are less stable than the trajectories themselves. At developmental branch points where multiple fates remain accessible, the coherence gap indicates that mature differentiated cells should be more susceptible to fate-switching perturbations than cells mid-differentiation. This generates testable hypotheses: time-resolved reprogramming experiments should show that mature cells respond more readily to reprogramming factors than cells caught mid-differentiation, and developmental stages where multiple fates have comparable basin sizes should exhibit higher phenotypic plasticity.

Cancer and disease. In normal tissues, regulatory networks typically produce one dominant cellular phenotype. However, in cancer, mutations can create multiple partially stable phenotypic states—epithelial, mesenchymal, and hybrid intermediates—each with non-negligible basin sizes [2,22,23]. When such multiattractor landscapes emerge, our results suggest enhanced phenotypic plasticity: cancer cells more readily undergo epithelial-mesenchymal transition, dedifferentiation, or therapy-induced state switching.

Cellular reprogramming. The relative fragility of attractors provides a mechanistic explanation for why reprogramming of mature cells is possible. Differentiated cells, although occupying "deep valleys" with high absolute coherence, often sit near basin boundaries that can be breached by

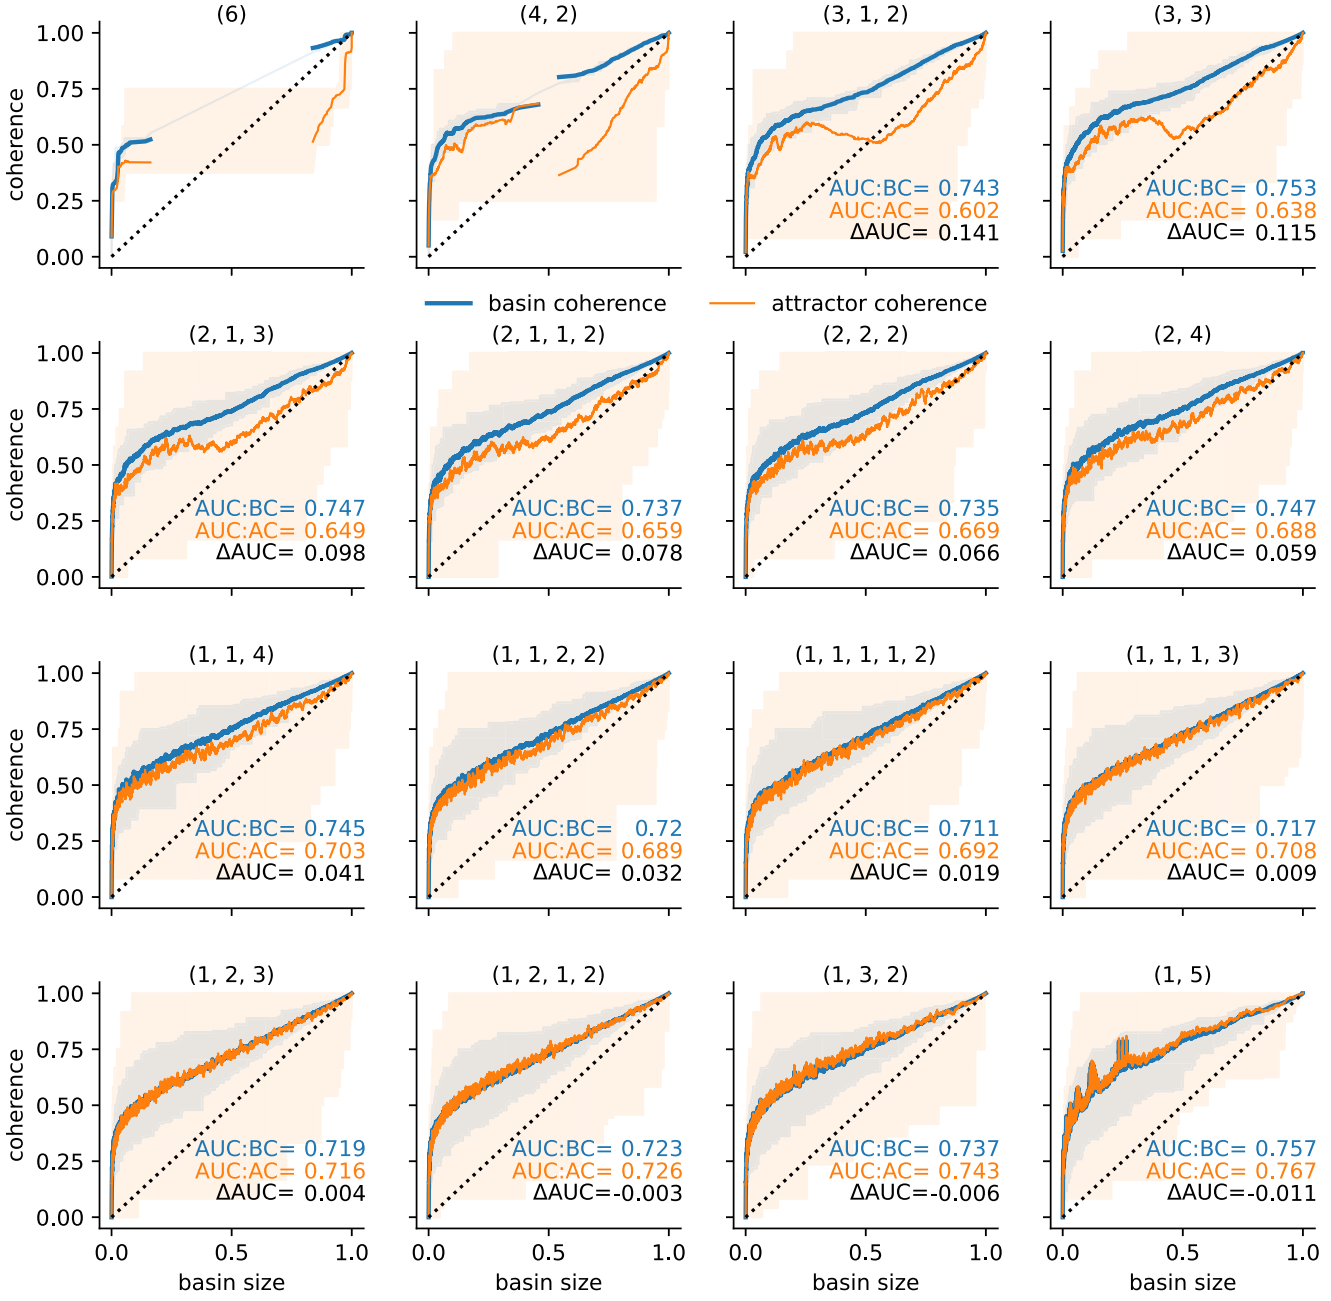


FIG. 13. Robustness of nested canalizing Boolean network attractors and their basins. For 10 000 random 12-node BNs governed by five-input nested canalizing update rules of different layer structure (specified above each panel), the coherence of each attractor (orange) and the coherence of its corresponding basin (blue) are stratified by basin size (x-axis). Shaded areas indicate the range of observed values. Lines depict rolling-window means with window size 50, and AUC is the area under the respective lines. The basin size distribution corresponding to each type of network is shown in Fig. 14.

suitable perturbations, such as transcription-factor cocktails [46,47]. This framework makes counterintuitive predictions: perturbations applied to fully differentiated cells may induce reprogramming more efficiently than perturbations applied during differentiation, and highly canalized cell types may, paradoxically, be easier to reprogram from their mature state than from intermediate differentiation stages.

Synthetic biology and bioengineering. For the design of synthetic gene circuits, our results highlight a potential trade-off between robustness and controllability: canalization increases absolute coherence but creates coherence gaps when multiple functional states exist. For single-state circuits (biosensors, memory devices), canalization and bias should be maximized to create a dominant attractor with high stability. For multistate circuits (toggle switches, oscillators), moderate canalization and bias may provide the best balance between robustness and controllability.

F. Limitations and future directions

The coherence gap metric (ΔAUC) removes confounding effects of basin size distribution but may overestimate the importance of the coherence gap in highly ordered networks

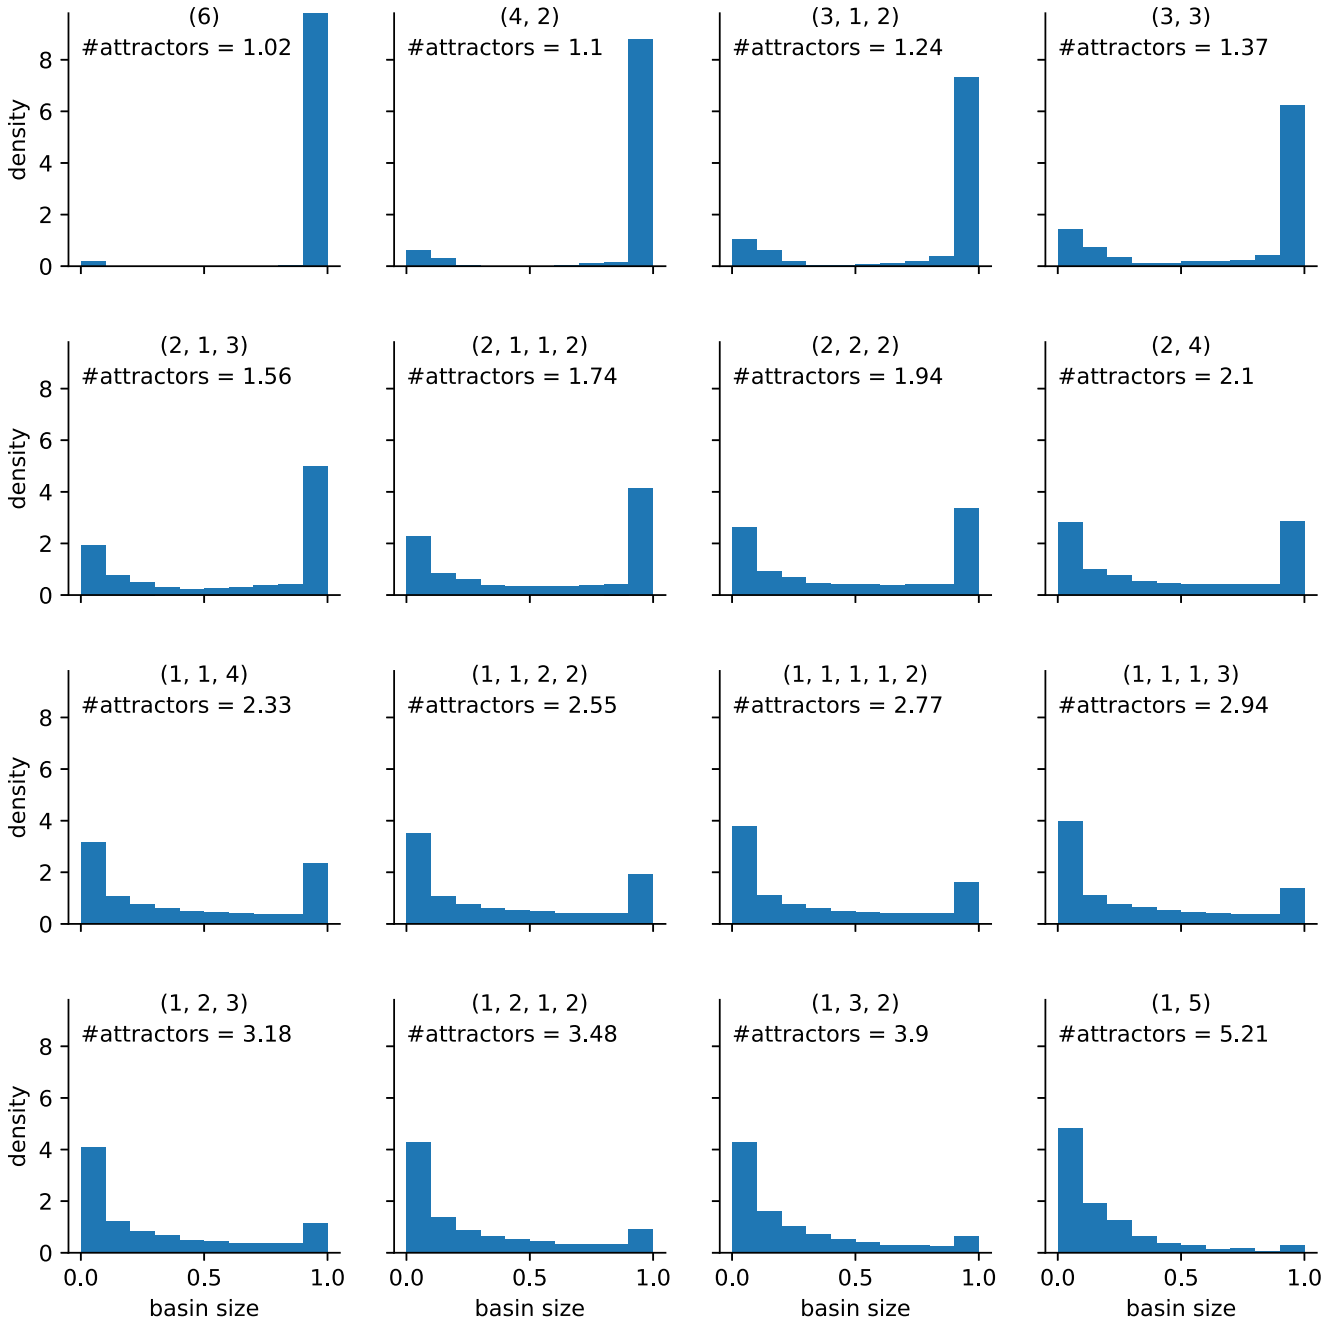


FIG. 14. Basin size distribution for different types of nested canalizing Boolean networks. For 10 000 random 12-node BNs governed by six-input nested canalizing update rules of different layer structure (specified above each panel), the distribution of the basin size is shown, in addition to the average number of network attractors.

where midsized basins (which exhibit the highest coherence gap) are rare. Alternative metrics that weight by biological relevance or phenotypic importance could provide complementary perspectives.

Additionally, synchronously updated Boolean networks, while capturing key regulatory logic, necessarily simplify quantitative details of gene expression dynamics, post-translational modifications, and metabolic constraints. Extensions to asynchronously updated, continuous or multistate models could reveal additional nuances of the stability gradient within basins. However, extending coherence measures to asynchronous or stochastic update schemes is nontrivial. Under stochastic dynamics, states may reach multiple attractors with different probabilities, leading to overlapping and probabilistic basins. This raises conceptual and definitional challenges: how should basin coherence be defined when basin membership is no longer deterministic? Moreover, stochastic update schemes can yield nontrivial attractors with more complicated structures, often lacking proper periodicity [48]. Developing a probabilistic extension of coherence that accounts for destination uncertainty is an important direction for future work. While such models

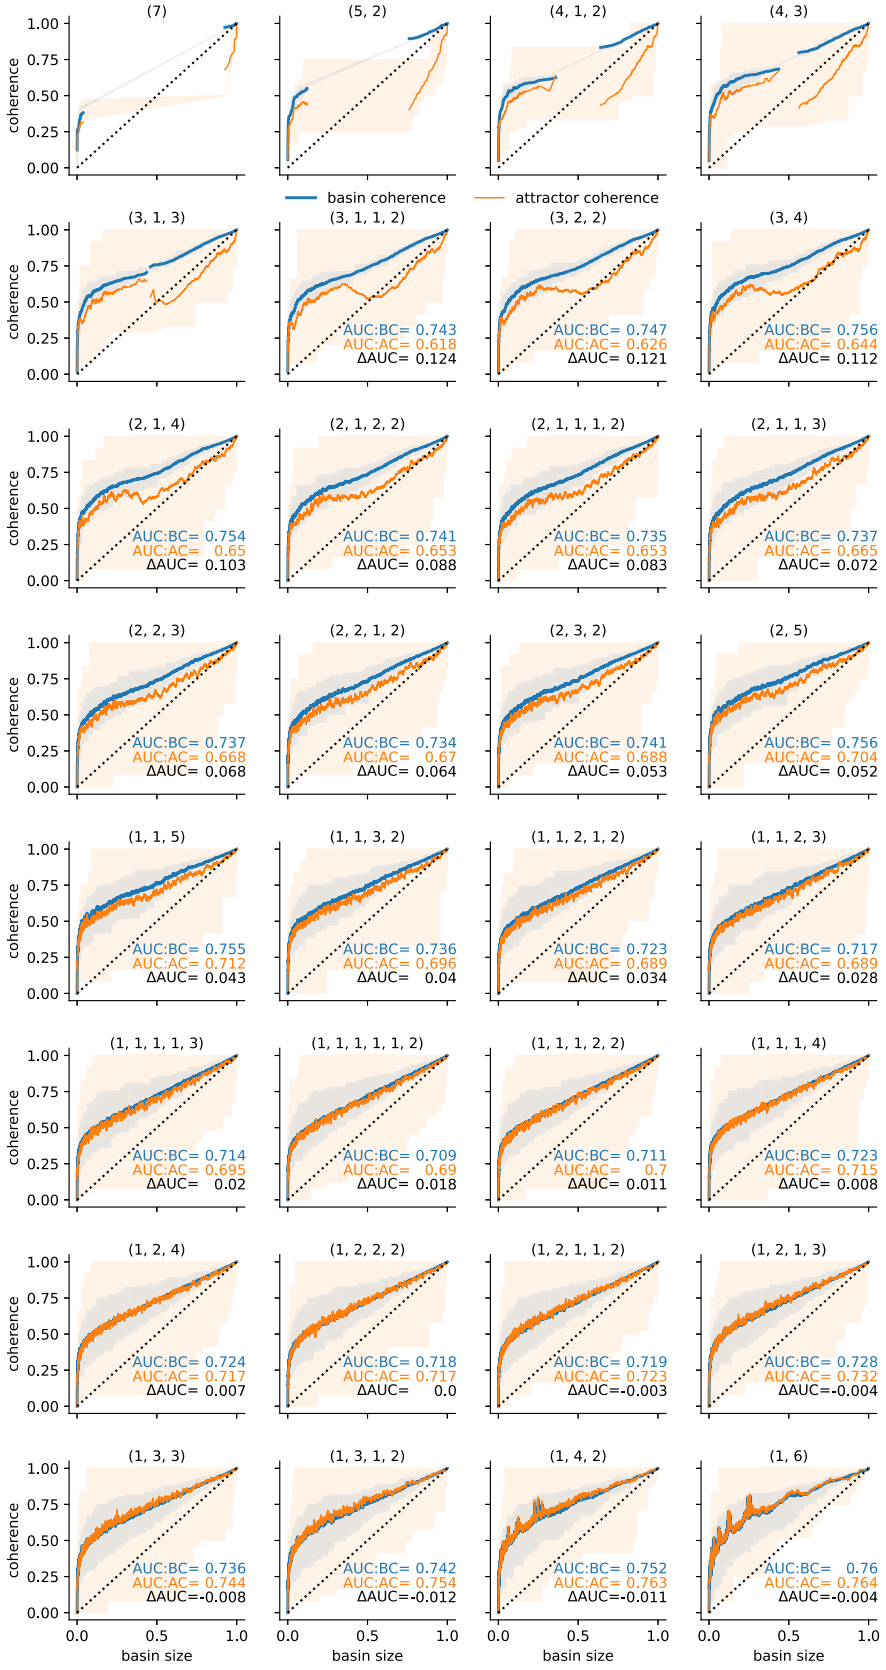


FIG. 15. Robustness of nested canalizing Boolean network attractors and their basins. For 10 000 random 12-node BNs governed by five-input nested canalizing update rules of different layer structure (specified above each panel), the coherence of each attractor (orange) and the coherence of its corresponding basin (blue) are stratified by basin size (x-axis). Shaded areas indicate the range of observed values. Lines depict rolling-window means with window size 50, and AUC is the area under the respective lines. The basin size distribution corresponding to each type of network is shown in Fig. 16.

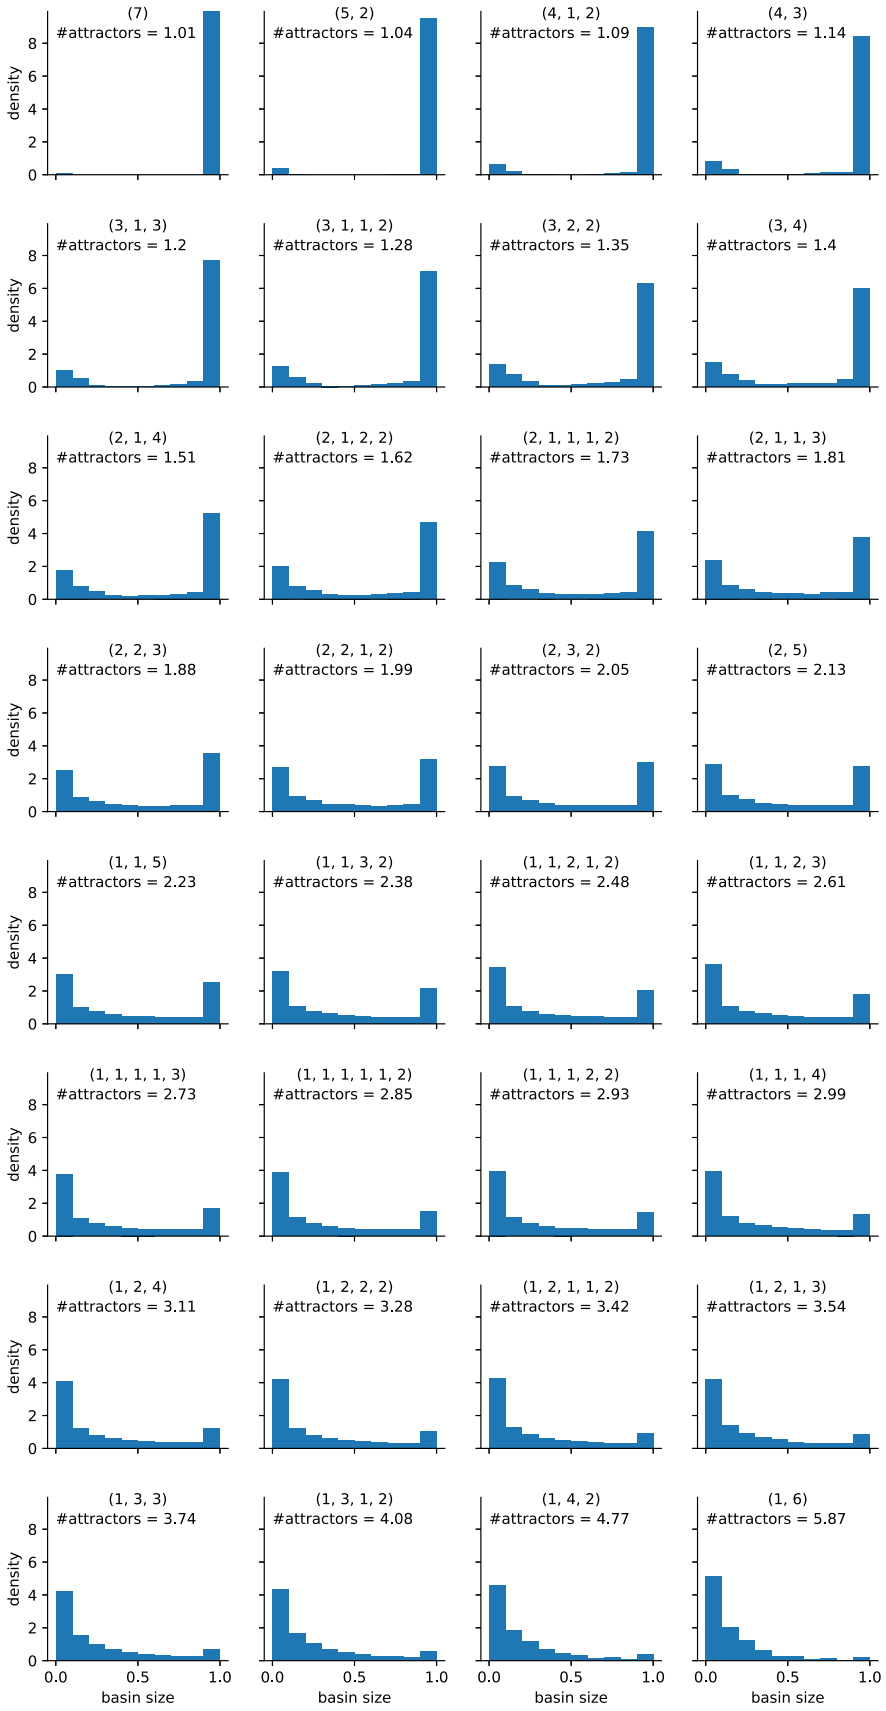


FIG. 16. Basin size distribution for different types of nested canalizing Boolean networks. For 10 000 random 12-node BNs governed by six-input nested canalizing update rules of different layer structure (specified above each panel), the distribution of the basin size is shown, in addition to the average number of network attractors.

would better capture biological stochasticity, the synchronous framework provides a well-defined baseline for formalizing attractor stability.

Moreover, developing a deeper theoretical understanding of how bias geometrically constrains attractor positions in state space remains an important direction. While our empirical results demonstrate the near-perfect predictive power of bias, a rigorous mathematical proof connecting bias to attractor boundary proximity would strengthen the mechanistic foundation.

G. A revised landscape metaphor

Waddington's landscape has endured as a conceptual framework precisely because it captures intuitions about developmental stability and canalization. Our findings refine rather than replace this metaphor (Fig. 1). The valleys are indeed deep, carved by canalization to ensure reliable development. However, because canalization also freezes large parts of the regulatory network, the remaining dynamical degrees of freedom are confined to a small active subspace. This effectively flattens the ridges near valley floors, creating a landscape where mature phenotypes are poised for potential transitions despite the protective valleys that brought them there.

This revised landscape provides a unifying framework for understanding how biological systems achieve both the stability needed for reliable function and the plasticity required for adaptation, wound healing, and evolution. The attractor-centric measure introduced here offers a path toward quantifying this balance in real biological systems, and predicting when and how phenotypic transitions will occur.

ACKNOWLEDGMENTS

The authors thank Maria Siskaki, Luis Sordo Vieira and Eric Deeds for useful comments and suggestions. MW was funded in part by the National Science Foundation, Grant No. DMS-2424635, and by the NIH, Grants No. R01 AI135128 and No. R01 HL169974-01. CK was funded in part by the Simons Foundation, Grant No. 712537, and the National Science Foundation, Grants No. DMS-2424632 and No. DMS-2451973.

DATA AVAILABILITY

The data that support the findings of this article are openly available [49].

[1] C. H. Waddington, *The Strategy of the Genes*, 1st ed. (Routledge, London, UK, 1957).

[2] S. Huang, G. Eichler, Y. Bar-Yam, and D. E. Ingber, Cell fates as high-dimensional attractor states of a complex gene regulatory network, Phys. Rev. Lett. **94**, 128701 (2005).

[3] S. Huang, The molecular and mathematical basis of Waddington's epigenetic landscape: A framework for post-Darwinian biology? BioEssays **34**, 149 (2012).

[4] S. Kauffman, Metabolic stability and epigenesis in randomly constructed genetic nets, J. Theor. Biol. **22**, 437 (1969).

[5] J. D. Schwab, S. D. Kühlwein, N. Ikonomi, M. Kühl, and H. A. Kestler, Concepts in Boolean network modeling: What do they all mean? Comput. Struct. Biotechnol. J. **18**, 571 (2020).

[6] B. Derrida and Y. Pomeau, Random networks of automata: A simple annealed approximation, Europhys. Lett. **1**, 45 (1986).

[7] E. Balleza, E. R. Alvarez-Buylla, A. Chaos, S. Kauffman, I. Shmulevich, and M. Aldana, Critical dynamics in genetic regulatory networks: Examples from four kingdoms, PLoS One **3**, e2456 (2008).

[8] B. C. Daniels, H. Kim, D. Moore, S. Zhou, H. B. Smith, B. Karas, S. A. Kauffman, and S. I. Walker, Criticality distinguishes the ensemble of biological regulatory networks, Phys. Rev. Lett. **121**, 138102 (2018).

[9] A. Roli, M. Villani, A. Filisetti, and R. Serra, Dynamical criticality: Overview and open questions, J. Syst. Sci. Complex. **31**, 647 (2018).

[10] C. Kadelka, T.-M. Butrie, E. Hilton, J. Kinseth, A. Schmidt, and H. Serdarevic, A meta-analysis of Boolean network models reveals design principles of gene regulatory networks, Sci. Adv. **10**, eadj0822 (2024).

[11] S. Squires, A. Pomerance, M. Girvan, and E. Ott, Stability of Boolean networks: The joint effects of topology and update rules, Phys. Rev. E **90**, 022814 (2014).

[12] C. Kadelka, J. Kuipers, and R. Laubenbacher, The influence of canalization on the robustness of Boolean networks, Physica D **353-354**, 39 (2017).

[13] C. Kadelka, Y. Li, J. Kuipers, J. O. Adeyeye, and R. Laubenbacher, Multistate nested canalizing functions and their networks, Theor. Comput. Sci. **675**, 1 (2017).

[14] K. Willadsen, J. Triesch, and J. Wiles, Understanding robustness in random Boolean networks, in *Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems (ALife XI)* (MIT Press, Winchester, UK, 2008).

[15] C. Kadelka, D. Murrugarra, and R. Laubenbacher, Stabilizing gene regulatory networks through feedforward loops, Chaos **23**, 025107 (2013).

[16] K. H. Park, F. X. Costa, L. M. Rocha, R. Albert, and J. C. Rozum, Models of cell processes are far from the edge of chaos, PRX Life **1**, 023009 (2023).

[17] C. Kadelka, M. Wheeler, A. Veliz-Cuba, D. Murrugarra, and R. Laubenbacher, Modularity of biological systems: A link between structure and function, J. R. Soc. Interface **20**, 20230505 (2023).

[18] S. E. Harris, B. K. Sawhill, A. Wuensche, and S. Kauffman, A model of transcriptional regulatory networks based on biases in the observed regulation rules, Complexity **7**, 23 (2002).

[19] C. Kadelka, Canalization as a stabilizing principle of gene regulatory networks: A discrete dynamical systems perspective, npj Syst. Biol. Appl. **12**, 30 (2026).

[20] I. Shmulevich, H. Lähdesmäki, E. R. Dougherty, J. Astola, and W. Zhang, The role of certain Post classes in Boolean network models of genetic networks, Proc. Natl. Acad. Sci. USA **100**, 10734 (2003).

[21] J. Davila-Velderrain, J. C. Martinez-Garcia, and E. R. Alvarez-Buylla, Modeling the epigenetic attractors landscape: Toward

a post-genomic mechanistic understanding of development, Front. Genet. **6**, 160 (2015).
[22] S. Huang and D. E. Ingber, A non-genetic basis for cancer progression and metastasis: Self-organizing attractors in cell regulatory networks, Breast Disease **26**, 27 (2007).
[23] C. Campbell and R. Albert, Stabilization of perturbed Boolean network attractors through compensatory interactions, BMC Syst. Biol. **8**, 53 (2014).
[24] V. S. N. Bavisetty, M. Wheeler, R. Laubenbacher, and C. Kadelka, Upper bound for the stability of Boolean networks, Phys. Rev. E **112**, 044310 (2025).
[25] B. Min, J. Choi, and R. Laubenbacher, Dynamics of attractor transitions in Boolean networks under noise, Chaos **36**, 023142 (2026).
[26] B. Derrida and H. Flyvbjerg, The random map model: A disordered model with deterministic dynamics, J. Phys. **48**, 971 (1987).
[27] S. Kauffman, C. Peterson, B. Samuelsson, and C. Troein, Genetic networks with canalyzing Boolean rules are always stable, Proc. Natl. Acad. Sci. USA **101**, 17102 (2004).
[28] S. Cook, C. Dwork, and R. Reischuk, Upper and lower time bounds for parallel random access machines without simultaneous writes, SIAM J. Comput. **15**, 87 (1986).
[29] I. Shmulevich and S. A. Kauffman, Activities and sensitivities in Boolean network models, Phys. Rev. Lett. **93**, 048701 (2004).
[30] S. A. Kauffman, *The Origins of Order: Self-organization and Selection in Evolution* (Oxford University Press, New York, USA, 1993).
[31] S. Kauffman, The large scale structure and dynamics of gene control circuits: An ensemble approach, J. Theor. Biol. **44**, 167 (1974).
[32] Q. He and M. Macauley, Stratification and enumeration of Boolean functions by canalizing depth, Physica D **314**, 1 (2016).
[33] E. Dimitrova, B. Stigler, C. Kadelka, and D. Murrugarra, Revealing the canalizing structure of Boolean functions: Algorithms and applications, Automatica **146**, 110630 (2022).
[34] C. O. Reichhardt and K. E. Bassler, Canalization and symmetry in Boolean models for genetic regulatory networks, J. Phys. A **40**, 4339 (2007).
[35] M. Marques-Pita and L. M. Rocha, Canalization and control in automata networks: Body segmentation in Drosophila melanogaster, PLoS One **8**, e55946 (2013).
[36] C. Kadelka, B. Keilty, and R. Laubenbacher, Collectively canalizing Boolean functions, Adv. Appl. Math. **145**, 102475 (2023).
[37] A. J. Gates, R. Brattig Correia, X. Wang, and L. M. Rocha, The effective graph reveals redundancy, canalization, and control pathways in biochemical regulation and signaling, Proc. Natl. Acad. Sci. USA **118**, e2022598118 (2021).
[38] C. Kadelka and B. Coberly, Boolforge Controlled generation and analysis of Boolean functions and networks, arXiv:2509.02496.
[39] C. Kadelka and K. Hari, Critical assessment of the ability of Boolean threshold models to describe gene regulatory network dynamics, PNAS Nexus **4**, pgaf228 (2025).
[40] S. Manicka, M. Marques-Pita, and L. M. Rocha, Effective connectivity determines the critical dynamics of biochemical networks, J. R. Soc. Interface **19**, 20210659 (2022).
[41] J. E. S. Socolar and S. A. Kauffman, Scaling in ordered and critical random Boolean networks, Phys. Rev. Lett. **90**, 068702 (2003).
[42] M. Möller and B. Drossel, The formation of the frozen core in critical Boolean networks, New J. Phys. **14**, 023051 (2012).
[43] U. Bastolla and G. Parisi, Relevant elements, magnetization and dynamical properties in Kauffman networks: A numerical study, Physica D **115**, 203 (1998).
[44] C. Kadelka and D. Murrugarra, Canalization reduces the nonlinearity of regulation in biological networks, npj Syst. Biol. Appl. **10**, 67 (2024).
[45] H. Kitano, Towards a theory of biological robustness, Mol. Syst. Biol. **3**, 137 (2007).
[46] H. Wang, Y. Yang, J. Liu, and L. Qian, Direct cell reprogramming: Approaches, mechanisms and progress, Nat. Rev. Mol. Cell Biol. **22**, 410 (2021).
[47] J. Cerneckis, H. Cai, and Y. Shi, Induced pluripotent stem cells (iPSCs): Molecular mechanisms of induction and applications, Sign. Transduct. Target. Ther. **9**, 112 (2024).
[48] A. Saadatpour, I. Albert, and R. Albert, Attractor analysis of asynchronous Boolean models of signal transduction networks, J. Theor. Biol. **266**, 641 (2010).
[49] V. S. N. Bavisetty, M. Wheeler, and C. Kadelka, Code repository, Zenodo (2026), https://doi.org/10.5281/zenodo.19700792.