ATTRACTOME: FROM THEORY TO GENERATIVE MODELS FOR CANCER DYNAMICS AND CONTROL

Arash Mehrjou
Max Planck Institute for Intelligent Systems & GSK.ai
arash@distantvantagepoint.com

ABSTRACT

Cancer progression couples intracellular gene-regulatory dynamics with eco-evolutionary selection, yet current approaches often model these axes separately. We introduce the Attractome, a unified framework that treats malignant phenotypes as attractors and their persistence via frequency-dependent selection. Effective barrier height between phenotypic basins links stochastic transition rates, mutational accessibility, and minimal control effort for reprogramming. We formalize a joint stability criterion: long-lived malignant phenotypes, including cancer stem cell states, must be both dynamically stable as attractors and evolutionarily stable under ecological feedback. We derive early-warning signatures of impending transitions, including critical slowing down, variance amplification, flickering, and spatial correlation, measurable in time-resolved single-cell and spatial genomic data. We outline a roadmap in which generative AI models learn quasi-potentials and barrier geometry from single-cell data for risk prediction and patient-adaptive intervention design.

1 INTRODUCTION

Tumor progression reflects a reiterative process of clonal expansion, genetic diversification, and selection within complex tissue ecosystems (Greaves & Maley, 2012; Gerlinger et al., 2012; Laplane & Maley, 2024). In parallel, cancer can be viewed as a disease of disturbed cell differentiation, wherein cells occupy abnormal stable states ("cancer attractors") in gene-regulatory landscapes (Kauffman & Fox, 1993; Huang et al., 2009; Ao et al., 2008; Moris et al., 2016). These perspectives, often studied separately, are fundamentally interlinked. Oncogenic mutations rewire regulatory circuits and reshape basin geometry of the underlying dynamical system, making malignant states more accessible; conversely, phenotypic states modulate effective mutation rates, stress responses, and selection pressures, altering evolutionary trajectories (Feinberg & Levchenko, 2023; Shaffer et al., 2017; Nowak et al., 2002; McFarland et al., 2017). Figure 1 provides a schematic depiction of how mutation-driven remodeling alters barrier height between healthy and malignant basins and thereby changes both autonomous transitions and intervention requirements.

We develop a unified eco-evolutionary dynamical framework that couples intracellular state dynamics with population composition. In this view, attractor geometry and frequency-dependent fitness co-determine initiation, persistence, and therapeutic response. A central construct is barrier height between phenotypic basins, which governs (1) noise-driven transition rates, (2) evolutionary accessibility via mutational remodeling, and (3) minimal control effort required for therapeutic escape. We show that Cancer Stem Cell (CSC)-like phenotypes exhibit joint dynamical-evolutionary stability and that durable control therefore requires coordinated disruption of both axes. We further derive early-warning signatures of approaching critical transitions in cell states and discuss dual control strategies that combine landscape engineering with evolutionary steering. To formalize basin removal for landscape engineering, we define the Minimal Sustaining Set (MSS) as the smallest collection of regulatory constraints whose removal eliminates a malignant attractor; MSS provides a principled target set for minimal circuit intervention.

In the context of generative AI for genomics, the Attractome provides a structured target space for *learning* and *sampling* cell-state dynamics from high-dimensional single-cell measurements. In particular, generative models can be used to infer low-dimensional landscape surrogates (quasi-

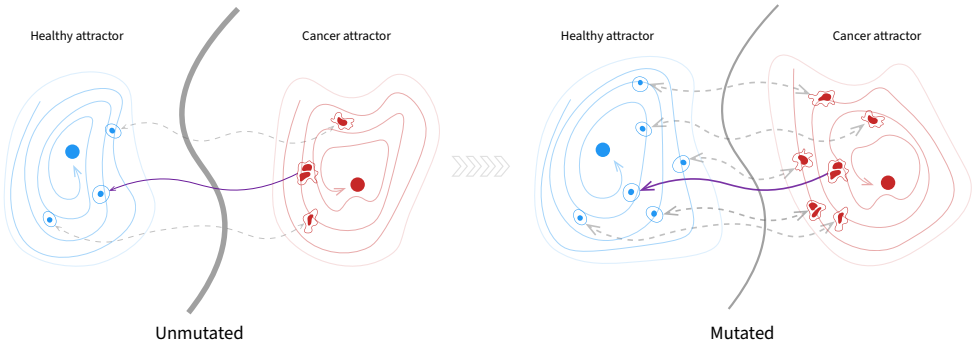

Figure 1: Basin geometry and mutation-driven remodeling. The left panel shows the unmutated state where healthy and cancer attractor basins are well separated. The barrier ridge between them is gray and high, denoted ΔE_{high}, so autonomous transitions are rare. Healthy basin is shown in blue and cancer basin in red. Dashed gray arrows depict noise-driven transitions, and the weak purple arrow indicates an intervention-induced transition (e.g., drug therapy) that requires substantial effort when the barrier is high. The right panel shows the mutated state where basins lie closer and the gray barrier is lower, denoted ΔE_{low} with $\Delta E_{\text{low}} \ll \Delta E_{\text{high}}$. Autonomous transitions become more frequent, with dashed gray arrows more pronounced, and the stronger purple intervention arrow indicates that therapeutic interventions require less effort to drive cells from the cancer attractor to the healthy attractor when the energy barrier is lower. CSC denotes a cancer stem cell, a stem-like malignant subpopulation with self-renewal and the capacity to generate differentiated progeny.

potentials), estimate barrier heights between phenotypic basins, and simulate counterfactual transitions under perturbations. This makes the framework directly actionable for time-resolved single-cell RNA-seq/ATAC-seq(Wang et al., 2009; Buenrostro et al., 2013), perturb-seq(Dixit et al., 2016), and spatial transcriptomics(Marx, 2021), where the goal is to anticipate malignant commitment and design interventions conditioned on patient-specific data.

1.1 Related work and positioning

The attractor perspective on cancer draws on epigenetic landscape metaphors and multistability analyses (Ferrell, 2012; Thomas, 1973), while evolutionary game-theoretic models clarify ecological interactions among tumor phenotypes (Weibull, 1997; Coggan & Page, 2022; Gatenby & Brown, 2020). We connect these traditions by treating genotype-driven parameter remodeling as geometric changes in basin topology while allowing phenotypic state to feed back on mutation and selection. Early-warning signatures align with tipping-point literature (Scheffer et al., 2009; Dakos et al., 2012; Kuehn, 2011) and recent efforts to quantify phenotypic plasticity (Burkhardt et al., 2022). We contribute a unified eco-evolutionary framework that provides a simple theoretical foundation for explaining diverse cancer phenomena, including initiation risk, persistent phenotypes, relapse, and frequency-dependent therapy response, through a common geometric language. Central to this unification is barrier height, which emerges as the unifying quantity governing stochastic transitions, mutational accessibility, and control effort across these different scales.

2 Unified eco-evolutionary dynamics

2.1 Intracellular dynamics and gene-regulatory landscapes

Consider a cell with genotype G and intracellular state vector $\boldsymbol{x}(t)$. Intracellular dynamics are governed by gene-regulatory network (GRN) dynamics

$$\frac{d\boldsymbol{x}}{dt} = F(\boldsymbol{x}; G), \tag{1}$$

where attractors $\boldsymbol{x}^*$ satisfy $F(\boldsymbol{x}^*; G) = \mathbf{0}$ and are asymptotically stable when the Jacobian spectrum has negative real parts (Khalil & Grizzle, 2002; Ferrell, 2012; Thomas, 1973; Zañudo & Albert,

2015). For a fixed genotype G, this defines a fixed landscape with stable attractors corresponding to distinct phenotypic states. Multistable landscapes admit multiple attractors separated by barriers, where cells can transition between attractors through noise-driven dynamics or external perturbations.

2.2 Evolutionary dynamics and landscape remodeling

Evolutionary processes modify the landscape by changing genotype-dependent parameters. Consider a population comprised of cells indexed by i, each with genotype G_i and intracellular state vector $\boldsymbol{x}_i(t)$. The population composition over genotypes changes via state-dependent birth-death processes and mutation (where the index i is dropped in the following equation for convenience):

$$\frac{dn_G}{dt} = n_G \left[b_G(\boldsymbol{x}) - d_G(\boldsymbol{x})\right] + \sum_{G'} n_{G'}\, \mu_{G' \to G}(\boldsymbol{x}), \tag{2}$$

where n_G is the number of cells with genotype G, $b_G(\boldsymbol{x})$ and $d_G(\boldsymbol{x})$ denote birth and death rates that depend on phenotypic state, and $\mu_{G' \to G}(\boldsymbol{x})$ denotes mutation rates, which can increase in stressed or unstable states (Galhardo et al., 2007; Labrie et al., 2022).

Evolutionary game theory(Weibull, 1997) provides the population-level criterion for ecological stability. Let phenotypic strategies be indexed by i with population composition $p = (p_1, \ldots, p_m)$ and frequency-dependent payoffs $\pi_i(p; \eta)$ that can depend on microenvironmental support η (e.g. resources, niche factors, etc). The population dynamics are captured by the replicator equation

$$\frac{dp_i}{dt} = p_i\big(\pi_i(p;\eta) - \bar{\pi}(p;\eta)\big), \quad \bar{\pi}(p;\eta) = \sum_j p_j\, \pi_j(p;\eta), \tag{3}$$

which formalizes selection among phenotypes with interactions (Weibull, 1997; Coggan & Page, 2022). Mutational inputs can be incorporated via replicator-mutator extensions, linking ecological selection with genetic diversification.

Evolutionary dynamics drive several key phenomena. As a few examples, *Mutation-driven landscape remodeling* arises when oncogenic alterations change parameters in F, creating or stabilizing malignant attractors and lowering barriers between basins (Shin & Cho, 2023). *State-dependent evolvability* occurs because cells in particular phenotypic states (e.g., highly proliferative or stressed) experience elevated mutation rates and altered selection, changing the speed and direction of evolution. *Dynamic fitness landscapes* emerge because birth-death rates depend on both genotype and phenotypic state, leading to frequency-dependent fitness and ecological feedbacks (Diaz-Uriarte & Vasallo, 2019).

To characterize which phenotypes persist at the population level under frequency-dependent selection, we recall here the standard definition of evolutionary stability, as it will be needed in what follows, complementing the dynamical stability of attractors that explains long-term persistence of malignant phenotypes.

Definition 1 (Evolutionarily stable strategy (ESS))**.** *A composition p^* is evolutionarily stable if, for all nearby mutants q, small invasions reduce mutant fitness, meaning $\pi(p^*; (1-\epsilon)p^* + \epsilon q) \geq \pi(q; (1-\epsilon)p^* + \epsilon q)$ for all sufficiently small $\epsilon > 0$. ESS implies local invadability resistance under frequency-dependent selection.*

2.3 Coupled dynamics and joint stability

The coupling between intracellular and evolutionary dynamics creates a unified system where genotype-dependent parameters shape the landscape while phenotypic states feed back on mutation and selection. This coupling enables *non-genetic phenotypic switching*, transitions between attractors without genetic change (e.g., drug-sensitive to drug-tolerant states), enabling rapid plastic responses to therapy (Kemper et al., 2014; Peyre et al., 2025). Barrier geometry affects invadability indirectly by altering phenotypic payoffs through state-dependent proliferation and death, and by modulating mutational accessibility; thus landscape engineering and evolutionary steering jointly manipulate payoffs and basin structure (Mehrjou et al., 2021).

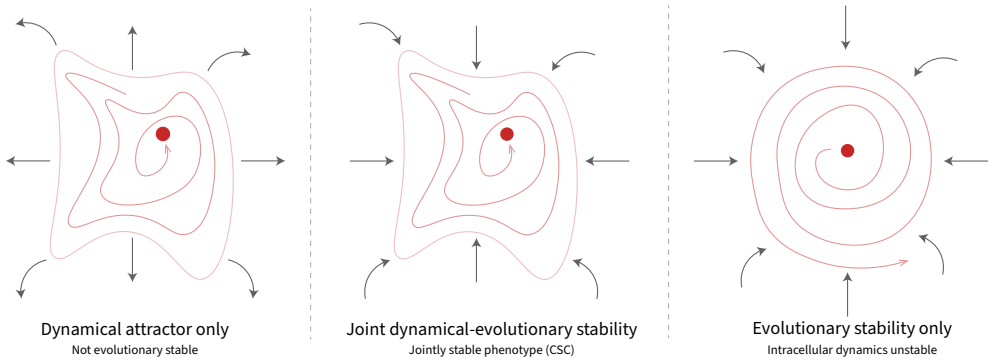

Figure 2: Joint dynamical and evolutionary stability. The middle panel shows a CSC-like state that satisfies both dynamical attractor stability and ESS stability, enabling long-term persistence. The left panel shows a dynamical attractor that is not an ESS, which is outcompeted over evolutionary time. The right panel shows an ESS that is not an attractor, which is lost to phenotypic drift.

Persistent malignant phenotypes, with cancer stem cell CSC-like states as canonical examples, must satisfy a dual stability criterion: they must be *dynamical attractors* of intracellular regulatory networks (with nonempty basins and asymptotic stability) while simultaneously implementing *evolutionarily stable strategies* (ESS) under frequency-dependent selection (Weibull, 1997; Coggan & Page, 2022; Gatenby & Brown, 2020). Attractors with insufficient ecological fitness are displaced over evolutionary timescales; ESS-like phenotypes lacking dynamical basins are lost to phenotypic drift. This biconditional requirement explains persistence and relapse, as residual CSC-like subpopulations re-expand after therapy because they remain stable on both axes.

Proposition 1 (Joint dynamical and evolutionary stability)**.** *A malignant phenotype persists over long timescales if and only if it is both (1) a dynamical attractor with a nonempty basin and asymptotic stability, and (2) an evolutionarily stable strategy under frequency-dependent selection. States satisfying only one of these conditions are transient and cannot sustain long-term tumor persistence. See Appendix B for detailed definitions and a proof sketch.*

Corollary 1 (Control complementarity)**.** *For phenotypes satisfying the above proposition, landscape-only control or evolutionary-only control is insufficient for eradication; durable control requires joint manipulation of basin geometry (to remove or destabilize the malignant attractor) and ecological payoffs (to break ESS invadability).*

Figure 2 illustrates the joint condition. In the middle panel, convergent state-space trajectories and inward population arrows indicate both dynamical and evolutionary stability, a persistent CSC-like state. In the left panel, trajectories converge (dynamical stability) but population arrows point outward (ecological instability), yielding transient states outcompeted in the long run. In the right panel, population arrows point inward (ecological stability) but trajectories diverge (no attractor), leading to eventual loss through phenotypic drift.

3 Barrier height as a common currency

Multistable gene-regulatory dynamical systems with some standard regulatory assumptions admit a quasi-potential $V(\boldsymbol{x})$ (Freidlin et al., 1970; Brackston et al., 2018) in which attractors correspond to local minima and saddle points represent transition states between basins. We define the *effective barrier height* between phenotypes as

$$\Delta E := \min_{\gamma:\boldsymbol{x}_A\to\boldsymbol{x}_B} \max_{\boldsymbol{x}\in\gamma} \left[V(\boldsymbol{x}) - V(\boldsymbol{x}_A)\right],$$

which measures the minimal mountain-pass height separating basins. We show that this single geometric quantity unifies three scales of tumor dynamics.

Molecular scale (stochastic transitions). Kramers-type escape dynamics imply that the noise-driven transition rate between basins scales as rate $\propto \exp(-\Delta E/\sigma^2)$, where σ is the effective noise mag-

nitude (Hänggi et al., 1990; Fiasconaro et al., 2010). Doubling ΔE suppresses spontaneous phenotypic switching exponentially; halving ΔE dramatically accelerates transitions.

Evolutionary scale (mutational accessibility). Oncogenic mutations reshape the landscape by modifying regulatory parameters, thereby lowering barriers to malignant attractors (Nowak et al., 2002; Komarova & Wodarz, 2004). When ΔE decreases along mutational trajectories, the probability of reaching malignant basins increases exponentially via the same escape mechanism, establishing a mechanistic link between genotype and phenotype-level transition dynamics.

Therapeutic scale (control effort). Optimal control considerations imply that the minimal intervention energy required to drive a system out of a malignant basin grows with ΔE; conversely, *barrier engineering* (e.g. reinforcing differentiation programs, restoring tumor suppressor function) reduces the control effort needed for successful reprogramming (Bardi & Capuzzo-Dolcetta, 2008). Control policies should therefore prioritize parameters that maximally increase ΔE between healthy and malignant basins and target Minimal Sustaining Set (MSS) members for attractor elimination.

Definition 2 (Minimal Sustaining Set (MSS))**.** *A minimal sustaining set for a malignant attractor is the smallest set of regulatory constraints (edges/parameters) whose removal eliminates that attractor while any strict subset does not.*

The three scales discussed above (molecular stochastic transitions, evolutionary mutational accessibility, and therapeutic control effort) are unified by a single geometric quantity. The effective barrier height ΔE provides a common currency that simultaneously governs all three phenomena, enabling predictions across scales from a single landscape measurement.

Proposition 2 (Barrier height as common currency)**.** *The effective barrier height ΔE separating phenotypic attractors simultaneously governs (1) noise-driven transition rates via Kramers-type escape scaling, (2) evolutionary accessibility through mutational remodeling of regulatory parameters that lower ΔE, and (3) the minimal control effort required to exit malignant basins. A single estimate of ΔE thus predicts spontaneous transition rates, invadability, and intervention requirements.*

Sketch of proof. The unification follows from recognizing that ΔE quantifies the minimal work required to traverse the mountain pass between basins. In stochastic dynamics, this work determines escape rates via large deviation principles. Mutations that lower ΔE increase accessibility through the same escape mechanism. For optimal control, the minimum-energy path crossing the barrier provides a lower bound on intervention effort. The common dependence on ΔE establishes the unified scaling across all three scales. Rigorous statements and detailed proofs are provided in Appendix B.

In the next section, we show how the proposed theoretical framework can lead to actionable biomarkers for early diagnosis and risk stratification.

4 Early-warning signatures of impending malignant transitions

A key practical consequence of the Attractome framework is the identification of measurable early-warning signatures that precede malignant transitions. These signatures emerge directly from the geometric structure of the landscape and the dynamical properties near basin boundaries, providing actionable biomarkers for early diagnosis, risk stratification, and intervention timing. By detecting these signatures in time-resolved single-cell and spatial genomics data, clinicians can anticipate malignant commitment before irreversible phenotypic changes occur, enabling proactive therapeutic intervention.

Systems approaching basin boundaries and bifurcations exhibit universal signatures(Lade & Gross, 2012) that provide anticipatory biomarkers for malignant commitment. These signatures arise from the geometry of the underlying landscape and the linearized dynamics near attractors.

Critical slowing down. Let $d\boldsymbol{x}/dt = F(\boldsymbol{x};\mu)$ and consider a stable attractor $\boldsymbol{x}^*(\mu)$ undergoing a bifurcation at $\mu = \mu_c$. The dominant eigenvalue of the Jacobian, $\lambda_{\min}(\mu)$, approaches zero as the bifurcation is approached, and the relaxation time (the characteristic time it takes a system to return to its stable state after a small perturbation) $\tau(\mu) = 1/|\lambda_{\min}(\mu)|$ diverges (Dakos et al., 2012; Scheffer et al., 2009). Empirically, increased temporal autocorrelation in stemness markers (quantifying the activation level of self-renewal and undifferentiated cellular programs) and delayed

recovery from perturbations signal proximity to a malignant basin boundary. See Theorem 1 in Appendix B for the formal statement and proof.

Variance amplification and flickering. In stochastic settings,

$$d\boldsymbol{x} = F(\boldsymbol{x}; \mu)\, dt + \sigma\, dW_t,$$

the stationary variance along the slowest mode scales as $\mathrm{Var} \sim \sigma^2/(2|\lambda_{\min}|)$, producing measurable increases in cell-to-cell variability (Chen et al., 2012; Fiasconaro et al., 2010). As barriers lower relative to noise, transient excursions into alternative basins ("flickering") become more frequent, with rates governed by $\exp(-\Delta E/\sigma^2)$. Formal statements and proofs are provided in Theorems 2 and 3 in Appendix B.

Spatial correlation and heterogeneity. In spatially extended populations, reaction-diffusion coupling yields increasing correlation length ξ near criticality, reflecting coordinated fluctuations across neighborhoods (Kuehn, 2011; Burkhardt et al., 2022). Bimodal or heavy-tailed distributions of stemness scores and expanding spatial correlation structures provide population-level indicators of landscape flattening and impending transitions. The formal statement and proof are provided in Theorem 4 in Appendix B.

Together, these signatures enable anticipatory detection of CSC emergence in experimental settings that resolve cellular states at single-cell resolution across time and space, thereby permitting estimation of temporal autocorrelation, recovery dynamics, and spatial correlation structure. These conditions can be satisfied, for example, by time-resolved single-cell assays, live-cell imaging with regulatory circuit reporters, and spatial statistical analysis applied to multiplexed molecular measurements. Beyond providing an interpretive lens for early detection, a central implication of this theoretical framework is that it motivates the targeted collection of specific data modalities by clarifying their diagnostic value and the dynamical quantities they are expected to reveal, rather than treating such measurements as purely descriptive.

5 Dual control strategies for landscape engineering and evolutionary steering

Durable control requires coordinated manipulation of intracellular dynamics and population-level selection. Landscape engineering *directly* acts on the right-hand side of intracellular dynamics (Equation 1), while evolutionary steering *indirectly* changes ecological payoffs and frequency-dependent fitness (Equation 3).

Landscape engineering. Therapeutic interventions can be understood as reshaping the regulatory landscape to control transitions between cellular states. Differentiation-promoting interventions and barrier-reinforcing strategies increase the effective barrier height ΔE surrounding healthy basins by weakening sustaining feedbacks of malignant states (e.g., self-renewal circuits) and restoring tumor suppressor pathways. Such interventions suppress autonomous phenotypic switching and reduce the accessibility of malignant attractors, but do not necessarily eliminate them.

Evolutionary steering. While landscape engineering targets the stability and accessibility of phenotypic states, it does not directly address population-level selection dynamics, which can re-establish malignant phenotypes through adaptive evolution. Evolutionary steering strategies are therefore required to reshape competitive interactions and selection pressures within the tumor ecosystem. Adaptive therapy exploits fitness costs of resistance to maintain a competitively suppressive population of drug-sensitive cells (Gatenby et al., 2009). For example, evolutionary double-binds (traps) (Basanta et al., 2012) deliberately select for phenotypic traits that confer short-term survival advantages but create liabilities exploited by subsequent interventions, while collateral sensitivity sequences leverage pleiotropic costs such that adaptation to one agent increases sensitivity to another (Gatenby & Brown, 2020).

Complementarity. The Attractome framework predicts that robust eradication requires the joint application of landscape engineering and evolutionary steering, as each modality compensates for the structural limitations of the other. Landscape-only control can transiently destabilize malignant basins but leaves ESS-like population equilibria intact, permitting re-invasion through selection. Conversely, evolutionary-only control can disrupt ESS configurations while leaving the phenotypic landscape structurally unchanged, enabling drift or noise-driven return to malignancy. Combined

policies create a coupled dynamical–evolutionary constraint (a "double barrier") that suppresses both state-space accessibility and evolutionary invasibility, yielding a control regime that neither mechanism can achieve alone.

6 ILLUSTRATIVE MINIMAL MODEL

We illustrate the Attractome framework using a parsimonious three-dimensional dynamical system with state variables S (stemness activity), D (differentiation activity), and C (control/integrator state). These abstract variables represent projections or aggregations of higher-dimensional gene expression states onto a low-dimensional manifold that captures essential multistability, where S and D correspond to collective activation levels of gene modules associated with self-renewal and differentiation programs, respectively. The system exhibits three interpretable attractors: differentiated, regulated stem, and uncontrolled stem/CSC-like states. The dynamics follow

$$\begin{aligned}\frac{dS}{dt} &= a_s^{\text{eff}}(S - S^3) - b_{sd}D + k_c^{\text{eff}} \cdot \text{control\_on} \cdot C \\ \frac{dD}{dt} &= a_d(D - D^3) - b_{ds}S \\ \frac{dC}{dt} &= c_{\text{gain}} \cdot \text{control\_on} \cdot (s_{\text{target}} - S) - c_{\text{decay}}C\end{aligned}$$

with mutation-adjusted parameters

$$\begin{aligned}a_s^{\text{eff}} &= a_s(1 + m \cdot \mu_{as}), \\ k_c^{\text{eff}} &= k_c(1 - m \cdot \mu_{kc}),\end{aligned}$$

where $m \in [0, 1]$ is mutation strength, and control_on gates homeostatic feedback when external or endogenous control is active.

Self-activation with saturation is captured by $(x - x^3)$, which reinforces small activities and saturates at high levels, creating stable states near $x = \pm 1$ for stemness and differentiation programs. Cross-inhibition terms $-b_{sd}D$ and $-b_{ds}S$ implement mutual repression that enforces fate commitment. The integrator C accumulates stemness error $s_{\text{target}} - S$ and feeds back with gain k_c^{eff} when control is on, representing homeostatic regulation that maintains moderate stemness.

The stemness score used to classify states is $\text{stemness}(\boldsymbol{x}) = 0.8 \cdot \sigma(2.5S) + 0.2 \cdot \sigma(-2.5D)$, where $\sigma(z) = 1/(1 + e^{-z})$, which emphasizes high S while providing a small bonus for low D to yield a continuous measure that separates CSC-like attractors from others.

Mutation effects increase a_s^{eff} and weaken k_c^{eff}, reflecting strengthened self-renewal and compromised control. These changes enlarge and deepen the CSC-like basin and lower barriers, increasing autonomous transitions and reducing intervention effort needed for escape. Moreover, mutations shift the system equilibria toward higher stemness values: as mutation strength m increases, the CSC-like equilibrium moves to higher S^* values, and the corresponding stemness score increases. A detailed analysis of this equilibrium shift for the reduced dynamical system is provided in Appendix B.

Mutation-driven changes in regulatory parameters produce different phenotypic states corresponding to distinct evolutionary strategies: higher mutation levels ($m \to 1$) yield CSC-like states with high stemness implementing aggressive self-renewal, while lower mutation levels favor differentiated states with proliferation-limited strategies. Under replicator dynamics (Equation equation 3), the payoff $\pi(s; p, \eta) = \eta \cdot s - \beta \cdot (1 - s) - \alpha \cdot p_s$ depends on stemness score s, microenvironmental support η, frequency p_s of similar cells, differentiation costs β, and competitive interactions α. CSC-like strategies become evolutionarily stable when η is sufficiently high and p_s is moderate, but invadable when η is low or p_s is high. The formal necessary and sufficient conditions for ESS stability and invadability are provided in Theorem 5 in Appendix B. This demonstrates the joint dynamical-evolutionary stability requirement: CSC-like attractors persist only when both dynamically stable (nonempty basins) and evolutionarily stable (resistant to invasion). Minimal control energy to exit the CSC basin increases with barrier height, supporting the common-currency view and informing dosing schedules and pulse timing under physiologic constraints. Additional numerical illustrations of this minimal GRN, including the bifurcation structure under mutation and the MSS-based intervention analysis, are provided in Appendix A (Figs. 4 and 5).

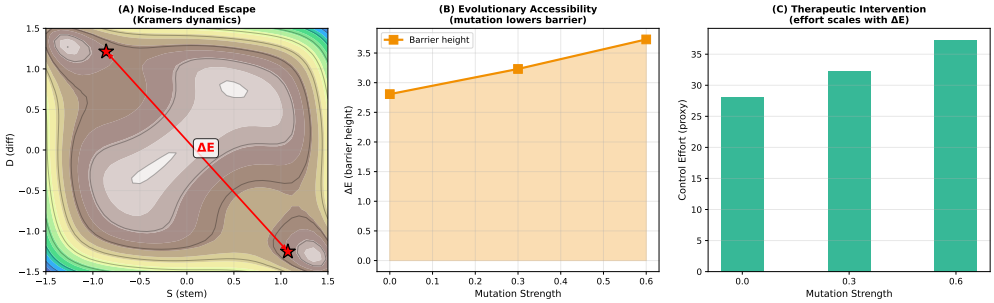

Figure 3: Barrier height ΔE as a common currency across scales. The left panel depicts Kramers escape over a barrier, where transition rates scale as $\exp(-\Delta E/\sigma^2)$. The center panel shows that mutations lower ΔE, increasing evolutionary accessibility to malignant basins. The right panel shows that minimal control effort to exit malignant basins grows with ΔE.

7 Generative AI roadmap for Attractome modeling

Generative AI provides the key mechanism for translating the theoretical Attractome framework into data-driven algorithms. Score-based/diffusion models(Song et al., 2020), normalizing flows(Rezende & Mohamed, 2015), or energy-based models(LeCun et al., 2006) learn landscape surrogates from single-cell measurements, with training informed by Attractome regularities (multistability constraints, barrier smoothness, quasi-potential structure) to capture the required geometric properties (Mehrjou et al., 2023). A practical roadmap proceeds in three stages: Stage 1 calibrates reduced regulatory circuit models from single-cell data, learning quasi-potentials encoding barrier geometry (Ferrell, 2012; Zañudo & Albert, 2015; Hänggi et al., 1990; Fiasconaro et al., 2010) and computing MSS via perturbation/sensitivity analysis; Stage 2 embeds learned dynamics within eco-evolutionary population models using replicator equations (Weibull, 1997; Gatenby et al., 2009), characterizes ESS regimes, and evaluates adaptive protocols using barrier height to budget control effort (Bardi & Capuzzo-Dolcetta, 2008); Stage 3 implements closed-loop deployment with generative models continuously updating patient-specific landscape estimates from time-resolved assays measuring early-warning signatures (Dakos et al., 2012; Chen et al., 2012; Kuehn, 2011; Burkhardt et al., 2022), triggering therapy adjustments before malignant commitment. A detailed pseudocode showing the complete pipeline is provided in Appendix B.

8 Discussion and conclusion

The Attractome abstracts complex molecular and ecological realities into reduced models and geometric surrogates to expose core control and detection principles. Its practical value lies in integrating scales, providing measurable signatures (autocorrelation, variance, flickering, spatial patterns), and unifying design rules (raise ΔE, target minimal sustaining sets, exploit resistance costs). Several limitations warrant emphasis. First, real tumors exhibit high-dimensional network dynamics and spatial structure; reduced models are surrogates that require calibration. Second, parameter uncertainty and context dependence challenge the inference of landscape geometry from data. Third, clinical constraints (toxicity, dosing windows, patient heterogeneity) restrict admissible controls, motivating robust and adaptive policies. Addressing these limitations requires longitudinal single-cell perturbation assays, patient-specific landscape reconstruction, and closed-loop therapy design integrating ecological and dynamical levers.

Cancer is an eco-evolutionary dynamical process. Barrier geometry provides a common currency linking stochastic transitions, mutational accessibility, and control effort. Joint dynamical and evolutionary stability clarifies persistence and relapse. Early-warning signatures enable anticipatory therapy. Dual control strategies that engineer landscapes and steer evolution are necessary for durable eradication. This unified framework offers practical guidance for detection and therapy design across tumor contexts and motivates patient-specific, data-driven implementations.

ACKNOWLEDGMENTS

We thank Nick Wreglesworth for valuable discussions and oncological perspectives that strengthened this work.

REFERENCES

Ping Ao, David Galas, Leroy Hood, and Xiaomei Zhu. Cancer as robust intrinsic state of endogenous molecular-cellular network shaped by evolution. *Medical Hypotheses*, 70(3):678–684, 2008.

Martino Bardi and Italo Capuzzo-Dolcetta. *Optimal Control and Viscosity Solutions of Hamilton–Jacobi–Bellman Equations*. Birkhäuser Boston, 2008.

David Basanta, Robert A Gatenby, and Alexander RA Anderson. Exploiting evolution to treat drug resistance: combination therapy and the double bind. *Molecular pharmaceutics*, 9(4):914–921, 2012.

Rowan D Brackston, Andrew Wynn, and Michael PH Stumpf. Construction of quasipotentials for stochastic dynamical systems: An optimization approach. *Physical Review E*, 98(2):022136, 2018.

Jason D Buenrostro, Paul G Giresi, Lisa C Zaba, Howard Y Chang, and William J Greenleaf. Transposition of native chromatin for fast and sensitive epigenomic profiling of open chromatin, dna-binding proteins and nucleosome position. *Nature methods*, 10(12):1213–1218, 2013.

Daniel B Burkhardt, Beatriz P San Juan, John G Lock, Smita Krishnaswamy, and Christine L Chaffer. Mapping phenotypic plasticity upon the cancer cell state landscape using manifold learning. *Cancer Discovery*, 12(8):1847–1859, 2022.

Li Chen, Rui Liu, Zhi-Ping Liu, M Li, and Kazuyuki Aihara. Detecting early-warning signals for critical transitions in complex diseases. *Scientific Reports*, 2:342, 2012.

Helena Coggan and Karen M Page. The role of evolutionary game theory in spatial and non-spatial models of the survival of cooperation in cancer: a review. *Journal of the Royal Society Interface*, 19(193):20220346, 2022.

Vasilis Dakos, Stephen R Carpenter, William A Brock, Aaron M Ellison, Vishwesha Guttal, Anthony R Ives, Sonia Kéfi, Valerie Livina, David A Seekell, Egbert H van Nes, et al. Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. *PLoS One*, 7(7):e41010, 2012.

Ramon Diaz-Uriarte and Claudia Vasallo. Every which way? on predicting tumor evolution using cancer progression models. *PLoS Computational Biology*, 15(8):e1007246, 2019.

Atray Dixit, Oren Parnas, Biyu Li, Jenny Chen, Charles P Fulco, Livnat Jerby-Arnon, Nemanja D Marjanovic, Danielle Dionne, Tyler Burks, Raktima Raychowdhury, et al. Perturb-seq: dissecting molecular circuits with scalable single-cell rna profiling of pooled genetic screens. *cell*, 167(7): 1853–1866, 2016.

Andrew P Feinberg and Andre Levchenko. Epigenetics as a mediator of plasticity in cancer. *Science*, 379(6632):eaaw3835, 2023.

James E Ferrell. Bistability, bifurcations, and waddington's epigenetic landscape. *Current Biology*, 22(11):R458–R466, 2012.

Alessandro Fiasconaro, Juan J Mazo, and Bernardo Spagnolo. Noise-induced enhancement of stability in a metastable system with damping. *Physical Review E*, 82(4):041120, 2010.

Mark I Freidlin et al. On small random perturbations of dynamical systems. *Russian Mathematical Surveys*, 25(1):1, 1970.

Rodrigo S Galhardo, Philip J Hastings, and Susan M Rosenberg. Mutation as a stress response and the regulation of evolvability. *Critical Reviews in Biochemistry and Molecular Biology*, 42(5): 399–435, 2007.

Robert A Gatenby and Joel S Brown. Integrating evolutionary dynamics into cancer therapy. *Nature Reviews Clinical Oncology*, 17(11):675–686, 2020.

Robert A Gatenby, Ariosto S Silva, Robert J Gillies, and B Roy Frieden. Adaptive therapy. *Cancer Research*, 69(11):4894–4903, 2009.

Marco Gerlinger, Andrew J Rowan, Stuart Horswell, James Larkin, David Endesfelder, Eva Gronroos, Pierre Martinez, Nicholas Matthews, Aengus Stewart, Patrick Tarpey, et al. Intratumor heterogeneity and branched evolution revealed by multiregion sequencing. *New England Journal of Medicine*, 366(10):883–892, 2012.

M. Greaves and C. C. Maley. Clonal evolution in cancer. *Nature*, 481(7381):306–313, 2012.

Peter Hänggi, Peter Talkner, and Michal Borkovec. Reaction-rate theory: fifty years after kramers. *Reviews of Modern Physics*, 62(2):251–341, 1990.

Sui Huang, Ingemar Ernberg, and Stuart Kauffman. Cancer attractors: a systems view of tumors from a gene network dynamics and developmental perspective. *Seminars in Cell & Developmental Biology*, 20(7):869–876, 2009.

Stuart Kauffman and Ronald F Fox. The origins of order: self-organization and selection in evolution. *Biophysical Journal*, 65(6):2698–2699, 1993.

Kristel Kemper, Pauline L de Goeje, Daniel S Peeper, and Renée van Amerongen. Phenotype switching: tumor cell plasticity as a resistance mechanism and target for therapy. *Cancer Research*, 74 (21):5937–5941, 2014.

Hassan K Khalil and Jessy W Grizzle. *Nonlinear Systems*, volume 3. Prentice Hall, 2002.

Natalia L Komarova and Dominik Wodarz. The optimal rate of chromosome loss for the inactivation of tumor suppressor genes in cancer. *Proceedings of the National Academy of Sciences*, 101(18): 7017–7021, 2004.

Christian Kuehn. A mathematical framework for critical transitions: bifurcations, fast–slow systems and stochastic dynamics. *Physica D: Nonlinear Phenomena*, 240(12):1020–1035, 2011.

Marilyne Labrie, Joan S Brugge, Gordon B Mills, and Ioannis K Zervantonakis. Therapy resistance: opportunities created by adaptive responses to targeted therapies in cancer. *Nature Reviews Cancer*, 22(6):323–339, 2022.

Steven J Lade and Thilo Gross. Early warning signals for critical transitions: a generalized modeling approach. *PLoS computational biology*, 8(2):e1002360, 2012.

Lucie Laplane and Carlo C Maley. The evolutionary theory of cancer: challenges and potential solutions. *Nature Reviews Cancer*, 24(10):718–733, 2024.

Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, Fujie Huang, et al. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.

Vivien Marx. Method of the year: spatially resolved transcriptomics. *Nature methods*, 18(1):9–14, 2021.

Christopher D McFarland, Julia A Yaglom, Jonathan W Wojtkowiak, Jacob G Scott, David L Morse, Michael Y Sherman, and Leonid A Mirny. The damaging effect of passenger mutations on cancer progression. *Cancer Research*, 77(18):4763–4772, 2017.

Arash Mehrjou, Mohammad Ghavamzadeh, and Bernhard Schölkopf. Neural lyapunov redesign. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, volume 144, pp. 459–470. PMLR, 2021.

Arash Mehrjou, Andrea Iannelli, and Bernhard Schölkopf. Learning dynamical systems using local stability priors. *Journal of Computational Dynamics*, 10(1):175–198, 2023.

Naomi Moris, Cristina Pina, and Alfonso Martinez Arias. Transition states and cell fate decisions in epigenetic landscapes. *Nature Reviews Genetics*, 17(11):693–703, 2016.

Martin A Nowak, Natalia L Komarova, Anirvan Sengupta, Prasad V Jallepalli, Ie-Ming Shih, Bert Vogelstein, and Christoph Lengauer. The role of chromosomal instability in tumor initiation. *Proceedings of the National Academy of Sciences*, 99(25):16226–16231, 2002.

Ludovic Peyre, Marielle Péré, Mickael Meyer, Benjamin Bian, Marina Moureau-Barbato, Walid Djema, Bernard Mari, Georges Vassaux, and Jérémie Roux. Transition between cell states of sensitivity reveals molecular vulnerability of drug-tolerant cells. *Molecular Systems Biology*, 21 (12):1702, 2025.

Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pp. 1530–1538. PMLR, 2015.

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. *Nature*, 461(7260):53–59, 2009.

Sydney M Shaffer, Margaret C Dunagin, Stefan R Torborg, Eduardo A Torre, Benjamin Emert, Clemens Krepler, Marilda Beqiri, Katrin Sproesser, Patricia A Brafford, Min Xiao, et al. Rare cell variability and drug-induced reprogramming as a mode of cancer drug resistance. *Nature*, 546(7658):431–435, 2017.

Dongkwan Shin and Kwang-Hyun Cho. Critical transition and reversion of tumorigenesis. *Experimental & Molecular Medicine*, 55(4):692–705, 2023.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

René Thomas. Boolean formalization of genetic control circuits. *Journal of Theoretical Biology*, 42 (3):563–585, 1973.

Zhong Wang, Mark Gerstein, and Michael Snyder. Rna-seq: a revolutionary tool for transcriptomics. *Nature reviews genetics*, 10(1):57–63, 2009.

Jörgen W Weibull. *Evolutionary Game Theory*. MIT Press, 1997.

Jorge GT Zañudo and Réka Albert. Cell fate reprogramming by control of intracellular network dynamics. *PLoS Computational Biology*, 11(4):e1004193, 2015.

A Numerical illustrations

This appendix provides additional numerical illustrations complementing the main text. All simulations use the minimal three-dimensional GRN model introduced in Section 6.

A.1 Basins, bifurcations, and critical slowing down

Our first numerical illustration of the minimal GRN instantiates the three-attractor model introduced in Section 6 and tracks how its long-time outcomes shift with mutation strength. For each value of m, we draw many random initial conditions, integrate the dynamics, and record the asymptotic endpoints. In the bifurcation diagram of Fig. 4, *each dot is the long-time endpoint of one such trajectory at a given* m: upper branches (high S and high stemness) correspond to CSC-like equilibria, while lower branches correspond to differentiated or regulated-stem states. The stemness projection panel makes explicit how a high-stemness branch emerges and separates from lower-stemness branches as m increases, before saturating at higher mutation levels.

A.2 Minimal sustaining sets

To probe which regulatory interactions are structurally essential, we perform a systematic perturbation analysis over all parameters in the minimal GRN, treating each perturbation as a targeted *intervention* on a specific regulatory interaction (e.g., weakening self-renewal feedback or control

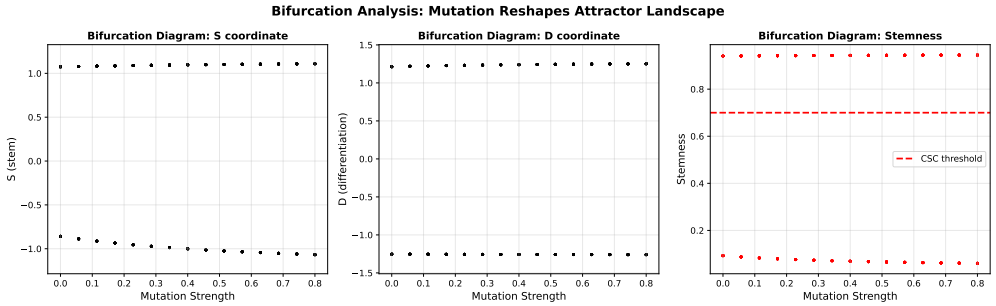

Figure 4: Bifurcation diagram for the minimal GRN as mutation strength m increases. For each value of m, we launch multiple trajectories from random initial conditions and plot their long-time endpoints as dots, projected onto (left) the S coordinate, (middle) the D coordinate (with a rescaled vertical axis to highlight the drift of the differentiated branch), and (right) the stemness score. The resulting point clouds trace out branches corresponding to distinct attractors: upper branches (high S, high stemness) are CSC-like equilibria, while lower branches are differentiated or regulated-stem states. Most of the qualitative change in branch positions occurs between $m = 0$ and moderate mutation levels, after which the branches stabilize, matching the idea that early mutation-driven remodeling deepens and separates the CSC basin while subsequent mutations produce smaller additional shifts.

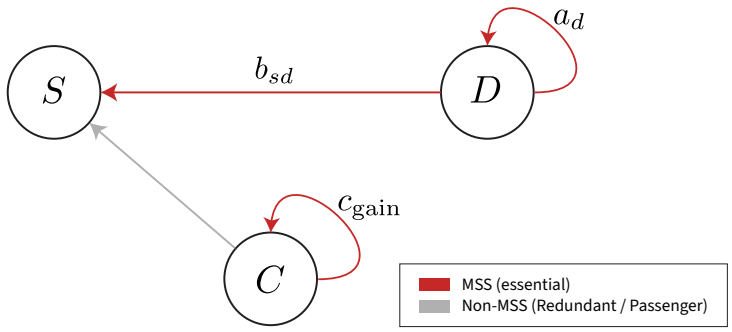

Figure 5: Minimal Sustaining Set (MSS) identification for the minimal GRN. Each "intervention" corresponds to reducing the strength of a particular regulatory interaction (self-activation, cross-inhibition, or control feedback) and recomputing the resulting dynamics from initial conditions near the CSC-like state to assess whether the malignant attractor and its basin persist. MSS members are those edges for which such interventions collapse or strongly erode the CSC basin, indicating that they materially reshape the local landscape, while non-MSS edges can be weakened without removing the malignant well, acting as passengers. This numerically instantiates the MSS and dual-control concepts introduced in Section 5, clarifying which specific feedbacks are structurally essential targets for landscape engineering.

gain). For a given parameter, we progressively reduce its strength, re-integrate trajectories starting near the original CSC-like state, and measure whether those trajectories still return to the same high-stemness attractor. In landscape terms, an intervention that eliminates the CSC attractor or drastically shrinks its basin is one that has successfully *reshaped the landscape* to remove or flatten the malignant well, whereas interventions that leave the CSC fixed point and its basin essentially intact correspond to landscape changes that are dynamically negligible at that attractor. The resulting classification of edges into essential versus redundant interactions is summarized schematically in Fig. 5, which highlights the minimal sustaining set (MSS) of interactions whose modification meaningfully changes the local landscape geometry around the CSC basin versus redundant passenger interactions that can be altered without eliminating the malignant state.

B MATHEMATICAL FORMALIZATION

B.1 PRELIMINARIES AND DEFINITIONS

Consider $d\boldsymbol{x}/dt = F(\boldsymbol{x};\theta)$ on $\mathcal{X} \subset \mathbb{R}^n$. An attractor $\boldsymbol{x}^*$ is asymptotically stable if $\forall\epsilon\,\exists\delta$ such that $\|\boldsymbol{x}(0) - \boldsymbol{x}^*\| < \delta \Rightarrow \|\boldsymbol{x}(t) - \boldsymbol{x}^*\| < \epsilon$ and $\boldsymbol{x}(t) \to \boldsymbol{x}^*$. The basin $\mathcal{B}(\boldsymbol{x}^*)$ is the set of initial conditions converging to $\boldsymbol{x}^*$. For quasi-gradient systems $F \approx -\nabla V(\boldsymbol{x};\theta)$, define the barrier height

$$\Delta E := \min_{\gamma:\, \boldsymbol{x}_A^* \to \boldsymbol{x}_B^*} \max_{\boldsymbol{x}\in\gamma} \big(V(\boldsymbol{x};\theta) - V(\boldsymbol{x}_A^*;\theta)\big),$$

where γ ranges over continuous paths from $\boldsymbol{x}_A^*$ to $\boldsymbol{x}_B^*$ crossing the saddle $\boldsymbol{x}_s$. In small-noise regimes, Kramers-type asymptotics yield mean first-passage times $\tau \propto \exp(\Delta E/\sigma^2)$ for noise magnitude σ.

An evolutionarily stable strategy (ESS) for frequency-dependent selection is a phenotype p^* such that no rare mutant q has higher fitness in the vicinity of p^*; formally $f(p^*, (1-\epsilon)p^* + \epsilon q) \geq f(q, (1-\epsilon)p^* + \epsilon q)$ for all small $\epsilon > 0$.

B.2 JOINT DYNAMICAL AND EVOLUTIONARY STABILITY

Proposition 3. *A malignant phenotype persists over long timescales if and only if it is both (1) a dynamical attractor with a nonempty basin and asymptotic stability, and (2) an evolutionarily stable strategy under frequency-dependent selection.*

Proof. ($\Rightarrow$) If the phenotype persists, intracellular dynamics must return trajectories to the phenotype under perturbations (asymptotic stability), and ecological interactions must prevent invasion by alternatives (ESS), else transient displacement occurs.

($\Leftarrow$) If the phenotype is a stable attractor, intracellular perturbations relax back. If it is an ESS, rare mutants cannot invade under replicator dynamics. Combining yields persistence under coupled dynamics where state trajectories are drawn to $\boldsymbol{x}^*$ and population composition resists invasion. A formalization models the population via $\dot{p}_i = p_i\big(\pi_i(p) - \bar{\pi}(p)\big)$ with payoffs depending on phenotypic state; dynamical stability ensures the p-dependent state remains in the basin, ESS ensures $\pi_i(p^*) \geq \pi_j(p^*)$ for invaders j in a neighborhood, prohibiting invasion. Stability of the joint fixed point follows from Lyapunov and invasion criteria. □

Corollary 2. *For phenotypes satisfying joint stability, landscape-only control or evolutionary-only control is insufficient for eradication; durable control requires joint manipulation of basin geometry and ecological payoffs.*

Proof. Landscape-only control can shrink or tilt the basin but leaves ESS intact; post-control, selection re-establishes the phenotype via re-invasion. Evolutionary-only control alters payoffs to break ESS but leaves the basin intact; plastic drift returns trajectories to the malignant attractor. A joint policy eliminates the attractor (or renders it unstable) and removes the selection advantage, preventing both drift and invasion. □

B.3 BARRIER HEIGHT AS COMMON CURRENCY

Proposition 4. *The barrier height ΔE determines (1) stochastic transition rates $r \propto \exp(-\Delta E/\sigma^2)$, (2) evolutionary accessibility via mutational barrier lowering, and (3) lower bounds on minimal control effort to cross the barrier.*

Proof. (1) Consider the stochastic dynamics $d\boldsymbol{x} = F(\boldsymbol{x})\,dt + \sigma\,dW_t$ where $F(\boldsymbol{x}) = -\nabla V(\boldsymbol{x})$ for a gradient system, and W_t is a standard Wiener process. By Freidlin–Wentzell large deviation theory, the mean first-passage time τ from attractor $\boldsymbol{x}_A^*$ to $\boldsymbol{x}_B^*$ satisfies

$$\lim_{\sigma\to 0} \sigma^2 \log \tau = \inf_{\gamma:\boldsymbol{x}_A^* \to \boldsymbol{x}_B^*} S[\gamma],$$

where the action functional is $S[\gamma] = \int_0^T \|\dot{\gamma}(t) - F(\gamma(t))\|^2 dt$ for paths γ connecting the attractors. The infimum is achieved along the minimum-action path γ^* that passes through the lowest saddle

point $\boldsymbol{x}_s$ between basins. Evaluating the action along γ^* yields $S[\gamma^*] = 2[V(\boldsymbol{x}_s) - V(\boldsymbol{x}_A^*)] = 2\Delta E$. Therefore, $\tau \propto \exp(2\Delta E/\sigma^2)$ in the small-noise limit, and the transition rate $r = 1/\tau \propto \exp(-\Delta E/\sigma^2)$ (where the factor of 2 is absorbed into the proportionality constant).

(2) Let the potential depend on parameters θ as $V(\boldsymbol{x};\theta)$. A mutation changes parameters from θ to $\theta + \delta\theta$, modifying the potential to $V(\boldsymbol{x};\theta+\delta\theta)$ and the barrier height to $\Delta E(\theta+\delta\theta) = \min_\gamma \max_{\boldsymbol{x}\in\gamma}[V(\boldsymbol{x};\theta+\delta\theta) - V(\boldsymbol{x}_A^*(\theta+\delta\theta);\theta+\delta\theta)]$. By continuity of V in θ, if $\Delta E(\theta + \delta\theta) < \Delta E(\theta)$, then applying part (1) with the post-mutation barrier, the transition rate becomes $r(\theta+\delta\theta) \propto \exp(-\Delta E(\theta+\delta\theta)/\sigma^2) > \exp(-\Delta E(\theta)/\sigma^2) \propto r(\theta)$. The ratio $r(\theta+\delta\theta)/r(\theta) = \exp([\Delta E(\theta) - \Delta E(\theta+\delta\theta)]/\sigma^2)$ shows exponential sensitivity: a decrease in barrier height by $\delta(\Delta E)$ increases the rate by a factor $\exp(\delta(\Delta E)/\sigma^2)$.

(3) For the controlled system $\dot{\boldsymbol{x}} = F(\boldsymbol{x}) + Bu$ with quadratic cost $J = \int_0^T \|u(t)\|^2 dt$, the optimal control problem is to minimize J subject to driving the system from $\boldsymbol{x}_A^*$ to $\boldsymbol{x}_B^*$. The Hamilton–Jacobi–Bellman (HJB) value function $V^*(x)$ satisfies $\min_u\{\|u\|^2 + \nabla V^*(\boldsymbol{x})^\top(F(\boldsymbol{x}) + Bu)\} = 0$ with boundary conditions. For any admissible control $u(t)$ achieving the transition, the HJB inequality gives $\|u(t)\|^2 \geq -\nabla V^*(\boldsymbol{x})^\top(F(\boldsymbol{x}) + Bu(t))$. Integrating along the optimal path γ^* from $\boldsymbol{x}_A^*$ to $\boldsymbol{x}_B^*$ crossing the saddle $\boldsymbol{x}_s$, and using that $F(\boldsymbol{x}) = -\nabla V(\boldsymbol{x})$ for gradient systems, we obtain $J^* = \int_{\gamma^*} \|u^*(t)\|^2 dt \geq \int_{\gamma^*} \|\nabla V(\boldsymbol{x})\|^2 dt \geq \int_{\gamma^*} dV = V(\boldsymbol{x}_s) - V(\boldsymbol{x}_A^*) = \Delta E$. For general F and control matrix B, the lower bound becomes $J^* \geq c\,\Delta E$ where $c = \inf_{\boldsymbol{x}\in\gamma^*} \lambda_{\min}((B^\top B)^{-1})$ depends on controllability and input constraints, establishing the result. □

B.4 Early-warning theorems and proofs

Theorem 1 (Critical slowing down)**.** *Let $\boldsymbol{x}^*(\mu)$ be a stable attractor of $d\boldsymbol{x}/dt = F(\boldsymbol{x};\mu)$ that loses stability at μ_c via a generic bifurcation. If the smallest-magnitude stable eigenvalue of the Jacobian satisfies $|\lambda_{\min}(\mu)| \to 0$ as $\mu \to \mu_c$, then the relaxation time $\tau(\mu) = 1/|\lambda_{\min}(\mu)|$ diverges as $\mu \to \mu_c$.*

Proof. Linearize near $\boldsymbol{x}^*$; solutions along the slowest mode decay as $e^{\lambda_{\min} t}$. As $\lambda_{\min} \to 0^-$ at criticality, the decay time diverges, producing increased autocorrelation and delayed recovery. □

Theorem 2 (Variance amplification)**.** *For the stochastic linearization $d\boldsymbol{y} = J(\mu)\boldsymbol{y}\,dt + \sigma\,dW_t$ near $\boldsymbol{x}^*(\mu)$, the stationary variance along the slow mode scales as $\mathrm{Var} \sim \sigma^2/(2|\lambda_{\min}(\mu)|)$, diverging as $\mu \to \mu_c$.*

Proof. Solve the Lyapunov equation $J\Sigma + \Sigma J^\top + \sigma^2 I = 0$. The dominant term along the slowest eigen-direction scales inversely with $|\lambda_{\min}|$. □

Theorem 3 (Flickering rate)**.** *In a bistable stochastic system with barrier height $\Delta E(\mu)$ separating attractors, the mean first-passage time satisfies $\tau \propto \exp(\Delta E(\mu)/\sigma^2)$; if $\Delta E(\mu) \downarrow 0$ as $\mu \to \mu_c$, the flickering rate $r = 1/\tau$ increases exponentially.*

Proof. Apply Kramers-type escape asymptotics along the minimum-energy path crossing the saddle; decreasing ΔE lowers the Arrhenius factor, increasing transition frequency. □

Theorem 4 (Spatial correlation)**.** *For reaction-diffusion linearizations $\partial_t \boldsymbol{y} = D\nabla^2\boldsymbol{y} + J(\mu)\boldsymbol{y} + \xi$ with diffusion $D > 0$ and noise ξ, the correlation length $\xi(\mu)$ increases as the system approaches criticality, typically $\xi \sim (D/|\lambda_{\min}|)^\alpha$ for some $\alpha > 0$ depending on geometry.*

Proof. Fourier-transform the linearized dynamics; the spectral density concentrates near low k as the gap $|\lambda_{\min}|$ closes, inflating the spatial correlation length obtained from the inverse decay scale. □

B.5 ESS invadability and evolutionary stability

Theorem 5 (ESS conditions for CSC-like strategies)**.** *Consider a population with phenotypes indexed by stemness score $s \in [0,1]$ and composition $p(s)$, where $p_s = \int_{s'\approx s} p(s')ds'$ denotes the frequency of cells with similar stemness. Let the payoff function be $\pi(s;p,\eta) = \eta\cdot s - \beta\cdot(1-s) - \alpha\cdot p_s$*

with parameters $\eta > 0$ (microenvironmental support), $\beta > 0$ (differentiation cost), and $\alpha > 0$ (competitive interaction strength). A CSC-like strategy with stemness s^* near 1 is evolutionarily stable (ESS) if and only if $\eta > \beta + \alpha p_{s^*}$ and $p_{s^*} < \frac{\eta-\beta}{\alpha}$. Conversely, CSC-like strategies are invadable by differentiated strategies (low s) when $\eta \leq \beta + \alpha p_{s^*}$ or when $p_{s^*} \geq \frac{\eta-\beta}{\alpha}$.

Proof. A composition p^* with dominant phenotype s^* is evolutionarily stable if, for all nearby mutants q with stemness $s' \neq s^*$, small invasions reduce mutant fitness: $\pi(s^*; (1-\epsilon)p^* + \epsilon q) \geq \pi(s'; (1-\epsilon)p^* + \epsilon q)$ for all sufficiently small $\epsilon > 0$.

For a CSC-like strategy $s^* \approx 1$, the payoff is $\pi(s^*; p^*, \eta) = \eta \cdot s^* - \beta \cdot (1 - s^*) - \alpha \cdot p_{s^*}$. For a differentiated mutant with $s' \approx 0$, the payoff is $\pi(s'; p^*, \eta) = \eta \cdot s' - \beta \cdot (1 - s') - \alpha \cdot p_{s'}$.

In the limit $s^* \to 1$ and $s' \to 0$, we have $\pi(s^*; p^*, \eta) \approx \eta - \alpha p_{s^*}$ and $\pi(s'; p^*, \eta) \approx -\beta - \alpha p_{s'}$. For the CSC-like strategy to resist invasion by differentiated mutants, we require $\pi(s^*; p^*, \eta) > \pi(s'; p^*, \eta)$, which yields $\eta - \alpha p_{s^*} > -\beta - \alpha p_{s'}$, or equivalently $\eta > \alpha p_{s^*} - \beta - \alpha p_{s'}$. Since $p_{s'} \approx 0$ for rare differentiated mutants, this simplifies to $\eta > \alpha p_{s^*} - \beta$, or $\eta > \beta + \alpha p_{s^*}$ when $p_{s^*} > 0$.

For stability against invasion by alternative CSC-like strategies with slightly different s, we consider a mutant s'' near s^* with frequency ϵ in the perturbed population. The payoff difference is $\Delta\pi = \pi(s^*; p_\epsilon, \eta) - \pi(s''; p_\epsilon, \eta) = \eta(s^* - s'') - \alpha(p_{s^*} - p_{s''})$ for small $|s^* - s''|$. Stability requires $\Delta\pi > 0$ for all small ϵ, which is satisfied when α is sufficiently large relative to η and the frequency gradient is not too steep, or when $\eta > \alpha p_{s^*}$ and p_{s^*} is moderate.

Combining these conditions, a CSC-like strategy with $s^* \approx 1$ is ESS if $\eta > \beta + \alpha p_{s^*}$ and $p_{s^*} < \frac{\eta-\beta}{\alpha}$, establishing the necessary and sufficient conditions. The converse (invadability conditions) follows immediately from the negation of these inequalities. □

B.6 Mutation-driven equilibrium shifts

For the reduced dynamical system with mutation-adjusted parameters $a_s^{\text{eff}} = a_s(1 + m\mu_{as})$ and $k_c^{\text{eff}} = k_c(1 - m\mu_{kc})$ where $m \in [0, 1]$ is mutation strength, we analyze how equilibria shift with mutations. For the case with control off (control_on $= 0$), the system has equilibria at (S^*, D^*, C^*) satisfying $a_s^{\text{eff}}(S^* - (S^*)^3) - b_{sd}D^* = 0$ and $a_d(D^* - (D^*)^3) - b_{ds}S^* = 0$. As mutation strength m increases, the CSC-like equilibrium (with $S^* > 0$ and high stemness) shifts to higher S^* values, and the basin of attraction for this equilibrium expands. Specifically, if (S_0^*, D_0^*) is an equilibrium for $m = 0$, then for $m > 0$ there exists an equilibrium (S_m^*, D_m^*) with $S_m^* > S_0^*$ when $S_0^* > 0$, and the stemness score stemness$(S_m^*, D_m^*) >$ stemness(S_0^*, D_0^*).

To establish this, we analyze the equilibrium conditions. The equilibria satisfy $F_S(S^*, D^*) = a_s^{\text{eff}}(S^* - (S^*)^3) - b_{sd}D^* = 0$ and $F_D(S^*, D^*) = a_d(D^* - (D^*)^3) - b_{ds}S^* = 0$.

Differentiating the equilibrium conditions with respect to m using the implicit function theorem, we have:

$$\frac{\partial F_S}{\partial S^*}\frac{dS^*}{dm} + \frac{\partial F_S}{\partial D^*}\frac{dD^*}{dm} + \frac{\partial F_S}{\partial m} = 0,$$
$$\frac{\partial F_D}{\partial S^*}\frac{dS^*}{dm} + \frac{\partial F_D}{\partial D^*}\frac{dD^*}{dm} = 0,$$

where $\frac{\partial F_S}{\partial m} = a_s\mu_{as}(S^* - (S^*)^3) > 0$ for $S^* > 0$ (since $S^* - (S^*)^3 > 0$ for $S^* \in (0, 1)$), and $\frac{\partial F_S}{\partial S^*} = a_s^{\text{eff}}(1 - 3(S^*)^2)$, $\frac{\partial F_S}{\partial D^*} = -b_{sd}$, $\frac{\partial F_D}{\partial S^*} = -b_{ds}$, $\frac{\partial F_D}{\partial D^*} = a_d(1 - 3(D^*)^2)$.

For a stable CSC-like equilibrium with $S^* > 0$ and $D^* < 0$ (mutual inhibition), the Jacobian has negative eigenvalues, implying $\frac{\partial F_S}{\partial S^*} < 0$ and $\frac{\partial F_D}{\partial D^*} < 0$ at the equilibrium. Solving the system for $\frac{dS^*}{dm}$ yields $\frac{dS^*}{dm} > 0$ when $\frac{\partial F_S}{\partial m} > 0$, establishing that S^* increases with m.

Since the stemness score stemness$(S, D) = 0.8 \cdot \sigma(2.5S) + 0.2 \cdot \sigma(-2.5D)$ is monotonically increasing in S and decreasing in D (for $D < 0$), and D^* decreases as S^* increases (due to mutual inhibition $-b_{ds}S^*$), we have stemness$(S_m^*, D_m^*) >$ stemness(S_0^*, D_0^*) for $m > 0$.

The expansion of the basin follows from analyzing how increasing a_s^{eff} affects the vector field. The self-activation term $a_s^{\text{eff}}(S - S^3)$ creates a cubic nullcline with stable fixed points near $S = \pm 1$

and an unstable point at $S = 0$. When a_s^{eff} increases, the magnitude of the self-activation force $|a_s^{\text{eff}}(S - S^3)|$ grows for all $S \in (0, 1)$, meaning trajectories are pulled more strongly toward the CSC-like attractor at $S \approx 1$.

Geometrically, this corresponds to the quasi-potential landscape $V(S, D)$ developing a deeper and wider valley around the CSC-like equilibrium. The basin boundary, defined by the stable manifold of the saddle point separating attractors, shifts outward as the CSC-like attractor becomes more stable (more negative eigenvalues in the Jacobian). Intuitively, a stronger self-activation term means that initial conditions with moderate stemness ($S \in (0, S_{\text{saddle}})$) that previously would have been attracted to the differentiated state ($S \approx -1$) or remained near the unstable point are now pulled into the CSC-like basin, effectively expanding the region of state space that converges to high stemness.

B.7 Learning-based Attractome reconstruction

The following pseudocode outlines how generative models can be trained to learn landscape surrogates and integrated into the Attractome framework:

Algorithm: Learning-based Attractome reconstruction

1. Data Collection:

1. Collect single-cell RNA-seq, ATAC-seq, or proteomics measurements $\mathcal{D} = \{(\boldsymbol{x}_i, \mathbf{u}_i)\}_{i=1}^N$, where $\boldsymbol{x}_i \in \mathbb{R}^n$ represents the high-dimensional gene expression state (e.g., transcript counts for n genes) of cell i, and $\mathbf{u}_i$ denotes microenvironmental or experimental conditions (e.g., growth factors, drug concentrations, oxygen levels) that modulate cellular dynamics.
2. Optionally include perturbation data: $\mathcal{D}_{\text{pert}} = \{(\boldsymbol{x}_i, \mathbf{u}_i, \delta\mathbf{u}_i, \boldsymbol{x}_i')\}$, where $\boldsymbol{x}_i$ is the pre-perturbation state, $\delta\mathbf{u}_i$ is the applied perturbation (e.g., drug dose, gene knockout, cytokine addition), and $\boldsymbol{x}_i'$ is the post-perturbation state after some time interval. Perturbation data enables direct learning of response functions $\boldsymbol{x}' = f(\boldsymbol{x}, \delta\mathbf{u})$ and gradient-like structure $\mathbf{F}_{\text{eff}}(\boldsymbol{x}) \approx (\boldsymbol{x}' - \boldsymbol{x})/\Delta t$, which constrains the learned quasi-potential to satisfy $-\nabla V(\boldsymbol{x}) \propto \mathbf{F}_{\text{eff}}(\boldsymbol{x})$ and improves identification of barrier locations and basin boundaries.
3. Optionally include temporal snapshots: $\mathcal{D}_{\text{time}} = \{(\boldsymbol{x}_i^{(t)}, t)\}$ (trajectories), which capture stochastic transitions between attractors and provide empirical estimates of transition rates that can be matched to Kramers-type predictions $\tau \propto \exp(\Delta E/\sigma^2)$ to validate learned barrier heights.

2. Training Parametric Model:

1. Initialize generative model $p_\phi(\boldsymbol{x}|\boldsymbol{z})$ (score-based/diffusion, normalizing flow, or energy-based)
2. Initialize encoder $q_\psi(\boldsymbol{z}|\boldsymbol{x})$ for latent state $\boldsymbol{z} \in \mathbb{R}^d$ where $d \ll n$
3. **Loss Function:**
   - Reconstruction: $\mathcal{L}_{\text{recon}} = \mathbb{E}_{(\boldsymbol{x},\boldsymbol{z})\sim p_{\text{data}}}[-\log p_\phi(\boldsymbol{x}|\boldsymbol{z})]$
   - Latent regularization: $\mathcal{L}_{\text{reg}} = \text{KL}(q_\psi(\boldsymbol{z}|\boldsymbol{x})||\mathcal{N}(0, I))$
   - **Attractome-informed regularizers:**
     - Multistability constraint: $\mathcal{L}_{\text{multistab}} = -\sum_{k=1}^K \log\det(\text{Cov}(\boldsymbol{z}|\text{attractor}_k))$ (encourage distinct clusters)
     - Barrier regularity: $\mathcal{L}_{\text{barrier}} = \lambda \sum_{\text{saddle } s} \|\nabla V(\boldsymbol{z}_s)\|^2$ (smooth barriers at saddles)
     - Quasi-potential consistency: $\mathcal{L}_{\text{pot}} = \mathbb{E}[(\partial V/\partial \boldsymbol{z} \cdot \mathbf{F}_{\text{eff}} + \|\mathbf{F}_{\text{eff}}\|^2)^2]$ (gradient-like structure)
   - Total: $\mathcal{L} = \mathcal{L}_{\text{recon}} + \beta\mathcal{L}_{\text{reg}} + \gamma_1\mathcal{L}_{\text{multistab}} + \gamma_2\mathcal{L}_{\text{barrier}} + \gamma_3\mathcal{L}_{\text{pot}}$
4. Optimize: $\phi^*, \psi^* = \arg\min_{\phi,\psi} \mathcal{L}$ using SGD/Adam

3. Deployment in Attractome Framework:

1. Extract learned quasi-potential: $V(\boldsymbol{z}; \phi^*) = -\log p_\phi(\boldsymbol{z})$ (up to normalization)
2. Infer effective dynamics: $\mathbf{F}_{\text{eff}}(\boldsymbol{z}) = -\nabla_{\boldsymbol{z}} V(\boldsymbol{z}; \phi^*) + \text{noise}$

3. Compute barrier heights: $\Delta E_{ij} = \min_{\gamma:\boldsymbol{z}_i^* \to \boldsymbol{z}_j^*} \max_{\boldsymbol{z}\in\gamma}(V(\boldsymbol{z}) - V(\boldsymbol{z}_i^*))$
4. Identify attractors: $\{\boldsymbol{z}_k^*\} = \{\boldsymbol{z} : \mathbf{F}_{\text{eff}}(\boldsymbol{z}) = 0, \lambda_{\max}(\nabla\mathbf{F}_{\text{eff}}) < 0\}$
5. Compute basin boundaries and MSS from sensitivity: MSS $= \arg\min_{|\theta'| \text{ s.t. } \Delta E(\theta') < \Delta E(\theta_0)} |\theta'|$
6. Integrate with replicator dynamics: $\dot{p}_k = p_k(f_k(p) - \bar{f}(p))$ where $f_k(p) = \pi(\text{stemness}(\boldsymbol{z}_k^*), p, \eta)$
7. Evaluate ESS conditions and control policies using learned ΔE values