

Identifying Essential Rule Sets in Agent-Based Models Through Systematic Ablation: A Tumor Evolution Case Study

Nhung Duong
Hanoi University of Pharmacy
N2TP Technology Solutions JSC
Ha Noi, Viet Nam
2312015@hup.edu.vn

Luong Doan
N2TP Technology Solutions JSC
Phenikaa University
Ha Noi, Viet Nam
21010345@st.phenikaa-uni.edu.vn

Anh Do
Hanoi University of Pharmacy
Ha Noi, Viet Nam
2381004@hup.edu.vn

Anh Truong
Hanoi University Of Pharmacy
Ha Noi, Viet Nam
2281005@hup.edu.vn

Ngoc Do
Hanoi University of Pharmacy
Ha Noi, Viet Nam
dominhngoc_hsgs19@hus.edu.vn

Tien Nguyen
N2TP Technology Solutions JSC
Ha Noi, Viet Nam
2023b003@oasis.inje.ac.kr

Tuan Do
Phenikaa University
Ha Noi, Viet Nam
tuando7758@gmail.com

ABSTRACT

We present a systematic methodology for identifying essential rule sets in agent-based models, demonstrated through tumor evolution under therapeutic pressure. Our framework addresses a critical challenge: determining which interaction rules are necessary for reproducing emergent spatiotemporal patterns versus those representing auxiliary complexity. We couple autonomous cancer cell agents with reaction-diffusion fields, then systematically ablate mechanisms including paracrine signaling, metabolic competition, phenotypic plasticity, and spatial interactions. Our approach establishes empirically-calibrated significance thresholds from baseline system stochasticity, providing principled criteria for mechanism classification. A key methodological insight distinguishes ablatable biological rules from non-ablatable framework requirements: removing spatial diffusion caused model failure, revealing it encodes physical constraints rather than testable hypotheses. We demonstrate that complex patterns emerge from minimal rule sets, with several commonly-modeled mechanisms contributing negligibly to tissue-scale behavior. This methodology advances multi-agent simulation by providing objective model reduction techniques that enhance mechanistic interpretability and facilitate parameter calibration.

KEYWORDS

Agent-based simulation; Tumor modeling; Complex systems; Model reduction; Multi-scale modeling

ACM Reference Format:

Nhung Duong, Luong Doan, Anh Do, Anh Truong, Ngoc Do, Tien Nguyen, and Tuan Do. 2026. Identifying Essential Rule Sets in Agent-Based Models



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/OQYW6891>

Through Systematic Ablation: A Tumor Evolution Case Study. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/OQYW6891>

1 INTRODUCTION

Agent-based models have become indispensable tools for understanding complex systems across domains including ecology, epidemiology, economics, and biology [3]. The agent-based modeling paradigm’s power lies in its ability to generate emergent system-level phenomena from simple individual-level interaction rules. However, contemporary agent-based models often incorporate dozens of mechanisms in pursuit of biological or social realism, creating models with hundreds of parameters whose individual contributions to emergent behavior remain unclear. This proliferation of rules complicates model interpretation, parameter calibration, and computational tractability while potentially obscuring the minimal mechanisms sufficient for pattern formation.

The fundamental question we address is: Given an agent-based model $\mathcal{M}_{\text{full}}$ with rule set $\mathcal{R} = \{r_1, r_2, \dots, r_n\}$ producing observable pattern distributions $P_{\mathcal{M}_{\text{full}}}(\Psi)$ over pattern space Ψ , what is the minimal subset $\mathcal{R}_{\text{min}} \subset \mathcal{R}$ such that the reduced model \mathcal{M}_{min} satisfies $D_W(P_{\mathcal{M}_{\text{min}}}(\Psi), P_{\mathcal{M}_{\text{full}}}(\Psi)) < \epsilon$, where D_W denotes the Wasserstein distance and ϵ represents acceptable tolerance? This question becomes particularly acute in biological systems where many plausible mechanisms exist but experimental evidence for each mechanism’s necessity remains limited.

We develop a systematic ablation methodology that addresses three key challenges in agent-based model reduction. First, we must distinguish genuine mechanistic contributions from stochastic variability inherent to agent-based systems. Second, we require metrics that capture emergent spatiotemporal patterns rather than instantaneous states. Third, we need principled statistical frameworks for testing significance when comparing high-dimensional pattern distributions. Our approach tackles these challenges through comprehensive baseline variability analysis, composite pattern fingerprints

aggregating spatial and temporal information, and empirically-calibrated divergence thresholds.

We demonstrate this methodology through a case study in tumor evolution under drug pressure, where cancer cell agents interact with continuous drug and nutrient fields in three-dimensional tissue environments. This domain provides an ideal testbed due to well-characterized cellular mechanisms, rich emergent dynamics including spatial heterogeneity and resistance evolution, and clinical relevance. Our findings reveal that despite modeling ten distinct biological mechanisms, only three prove essential for reproducing observed patterns, with several widely-studied processes, including cell-cell communication and immune surveillance, contributing negligibly to emergent behavior at tissue scales.

2 RELATED WORK

Agent-based modeling in oncology has evolved from simple cellular automata to sophisticated multi-scale frameworks capturing genetic, cellular, and tissue-level processes. Gallaher et al. [6] and colleagues pioneered coupling discrete cell agents with continuous reaction-diffusion equations for nutrients and signaling molecules, enabling realistic spatial gradients. The PhysiCell framework [7] advanced this paradigm by incorporating detailed biophysical constraints and cell cycle dynamics. The Hybrid Automata Library [2] provides a flexible platform for building such hybrid models with real-time visualization. However, these comprehensive models conflate essential mechanisms with auxiliary complexity, limiting biological insight.

Model reduction in systems biology traditionally focuses on parameter sensitivity analysis and sloppy model concepts introduced by Gutenkunst et al. [8], which identify parameter combinations that minimally affect model predictions. Complementary work [13] demonstrated that many biological models exhibit flat likelihood surfaces in most parameter directions while remaining sensitive to specific combinations. Our work extends these ideas from parameter-level to rule-level reduction, asking which entire mechanisms can be removed rather than which parameter values matter least.

The concept of using ablation studies to understand complex systems originates in neuroscience, where lesion studies identify brain regions essential for specific functions [12]. Computational implementations include knockout experiments in Boolean network models [11] and edge removal in network dynamics. In agent-based modeling, sensitivity analysis typically varies parameters continuously rather than discretely removing mechanisms. Our contribution systematically applies ablation methodology to multi-agent systems with rigorous statistical frameworks for significance testing.

Pattern analysis in spatial agent-based models commonly employs summary statistics including spatial correlation functions, clustering metrics, and nearest-neighbor distributions. Recent work has incorporated topological data analysis via persistent homology to capture multiscale spatial organization [1, 15]. Our approach synthesizes these methods into composite pattern fingerprints while addressing the crucial question of how large pattern differences must be to exceed baseline stochasticity. This question is rarely addressed in agent-based model comparison studies [10].

3 METHODOLOGY

3.1 Hybrid Agent-Based and PDE Framework

Our model couples autonomous cell agents with continuous fields governed by reaction-diffusion partial differential equations, a hybrid architecture applicable to any domain where discrete entities interact with continuous environmental fields [2]. We employ this approach rather than pure agent-based or pure continuum models for generalizable reasons: (1) PDEs efficiently represent diffusing molecular species, chemical signals, or resource distributions without tracking individual molecules, with computational cost scaling as $O(N_{\text{grid}}^3)$ versus $O(N_{\text{molecules}} \cdot N_{\text{agents}})$ for particle-based approaches; (2) discrete agents preserve stochastic effects critical for rare events; (3) individual-level heterogeneity in responses to environmental signals is naturally captured; (4) direct correspondence to experimental measurements spanning both scales.

The complete system state at time t consists of continuous fields $\{C_{\text{drug}}^t(\mathbf{r}), C_{\text{nutrient}}^t(\mathbf{r})\}_{\mathbf{r} \in \Omega}$ representing drug and nutrient concentrations over spatial domain $\Omega \subset \mathbb{R}^3$, along with discrete agents $\{C_{\ell}^t\}_{\ell=1}^{N_t}$ each characterized by position $\mathbf{p}_{\ell} \in \Omega$, internal state vector $\mathbf{x}_{\ell} \in \mathbb{R}^{75}$ encoding gene expression and signaling, and phenotype $\tau_{\ell} \in \{0, 1, 2, 3\}$ representing sensitive, tolerant, resistant, or immune cell types.

Fields evolve according to reaction-diffusion equations:

$$\frac{\partial C}{\partial t} = D \nabla^2 C - \kappa C + S(\mathbf{r}, t) - U(\mathbf{r}, t) \quad (1)$$

where D represents diffusion coefficient ($D_{\text{drug}} = 0.15$, $D_{\text{nutrient}} = 0.20$ spatial units²/hour), κ decay rate ($\kappa_{\text{drug}} = 0.02$, $\kappa_{\text{nutrient}} = 0.01$ hour⁻¹), S external sources, and U cellular uptake. We discretize using Crank-Nicolson alternating direction implicit schemes with grid resolution 25^3 (grid spacing $h = 2.0$ units) and timestep $\Delta t = 1$ hour, ensuring second-order accuracy and unconditional stability. Cellular uptake is distributed via cloud-in-cell interpolation, where each cell at position \mathbf{p}_{ℓ} deposits uptake to eight surrounding grid points with weights $w_{i,j,k}(\mathbf{p}_{\ell}) = \prod_{d \in \{x,y,z\}} \max(0, 1 - |p_{\ell,d} - d_{i,j,k}|/h)$, normalized such that $\sum_{i,j,k} w_{i,j,k} = 1$. Mass-conserving uptake limits withdrawal by available mass: $U_{\text{effective}} = \min(U_{\text{requested}}, C_{\text{current}}/\Delta t)$.

Cell agents undergo biased random walks combining persistence, diffusion, and chemotaxis [14]:

$$\mathbf{p}_{\ell}(t + \Delta t) = \mathbf{p}_{\ell}(t) + \mathbf{v}_{\text{persist},\ell}(t) \cdot \Delta t + \sqrt{2D_{\text{cell}}\Delta t} \xi_{\ell} \quad (2)$$

where persistence velocity evolves as:

$$\mathbf{v}_{\text{persist},\ell}(t + \Delta t) = \alpha \mathbf{v}_{\text{persist},\ell}(t) + (1 - \alpha) \boldsymbol{\eta}_{\ell} + \beta \nabla C_{\text{nutrient}}(\mathbf{p}_{\ell}) \quad (3)$$

with $\alpha = 0.7$ controlling memory timescale and $\beta = 0.1$ determining chemotactic strength.

Proliferation occurs stochastically with rates multiplicatively modulated by local nutrient availability via Michaelis-Menten kinetics $\phi_{\text{nutrient}}(C) = C^2/(K_{\text{nut}}^2 + C^2)$ and local density-dependent suppression $\phi_{\text{density}}(n_{\text{local}}) = \max(0, 1 - n_{\text{local}}/K_{\text{local}})$, where n_{local} counts neighbors within radius $R_{\text{density}} = 5$ spatial units. Death combines baseline apoptosis, drug-induced cytotoxicity following Hill equations with phenotype-specific IC₅₀ values (sensitive: IC_{50,0} = 1.0, resistant: IC_{50,2} = 15.0 concentration units), and nutrient starvation below critical thresholds. Phenotypic transitions follow continuous-time Markov chains with drug-dependent rates.

The model captures ten candidate mechanisms: (R1) nutrient-dependent proliferation, (R2) drug-induced death, (R3) stochastic phenotypic transitions, (R4) spatial diffusion creating gradients, (R5) paracrine signaling via coupling matrices, (R6) cell migration with chemotaxis, (R7) density-dependent contact inhibition, (R8) metabolic competition through nutrient uptake, (R9) immune surveillance by specialized agents, and (R10) cooperative resistance where resistant cells protect neighbors. Each mechanism represents a candidate for ablation studies testing its necessity for pattern formation.

Computational Tractability: Neighbor searches for rules R5, R7, R9, and R10 require identifying agents within interaction radius $R_{\text{interact}} = 8$ spatial units. For populations $N < 10^4$, k-d trees provide $O(N \log N)$ construction and $O(\log N)$ per query. Overall algorithmic complexity per timestep: $O(N_{\text{grid}}^3)$ for PDE solver, $O(N_{\text{agents}})$ for agent state updates, and $O(N_{\text{agents}} \log N_{\text{agents}})$ for neighbor-dependent rules, yielding approximately 6 minutes per 720-timestep simulation on typical workstations.

3.2 Ablation Study Design

We designed a comprehensive ablation study testing ten candidate mechanisms. Rules R1 (nutrient-dependent proliferation) and R2 (drug-induced death) were not ablated as they constitute definitional model requirements. We generated variants systematically removing each of the remaining eight mechanisms: M3 (no phenotypic transitions), M4 (no spatial diffusion), M5 (no paracrine signaling), M6 (no cell migration), M7 (no density effects), M8 (no metabolic competition), M9 (no immune surveillance), M10 (no cooperative resistance), plus minimal model M22 removing R7-R10 simultaneously.

Preliminary screening with small-sample pilot studies identified M4 as pathological: removing spatial diffusion produced 23-fold population explosion (mean final population 27,591 versus baseline 1,184 cells, median 24,452 versus 1,212 cells, Mann-Whitney $p < 0.001$) due to violated mass conservation. Population distributions showed extreme variance with 78% of M4 simulations (39/50) exceeding 10,000 cells compared to 0% (0/50) in baseline. This catastrophic failure stems from eliminated spatial gradients: without diffusion-generated concentration heterogeneity, all cells experience uniform nutrient availability regardless of spatial position, eliminating carrying capacity. This revealed spatial diffusion as a physical framework requirement encoding Fick’s law and conservation principles rather than an ablatable biological hypothesis. We excluded M4 from subsequent pattern analysis, focusing statistical comparison on the eight remaining variants where ablation tests genuine biological mechanism necessity.

Each variant was simulated for 50 independent replicates with parameters drawn from Latin hypercube sampling over 15-dimensional parameter space including diffusion coefficients, transition rates, drug efficacies, proliferation rates, and initial conditions. Simulations ran for $T = 720$ hours (30 days) on spatial domains $[0, 50]^3$ with snapshots recorded every 6 hours. This design yielded 450 total simulations enabling robust statistical comparison.

3.3 Pattern Characterization

Our feature set was chosen to mirror observables used to characterize tumor tissue phenotypes. Ripley’s K-function and the Morisita index quantify spatial clustering and heterogeneity, analogous to histological assessment of tumor architecture. Persistent homology captures multiscale topological organization, specifically whether cells form connected masses or fragmented clusters, that pathologists evaluate qualitatively in tissue sections. Phenotypic fractions and Shannon entropy correspond to subpopulation composition measurements obtained through flow cytometry or single-cell assays. Graph-based clustering coefficient and modularity characterize the local connectivity and community structure of the spatial cell arrangement, reflecting how tightly packed and spatially partitioned the population is. Population size and nearest-neighbor distance track standard growth curve endpoints and cell packing density. Together these seven categories span the spatial, compositional, and topological axes along which tumor phenotypes are distinguished in practice. We note that this fingerprint was not optimized via formal feature selection; whether a leaner subset yields equivalent mechanism classifications is a natural direction for future work. We constructed 85-dimensional pattern fingerprints aggregating seven metric categories at each of five timepoints ($t \in \{120, 240, 360, 480, 720\}$ hours). These timepoints sample the simulation at approximately regular intervals from early growth through the final state, ensuring that ablation effects are detected whether they manifest as transient disruptions or persistent alterations to long-term organization. Including multiple timepoints rather than only the endpoint prevents misclassification of mechanisms whose removal causes early divergence followed by reconvergence, or vice versa. For each snapshot, we computed:

(1) **Spatial statistics** via Ripley’s K-function [9] evaluated at radii $r \in \{2, 5, 10, 15, 20\}$ spatial units with isotropic edge correction, normalized as $\hat{L}(r) = \left(\frac{3\hat{K}(r)}{4\pi}\right)^{1/3} - r$ for three-dimensional point patterns (5 features).

(2) **Morisita index** $I_M = Q \sum_q n_q(n_q - 1) / (N(N - 1))$ measuring spatial heterogeneity through cell counts n_q in $Q = 5^3 = 125$ cubic quadrats (1 feature).

(3) **Topological persistence** [1, 15] via Vietoris-Rips complexes on farthest-point subsampled point clouds (500 landmarks, maximum radius 15 units), summarized through total persistence for H_0 (connected components) and H_1 (loops/voids) homology groups (2 features).

(4) **Phenotypic fractions** $\{f_S, f_T, f_R, f_I\}$ enumerating proportions of sensitive, tolerant, resistant, and immune cell types (4 features).

(5) **Shannon entropy** $H_{\text{pheno}} = -\sum_{\tau=0}^3 p_\tau \log p_\tau$ quantifying phenotypic diversity (1 feature).

(6) **Graph metrics** from spatial graphs constructed via hybrid k-nearest-neighbor ($k = 6$) and radius-based ($d_{\text{connect}} = 8$ units) connectivity, computing clustering coefficient and modularity via Louvain community detection (2 features).

(7) **Population statistics** including total cell count N^t and mean nearest-neighbor distance (2 features).

This yielded 17 features per timepoint, concatenated across five timepoints to produce $\Psi \in \mathbb{R}^{85}$. All fingerprints were normalized by subtracting mean vector and dividing by standard deviation

vector computed exclusively from 50 baseline replicates, ensuring comparisons measured divergence in standardized units relative to the baseline reference distribution.

3.4 Empirical Threshold Calibration

A critical innovation addresses the question: how large must pattern divergence be to confidently attribute differences to mechanism removal rather than stochastic variation? We established empirically-calibrated thresholds by measuring baseline-to-baseline variability [10]. The 50 baseline replicates were randomly split into two groups of 25 samples each, and Sliced Wasserstein distance was computed between group distributions using 500 random projections. Repeating this procedure 100 times yielded a null distribution with mean $\mu = 0.364$ and standard deviation $\sigma = 0.038$, demonstrating that identical model parameterizations diverge considerably (95% of baseline comparisons falling between $SW = 0.29$ and $SW = 0.44$) due solely to stochastic agent dynamics.

We classified mechanisms as **auxiliary** if divergence fell within one standard deviation of baseline noise ($SW < \mu + 1\sigma = 0.402$), **intermediate** if divergence exceeded noise but remained within three standard deviations ($0.402 < SW < \mu + 3\sigma = 0.478$), and **core** if divergence exceeded three standard deviations ($SW > 0.478$). These data-driven thresholds replace arbitrary cutoffs and provide principled statistical interpretation. Table 1 summarizes the complete threshold calibration results.

3.5 Distributional Comparison via Sliced Wasserstein Distance

We quantified pattern divergence using Sliced Wasserstein distance [4, 5], which measures dissimilarity between probability distributions by integrating one-dimensional Wasserstein distances over random projections:

$$SW_2(P, Q) = \left[\int_{\mathbb{S}^{d-1}} W_2^2(P^\theta, Q^\theta) d\theta \right]^{1/2} \quad (4)$$

approximated via Monte Carlo sampling. For each iteration, we drew a random unit vector $\theta \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ normalized to unit length, projected both sample sets onto this direction via matrix multiplication, computed the one-dimensional Wasserstein distance between projected samples, and averaged over 500 projections. This metric provides computational efficiency ($O(n \log n)$ per projection versus $O(n^3)$ for full optimal transport) and robustness to dimensionality.

All comparisons used consistent sample sizes: groups of 25 fingerprints were randomly sampled without replacement from each variant’s 50 replicates before computing distances. Bootstrap confidence intervals were constructed by resampling with replacement 1000 times. Statistical significance was assessed via two complementary non-parametric tests: Mann-Whitney U testing differences in distances from baseline centroid, and Kruskal-Wallis testing differences in first principal component scores, with Bonferroni correction for multiple testing ($\alpha = 0.05/8 = 0.0063$ for eight variant comparisons).

4 RESULTS

4.1 Spatial Diffusion is a Framework Requirement, Not an Ablatable Mechanism

Preliminary analysis of variant M4 (no spatial diffusion) revealed catastrophic model failure, providing a crucial methodological lesson: not all model components constitute ablatable biological hypotheses. Removing spatial diffusion produced 23-fold population explosion compared to baseline (mean final population 27,591 versus 1,184 cells, Mann-Whitney $p < 0.0001$). Population distributions showed extreme variance with 78% of simulations exceeding 10,000 cells and 18% exceeding 50,000 cells, compared to 0% in baseline.

This pathology stems from violated physical constraints. Spatial diffusion encodes fundamental conservation laws: nutrients must be transported from sources to consumers via concentration gradients following Fick’s law ($J = -D\nabla C$), and local consumption must create depletion zones. Without gradients, all cells experience identical resource availability regardless of spatial position, eliminating the carrying capacity that normally limits growth through interior cell starvation. The resulting dynamics violate mass conservation; cells consume nutrients without corresponding spatial depletion creating concentration heterogeneity, effectively generating infinite resources.

This finding establishes a critical distinction between two model component classes: (1) **ablatable biological mechanisms** representing testable hypotheses about behavioral rules (cell migration strategies, phenotypic transition rates, signaling pathways), and (2) **non-ablatable framework requirements** encoding physical laws (mass conservation, diffusion, thermodynamic constraints). We therefore excluded M4 from quantitative pattern comparisons, as the extreme population explosion makes distributional distance metrics ill-defined (Figure 1). Two-panel figure showing population dynamics comparison. Panel A shows population trajectories over 720 hours with baseline model maintaining low populations around 1,000-2,000 cells while M4 variant explodes to 10,000-70,000 cells. Panel B shows histogram of final populations demonstrating complete separation between distributions with 23-fold mean difference.

4.2 Baseline Variability Establishes Empirical Thresholds

Analysis of baseline-to-baseline comparisons revealed substantial stochastic variability: randomly splitting 50 baseline replicates into two groups of 25 and computing Sliced Wasserstein distance yielded mean $SW = 0.364 \pm 0.038$ over 100 random splits. This finding demonstrates that pattern distributions from identical model parameterizations diverge considerably due solely to stochastic agent dynamics, with 95% of baseline comparisons falling between $SW = 0.29$ and $SW = 0.44$. Any ablation producing divergence within this range cannot be confidently distinguished from noise.

The empirically-derived thresholds (auxiliary < 0.402 , intermediate $0.402-0.478$, core > 0.478) differ substantially from arbitrary conventions. In particular, the auxiliary threshold of 0.402 is four times larger than originally proposed (0.1). This calibration proves crucial for correct interpretation; many mechanisms producing

Table 1: Empirical Threshold Calibration from Baseline Variability

Statistic	Value	Interpretation
Baseline splits	100	Random 25 vs 25 samples
Mean SW (μ)	0.364	Central baseline variability
Std Dev (σ)	0.038	Spread of stochastic noise
95% Range	[0.29, 0.44]	Typical noise-level divergence
Auxiliary threshold ($\mu + 1\sigma$)	0.402	84th percentile, conservative
Intermediate threshold ($\mu + 3\sigma$)	0.478	99.7th percentile, high confidence

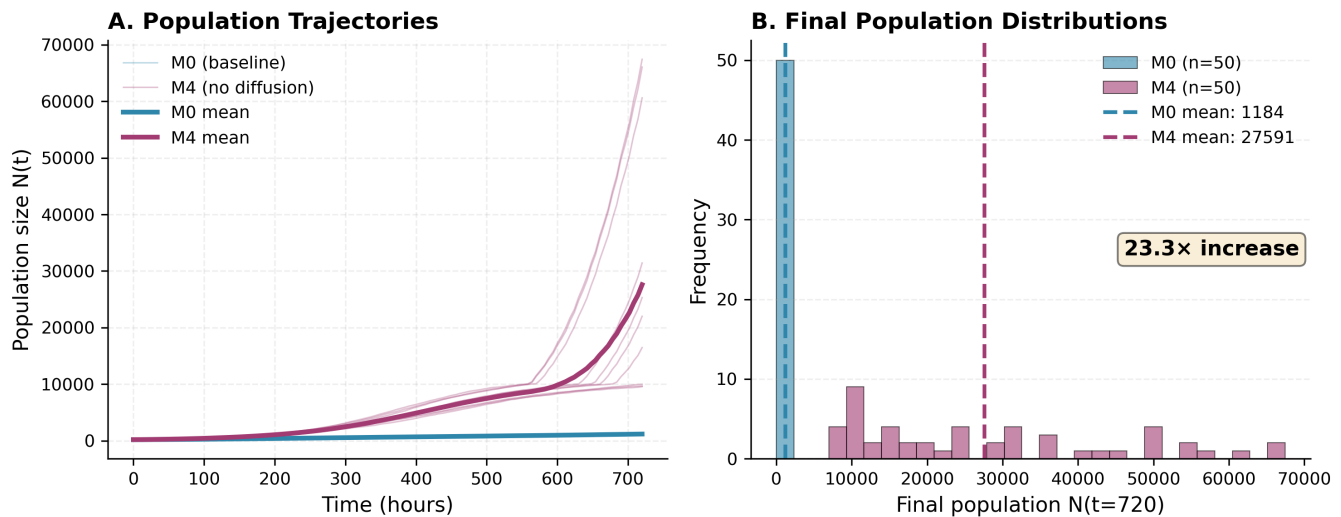


Figure 1: Framework requirements versus ablatable mechanisms. (A) Population trajectories for baseline model M0 (blue, n=50) versus variant M4 removing spatial diffusion (purple, n=50) (B) Final population distributions at t=720 hours.




Figure 2: Pairwise pattern divergence matrix quantified by Sliced Wasserstein distance between all model variants.

divergence in the range 0.25-0.40 would be misclassified without empirical threshold determination.

4.3 Cell Migration is the Sole Empirically Essential Biological Mechanism

Among ablatable biological mechanisms, variant M6 removing cell migration produced catastrophically different patterns with SW distance 32.89 from baseline, exceeding the core threshold by a factor of 69. This mechanism’s removal increased distance from the baseline centroid by 34-fold (312.5 versus 9.1) and achieved statistical significance with $p < 0.0001$ surviving Bonferroni correction. No other single biological mechanism approached this effect magnitude.

The dramatic importance of cell migration [14] derives from multiple interacting factors (Figure 2). Without motility, cells cannot escape nutrient-depleted microenvironments, leading to widespread starvation in tumor interiors despite abundant nutrients at boundaries. Lack of spatial exploration prevents invasion front formation and limits tumor expansion beyond initial seeding volumes.

Phenotypic mixing becomes impossible as resistant and sensitive cells remain segregated by initial spatial positions. These effects compound over the 30-day simulation horizon, producing population crashes (observed mean populations 10-20% of baseline) and fundamentally altered spatial architectures.

The timescale argument reinforces migration’s centrality: cell migration operates on fast timescales (cells traverse multiple spatial units per day) while evolutionary transitions operate on slow timescales (mean transition times $\sim 10^4$ - 10^5 hours). Without fast spatial processes allowing cells to reach favorable microenvironments, slow evolutionary adaptation cannot rescue dying populations.

4.4 Phenotypic Transitions Show Weak Intermediate Effects

Variant M3 removing phenotypic transitions produced SW distance 0.448, falling in the intermediate range (above 0.402 but below 0.478). Statistical testing yielded $p = 0.163$ (Mann-Whitney U) and $p = 0.689$ (Kruskal-Wallis), failing to achieve significance after Bonferroni correction.

The weak effect likely reflects pre-existing heterogeneity sufficiency: with initial resistant cell fractions around 2% (4 resistant cells among 200 initial cells), selection on standing variation drives resistance dynamics over 30-day timescales. Resistant cells possess modest fitness advantages (10% higher proliferation rates) enabling clonal expansion without requiring de novo transitions. This mechanism would likely become essential at longer timescales (90+ days) when initial resistant clones are depleted or at lower initial resistant fractions ($f_R < 10^{-4}$) where standing variation becomes inadequate.

4.5 Six Mechanisms are Genuinely Auxiliary

Mechanisms R5 (paracrine signaling), R7 (density-dependent effects), R8 (metabolic competition), R9 (immune surveillance), R10 (cooperative resistance), and minimal model M22 (simultaneous removal of R7-R10) all produced divergence below the auxiliary threshold with SW distances ranging 0.349-0.387. None achieved statistical significance ($p > 0.069$ for all comparisons). These results demonstrate that six of ten modeled mechanisms contribute negligibly to macroscopic spatiotemporal patterns over 30-day horizons.

Several factors explain this extensive redundancy. First, many explicitly-modeled mechanisms emerge implicitly from simpler components: carrying capacity arises from nutrient gradients and local depletion without requiring explicit density-dependent rules; competition emerges from spatial diffusion and uptake without cell-level Michaelis-Menten kinetics. Second, weak effect sizes at realistic parameter values: immune cells comprise 0% of the population (precluding immune effects), protection factors are modest (0.3), and paracrine coupling is weak (matrix norm 0.01). Third, coarse-grained observables may miss fine-scale effects: our fingerprints aggregate spatial statistics over 2-20 unit scales and temporal snapshots every 6 hours.

The finding that auxiliary mechanisms contribute negligibly at macroscopic scales does not imply biological unimportance. These processes may critically affect unmeasured endpoints (single-cell state distributions, molecular heterogeneity), dominate in alternative parameter regimes (high immune infiltration, strong cooperation), or become essential at different scales (molecular assays

versus tissue imaging). Table 2 summarizes the complete classification of all ten mechanisms based on empirical ablation results.

4.6 Minimal Model Validates Parsimony Principle

The minimal model M22 removing 40% of rules (R7-R10) simultaneously achieved SW distance 0.367 from baseline, falling well within the auxiliary threshold (0.402). Statistical testing yielded $p = 0.132$ (Mann-Whitney U), failing to detect significant differences. Distance from the baseline centroid increased only 7% (9.73 versus 9.11).

This result provides compelling evidence for the parsimony principle: a model retaining six mechanisms (R1-R6) produces patterns statistically indistinguishable from the full ten-mechanism model despite 40% rule reduction. Empirical profiling with 20 replicate full-duration simulations revealed minimal computational gains: M0 required 494 ± 69 seconds per simulation versus M22 at 472 ± 59 seconds, yielding speedup factor 1.048 \times (4.8% faster, Mann-Whitney $p = 0.24$, not statistically significant).

Component-level timing analysis demonstrates that this negligible improvement stems from architectural realities, as detailed in Table 3. PDE solver operations consumed 28-30% of total runtime, agent updates 64-66%, and birth-death processes 6%. The removed rules (R7-R10) represented only 7.3% of agent processing time (24 seconds per simulation), while expensive retained rules dominated agent costs. The theoretical maximum speedup from removing 7.3% of the 70% ABM component is 1.053 \times , closely matching observed 1.048 \times .

Model reduction benefits therefore accrue primarily through enhanced interpretability (focusing on three essential mechanisms), improved parameter calibration (fewer free parameters reducing likelihood surface complexity), and potential for analytical progress rather than computational acceleration.

5 DISCUSSION

5.1 Distinguishing Ablatable Mechanisms from Framework Requirements

The M4 result provides a crucial methodological contribution generalizable beyond this specific case study: systematic ablation studies must distinguish testable behavioral hypotheses from fundamental framework requirements. This distinction generalizes universally across agent-based modeling domains. In **ecological systems**, mass-energy conservation and thermodynamic constraints are framework requirements while specific interaction rules are ablatable mechanisms. In **economic models**, budget constraints and market clearing conditions encode fundamental accounting identities while behavioral decision rules constitute testable hypotheses. In **social simulations**, graph connectivity properties represent framework choices while opinion update mechanisms are ablatable. In **epidemiological models**, mass action contact structures and compartmental conservation represent framework assumptions while specific transmission pathways are ablatable.

The practical implication is that preliminary screening must precede full ablation analysis across all application domains. Testing each proposed ablation with small pilot studies (10-20 replicates

Table 2: Rule Classification Summary

Rule	Mechanism	Variant	SW Distance	95% CI	p-value	Classification
R1	Nutrient-dependent proliferation	(not tested)	-	-	-	Core (definitional)
R2	Drug-induced death	(not tested)	-	-	-	Core (definitional)
R4	Spatial diffusion	M4	-	-	<0.001	Core (framework)
R6	Cell migration	M6	32.89	[31.62, 35.91]	<0.001	Core (empirical)
R3	Phenotypic transitions	M3	0.448	[0.37, 0.55]	0.163	Intermediate
R8	Metabolic competition	M8	0.387	[0.37, 0.63]	0.117	Auxiliary
R5	Paracrine signaling	M5	0.378	[0.37, 0.56]	0.850	Auxiliary
R7-R10	(Combined removal)	M22	0.367	[0.37, 0.62]	0.132	Auxiliary
R10	Cooperative resistance	M10	0.362	[0.36, 0.54]	0.953	Auxiliary
R9	Immune surveillance	M9	0.349	[0.37, 0.54]	0.069	Auxiliary
R7	Density-dependent effects	M7	0.349	[0.37, 0.56]	0.372	Auxiliary

Table 3: Computational Profiling of Full versus Minimal Model

Component	M0 (full model)	M22 (minimal model)	Absolute Savings	Relative Change
PDE solver	140±20 s (28.4%)	141±19 s (30.0%)	-1 s	+0.6%
Agent updates	325±48 s (65.8%)	302±43 s (63.9%)	+24 s	-7.3%
Birth-death	29±4 s (5.8%)	29±4 s (6.1%)	0 s	0%
Spatial indexing	<1 s (0.0%)	<1 s (0.0%)	0 s	0%
Total runtime	494±69 s	472±59 s	+22 s	-4.6%
Speedup factor	-	1.048×	<i>(p</i> = 0.24, not significant)	

examining aggregate observables) can identify pathological variants before investing computational resources in comprehensive pattern analysis.

5.2 Implications for Agent-Based Modeling Practice

Our findings challenge conventional agent-based modeling approaches that maximize mechanistic detail under the assumption that comprehensive mechanism inclusion improves realism [3]. We demonstrate that six of ten modeled mechanisms contribute negligibly to emergent patterns, suggesting widespread redundancy in contemporary agent-based models. This redundancy carries multiple costs: increased parameter dimensionality complicating calibration [8], higher computational expense limiting ensemble sizes, and obscured interpretation as essential mechanisms are buried among auxiliary complexity.

However, minimal models provide substantial non-performance benefits that justify model reduction despite negligible computational gains. First, **exhaustive parameter exploration** becomes feasible: the minimal model with 6 mechanisms involves fewer free parameters enabling finer sampling of parameter space. Second, **improved calibration to empirical data** through reduced likelihood surfaces: fewer parameters create lower-dimensional optimization landscapes with fewer local minima. Third, **enhanced interpretability**: mechanistic understanding focuses on three essential processes rather than ten competing mechanisms. Fourth,

potential for analytical progress: simplified dynamics may permit mathematical analysis via moment closure approximations or asymptotic methods.

5.3 Beyond Minimality: Informed Rule Set Selection

While our framework identifies the minimal rule set sufficient for reproducing observed patterns, practitioners may reasonably prefer non-minimal models. A modeler might retain phenotypic transitions (R3, classified as intermediate with $SW_{i=}$ 0.448) because the mechanism has direct experimental correlates that facilitate communication with domain experts, even though the minimal model omits it without statistical penalty. More generally, if there is reason to suspect that the minimal set’s behavior is sensitive to parameter perturbations (for instance, that small changes to key parameters could shift the system into a qualitatively different regime), retaining additional mechanisms may provide a buffer against such fragility under conditions not yet tested. Our framework supports these decisions quantitatively. The SW distances, confidence intervals, and p-values in Table 2 provide a continuous measure of each mechanism’s contribution rather than a binary keep-or-discard verdict. A modeler targeting maximum parsimony applies the auxiliary threshold strictly; one prioritizing robustness or interpretability can retain intermediate mechanisms at known quantitative cost. The contribution of this work is not to prescribe a single reduced model but to replace intuition-driven mechanism inclusion with evidence-based rule set selection, where the tradeoff between parsimony and other modeling goals is made explicit and informed by data.

5.4 Scale-Dependent Mechanism Importance

A crucial insight generalizes across all agent-based modeling domains: mechanism importance depends on observational scale. In our tumor evolution case study, migration proves essential at tissue scales (millimeters, days-weeks) where spatial organization dominates, while paracrine signaling and phenotypic transitions prove auxiliary despite potential importance at molecular scales. This scale dependence manifests universally across ecological, economic, social, and epidemiological systems.

This scale dependence suggests a universal hierarchy: different mechanisms govern system behavior at different spatiotemporal resolutions. Agent-based models implicitly choose observational scales through their output metrics. Our spatial statistics and population counts emphasize tissue-level patterns, naturally highlighting mechanisms operating at those scales. Alternative metrics might reveal different essential mechanism sets.

5.5 Emergent Phenomena from Minimal Rule Sets

A striking finding is that complex spatiotemporal patterns (resistance foci, necrotic cores, spatial heterogeneity, invasion fronts) emerge from only three essential mechanisms plus spatial diffusion. This observation supports the principle that emergent complexity arises from simple interactions rather than complicated rules [6]. The tumor evolution patterns we observe, which mirror clinical imaging and experimental observations qualitatively, require no explicit programming of tissue architecture, competitive dynamics, or spatial structure. These features emerge automatically from agent-environment coupling.

5.6 Generalizability Across Domains

While demonstrated through tumor evolution, our methodology generalizes to any agent-based system where rule essentiality is uncertain. **Ecological systems** modeling predator-prey dynamics could test whether certain interaction types are dispensable. **Economic agent-based models** incorporating numerous behavioral rules could identify minimal rationality assumptions. **Epidemiological models** with age structure and spatial contact networks could determine essential transmission pathways. **Social simulation models** with opinion dynamics could isolate critical social processes. **Traffic models** could identify essential mechanisms for reproducing congestion patterns.

The general protocol remains consistent across domains: (1) establish baseline variability from repeated simulations to quantify stochastic noise, (2) design comprehensive ablations systematically removing candidate mechanisms, (3) quantify emergent pattern divergence using domain-appropriate distributional metrics, (4) test statistical significance with empirically-calibrated thresholds, (5) classify mechanisms by contribution magnitude relative to noise thresholds.

6 CONCLUSIONS

We have presented a systematic methodology for identifying essential rule sets in agent-based models through computational ablation studies with empirically-calibrated statistical frameworks. Applied to tumor evolution under drug pressure, we established critical

methodological distinctions and empirical findings generalizable across domains. Most importantly, we demonstrated that not all model components are ablatable: removing spatial diffusion caused 23-fold population explosion, revealing it as a physical framework requirement rather than a testable biological hypothesis. Among genuinely ablatable mechanisms, complex spatiotemporal patterns emerge from only three essential processes while six commonly-modeled mechanisms contribute negligibly at the observational scale studied.

Our key methodological contributions generalize across all agent-based modeling domains: (1) distinguishing ablatable behavioral mechanisms from non-ablatable framework requirements through preliminary screening; (2) empirical threshold calibration [10] quantifying baseline stochasticity; (3) composite pattern fingerprints aggregating domain-appropriate metrics [1, 5]; (4) rigorous statistical testing with bootstrap confidence intervals and multiple testing corrections; (5) systematic exploration of model reduction space through strategic ablation design. These contributions advance agent-based simulation methodology by making rule essentiality an empirical question amenable to objective measurement.

The findings carry universal practical implications. Comprehensive mechanism inclusion does not guarantee improved realism and may obscure essential processes beneath auxiliary complexity [8]. Systematic ablation testing should precede claims that specific mechanisms are necessary for observed patterns. Minimal models retaining only essential mechanisms improve interpretability and facilitate parameter calibration while maintaining predictive accuracy. The scale-dependent nature of mechanism importance necessitates matching model complexity to research question scales rather than maximizing detail uniformly.

Future work should extend our methodology in several directions: testing alternative observational scales to map essentiality landscapes across resolutions, systematically varying parameter regimes to identify conditions promoting auxiliary mechanisms to core status, validating minimal models against empirical data, and applying the framework to diverse agent-based systems beyond tumor evolution. The ultimate goal is a mature methodology for agent-based model reduction that balances realism with computational tractability and mechanistic interpretability, advancing multi-agent simulation from intuitive art to rigorous quantitative science through empirically-grounded principles applicable universally.

ACKNOWLEDGMENTS

We thank NVIDIA Corporation for providing computational resources through the NVIDIA Academic Grant Program. The NVIDIA RTX A6000 GPU awarded through this grant enabled the extensive simulations presented in this study.

REFERENCES

- [1] Dhananjay Bhaskar, William Y Zhang, and Ian Y Wong. 2021. Topological data analysis of collective and individual epithelial cells using persistent homology of loops. *Soft matter* 17, 17 (2021), 4653–4664.
- [2] Rafael R Bravo, Etienne Baratchart, Jeffrey West, Ryan O Schenck, Anna K Miller, Jill Gallaheer, Chandler D Gatenbee, David Basanta, Mark Robertson-Tessi, and Alexander RA Anderson. 2020. Hybrid Automata Library: A flexible platform for hybrid modeling with real-time visualization. *PLoS computational biology* 16, 3 (2020), e1007635.

- [3] Arthur Brugière, Doanh Nguyen-Ngoc, and Alexis Drogoul. 2022. Handling multiple levels in agent-based models of complex socio-environmental systems: A comprehensive review. *Frontiers in Applied Mathematics and Statistics* 8 (2022), 1020353.
- [4] Ishan Deshpande, Yuan-Ting Hu, Ruoyu Sun, Ayis Pyrros, Nasir Siddiqui, Sanmi Koyejo, Zhizhen Zhao, David Forsyth, and Alexander G Schwing. 2019. Max-sliced wasserstein distance and its use for gans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10648–10656.
- [5] Ishan Deshpande, Ziyu Zhang, and Alexander G Schwing. 2018. Generative modeling using the sliced wasserstein distance. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3483–3491.
- [6] Jill A Gallaher, Susan C Massey, Andrea Hawkins-Daarud, Sonal S Noticewala, Russell C Rockne, Sandra K Johnston, Luis Gonzalez-Cuyar, Joseph Juliano, Orlando Gil, Kristin R Swanson, et al. 2020. From cells to tissue: How cell scale heterogeneity impacts glioblastoma growth and treatment response. *PLoS computational biology* 16, 2 (2020), e1007672.
- [7] Ahmadreza Ghaffarizadeh, Randy Heiland, Samuel H Friedman, Shannon M Mumenthaler, and Paul Macklin. 2018. PhysiCell: An open source physics-based cell simulator for 3-D multicellular systems. *PLoS computational biology* 14, 2 (2018), e1005991.
- [8] Ryan N Gutenkunst, Joshua J Waterfall, Fergal P Casey, Kevin S Brown, Christopher R Myers, and James P Sethna. 2007. Universally sloppy parameter sensitivities in systems biology models. *PLoS computational biology* 3, 10 (2007), e189.
- [9] Mehrdad Jafari Mamaghani, Mikael Andersson, and Patrik Krieger. 2010. Spatial point pattern analysis of neurons using Ripley’s K-function in 3D. *Frontiers in neuroinformatics* 4 (2010), 1285.
- [10] Jeon-Young Kang and Jared Aldstadt. 2021. Using multiple scale space-time patterns to determine the number of replicates and burn-in periods in spatially explicit agent-based modeling of vector-Borne disease transmission. *ISPRS International Journal of Geo-Information* 10, 9 (2021), 604.
- [11] Aditya Lahiri, Haswanth Vundavilli, Madhurima Mondal, Pranabesh Bhattacharjee, Brian Decker, Giuseppe Del Priore, N Peter Reeves, and Aniruddha Datta. 2023. Drug target identification in triple negative breast cancer stem cell pathways: A computational study of gene regulatory pathways using Boolean networks. *IEEE Access* 11 (2023), 56672–56690.
- [12] Richard Meyers, Melanie Lu, Constantin Waubert De Puiseau, and Tobias Meisen. 2019. Ablation studies in artificial neural networks. (2019).
- [13] Francisco J Rubio, Jorge A Espindola, and José A Montoya. 2023. On near-redundancy and identifiability of parametric hazard regression models under censoring. *Biometrical Journal* 65, 8 (2023), 2300006.
- [14] Luis Saucedo-Mora, Miguel Ángel Sanz, Francisco Javier Montáns, and José María Benítez. 2024. A simple agent-based hybrid model to simulate the biophysics of glioblastoma multiforme cells and the concomitant evolution of the oxygen field. *Computer Methods and Programs in Biomedicine* 246 (2024), 108046.
- [15] Iris HR Yoon, Robert Jenkins, Emma Colliver, Hanyun Zhang, David Novo, David Moore, Zoe Ramsden, Antonio Rullan, Xiao Fu, Yinyin Yuan, et al. 2024. Deciphering the diversity and sequence of extracellular matrix and cellular spatial patterns in lung adenocarcinoma using topological data analysis. *bioRxiv* (2024), 2024–01.