A Computational Agent-Based Model for Cancer–Immune Dynamics

Jiayuan Zhu

Mentor: Nizhum Rahman

Abstract

We develop a computational agent-based model to study cancer immune dynamics in a spatially structured environment. In this framework, individual tumor and immune cells are modeled as discrete agents that move, divide, die, and interact locally over time.

Immune cells can recognize and kill nearby tumor cells. The model also includes periodic immune cell injections to simulate immunotherapy. By incorporating spatial heterogeneity and local interactions, the model allows exploration of how treatment timing and immune influx influence tumor progression. This model enables exploration of tumor–immune interactions and the potential effects of immune cell–based therapies.

1 Introduction

Cancer is a disease that causes millions of deaths worldwide each year [1]. One of the main challenges in cancer treatment is that the effectiveness of therapies can vary greatly from person to person [2]. Over the past few decades, several advanced treatment approaches have been developed, among which immunotherapy stands out as a major breakthrough [3]. Immunotherapy is now used to treat a variety of cancers, including bladder, breast, liver, and lung cancers [4, 5].

Despite this progress, cancer remains one of the leading causes of mortality worldwide and continues to be a primary focus of medical research [6]. Recent advances have revealed that cancer is not a single disease, but rather a complex set of conditions involving intricate multiscale interactions within the tissue environment [7, 8]. Tumor growth is powered by specific oncogenic drivers [9], and identifying these drivers remains a critical challenge for effective therapies.

To unravel this complexity and guide therapeutic development, mathematical oncology has become a well-established field in both the mathematical and biomedical sciences. Mathematical models of cancer generally fall into two categories: continuous and discrete frameworks. Classical continuous models, based on ordinary or partial differential equations, are powerful tools for quantitatively capturing tumor growth kinetics at the macroscopic level [10]. Discrete approaches, such as cellular automaton models and agent-based modeling, provide a more detailed, per-cell description of cancer dynamics [11].

Both approaches, however, have important limitations. Continuous models are often restricted to population-level behavior and can suffer from irreproducibility due to differences in parameter estimation [10]. Discrete models, while capable of simulating cell-level interactions, can struggle to provide quantitative predictions, and frameworks like cellular automata may not smoothly capture chemical gradients or complex cell movements [12].

Agent-based modeling (ABM) in particular has become increasingly popular for exploring tumor–immune dynamics, but many existing ABMs focus on only one cancer type [13] or a specific treatment [14, 15], or simulate basic cellular behavior without considering how therapies influence tumor–immune interactions [16]. Furthermore, there is a lack of agent-based models that systematically address three-dimensional (3D) cancer–immune interactions with treatment interventions, using transparent and biologically meaningful rules.

In this research, we develop a three-dimensional ($3D$) agent-based model of the cancer–immune system that incorporates basic cellular behaviors and T-cell injection therapy [17]. We introduce simple yet biologically meaningful rules for simulating cancer–immune interactions, present key results from our model, and discuss how these findings can inform the development of more rigorous and robust modeling approaches in the future.

2 Methods

Our agent-based model (ABM) is implemented under a three-dimensional on-lattice framework, providing a straightforward ground base for tracking individual cellular behavior [18] and visualizing clearly. Each lattice site can hold a cancer cell and an immune cell.

Fig. 1 is a schematic diagram of our ABM workflow. We use uniform time steps Δt to set the frequency of cellular phenotype behaviors and update the lattice, as a small time step helps lower the relative errors [19]. The total simulation run time T can be chosen freely as desired to demonstrate behavior in different time scales; the number of steps in a single trial is $s = T \,/\, dt$.

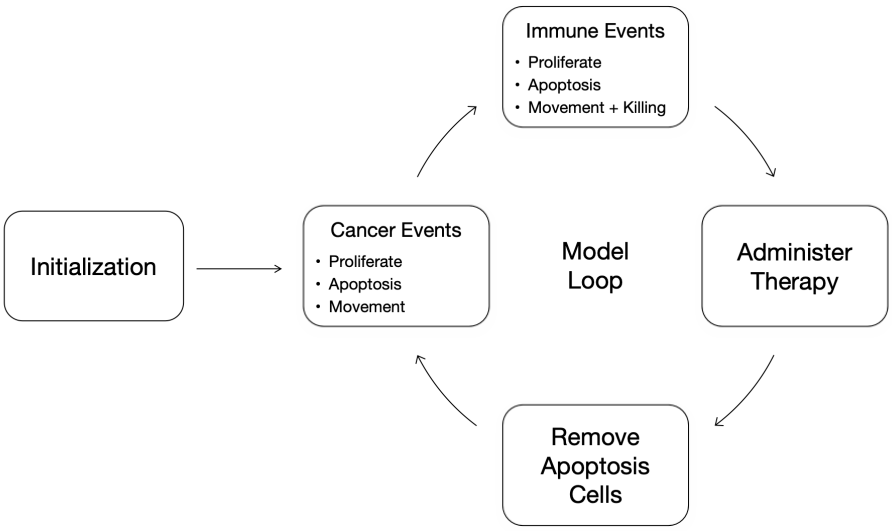

Figure 1: A flowchart describing the simulation algorithm

The Gillespie Stochastic Simulation Algorithm [20] provides a rigorous framework for discrete event implementation relying on the Poisson Process. This gives us a good insight into describing k events occurring in a fixed interval by a probability mass function:

$$\Pr\big(K\big) = \frac{\lambda^k \, e^{-\lambda(x)}}{k!}$$

with an event-specific rate parameter λ. For every small time interval Δt, the probability approximation of a single cell event occurring follows:

$$\Pr\big(K = 1\big) = \lambda \, \Delta t$$

which implies that at most one cell event could take place during Δt. This approximation ensures simplified implementation and computation efficiency [21].

2.1 Cancer cell event

Our agent-based model primarily focuses on certain key hallmarks of cancer [7] and describes them with discrete, rule-based behaviors.

2.1.1 Initialization

The cancer cells are initialized into the lattice with an initial population C_0. A single cell is placed at the geometric center of the grid. Remaining cells are then added around it by selecting empty sites within the Moore neighborhood. This initialization procedure generates a compact tumor mass, closely mimicking the real-life scenario of tissue architecture.

2.1.2 Proliferate

Cancer proliferation event is modeled by the Poisson distribution with a base division rate parameter α_C. If a proliferation event is selected for a specific cancer cell, it will

detect empty spaces in its Moore neighborhood that are not currently occupied by another cancer cell, and a new cancer cell will be generated and randomly occupy one of the empty spots. If there is no space in its Moore neighborhood, no cell proliferation would take place.

2.1.3 Movement

Cancer cells' movement is implemented as a stochastic process at a rate μ_C. While executing movement events, the cell would check for sufficient spaces in its Moore neighborhood. If one or more empty sites are available, the cell relocates to a randomly selected spot; if no adjacent spaces are free, the cell remains in place.

2.1.4 Apoptosis

Cancer cells would undergo apoptosis events based on their natural death rate parameter δ_C as a Poisson process. When an apoptosis event occurs, the cell is immediately removed from the lattice and is excluded from all subsequent events.

2.2 Immune cell event

2.2.1 Initialization

The initial immune cell population is specified by the parameter I_0, which can be set to zero to simulate a therapy-only scenario. While we could still consider some initial immune behavior inside the grid, placed in random spaces. This randomized distribution reflects the reality situation that immune cells patrol and may reside in any location within the human body.

2.2.2 Proliferate

Immune proliferation event is decided by a Poisson distribution using the base rate parameter α_I as its division rate. If a proliferation event is selected for a given immune cell, the cell scans its Moore neighborhood for empty lattice sites (not occupied by any immune or tumor cells) and, if at least one is found, generates a daughter immune cell in one randomly chosen empty spot. If no empty site exists, no proliferation occurs.

2.2.3 Apoptosis

Immune cells undergo apoptosis events based on their natural death rate parameter δ_I. When an apoptosis event occurs, the cell is immediately removed from the lattice and is excluded from all subsequent events.

2.2.4 Movement

Immune movement is a key determinant of immune functionality, since it directly governs the frequency and efficiency of immune–cancer interaction. To systematically investigate its impact on therapeutic performance, we implemented three distinct movement algorithms, each callable on demand, to explore how alternative rules affect killing efficiency.

Random Movement The first immune movement function is simply moving randomly using the rate parameter μ_I. The immune cell will detect all possible empty spots without an existing immune cell and move to one of the spots. If the spot is occupied by a cancer cell, the immune cell will try to kill it with its natural killing rate parameter κ_I [22].

Random Movement with Hunting The second movement function provides a hunting option; each immune cell will first search for adjacent cancer cells in its Moore neighborhood. If one or more cancer cells are detected, the cell executes hunting action based on the killing parameter κ_I. If there is no cancer cell present in the immune neighborhood, the immune cell will choose to move randomly using the rate parameter μ_I, as the first movement function indicates.

Search Movement with Hunting The final movement strategy implements a more advanced and realistic algorithm that conducts both searching and hunting. At each time step, every immune cell in the lattice will search for the cancer cell within the closest distance, and move towards it with the searching rate parameter ξ_I. Once an immune cell has moved into the Moore neighborhood of cancer, it will search for cancer cells locally, and execute hunting action based on the killing rate κ_I. This strategy provides a solution to long-range targeting along with close-range killing efficiency.

2.3 Therapy Implementation

The primary objective of this agent-based model is to quantify and assess how immunotherapy alters tumor dynamics. We emulate CAR T-cell treatment by periodically injecting immune cells into the spatial lattice and examining the population dynamics. Two dosing protocols, fixed-dose and tumor-burden-proportional, are applied under otherwise identical conditions to compare their effects on treatment efficacy.

Fixed-dose Injections A constant quantity of immune cells, N_{inj}, is administered at uniform time intervals of length Δ_I. At each injection time $T_{\text{inj}} = t\ \Delta_I$, N_{inj} cells are injected into randomly selected, unoccupied lattice sites. Immediately upon injection, these agents participate in all immune events based on the given rules and parameters.

Tumor-adaptive Injections At each injection time $T_{\text{inj}} = t\ \Delta_I$, the number of immune cells injected $N_{\text{inj}}(t)$ is scaled proportionally to the current cancer burden. Let $N_C(t)$ denote the cancer cells population at time t and the proportional injection rate by

R_I, then

$$N_{\text{inj}}(t) \;=\; \lceil R_I \times N_C(t) \rceil$$

Immune cells are injected into random empty sites.

3 Results

3.1 Experimental design

We aim to systematically explore how therapeutic dosing impacts cancer–immune dynamics. Specifically, for each of the three immune movement algorithms, we evaluate two dosing options with different dosing sizes: fixed-dose and proportional-dose, which are administered every two days. All other model parameters remain strictly identical to those listed in **Table 1**.

Name	Description	Value	Source/Notes
C_0	Initial cancer cell count	100	Assumption
α_C	Cancer proliferation rate (min^{-1})	2 / 1440	Talkington and Durrett (2015)
μ_C	Cancer movement rate (min^{-1})	2 / 1440	Bergman et al. (2024)
δ_C	Cancer apoptosis rate (min^{-1})	0.005 / 1440	Bergman et al. (2024)
I_0	Initial immune cell count	50	Assumption
Δ_I	Immune dose schedule (min)	2×1440	Chosen
α_I	Immune proliferation rate (min^{-1})	0.5 / 1440	Bergman et al. (2024)
μ_I	Immune movement rate (min^{-1})	2 / 1440	Bergman et al. (2024)
ξ_I	Immune search rate (min^{-1})	6 / 1440	Assumption
κ_I	Immune kill rate (min^{-1})	6 / 1440	Bhat and Watzl (2007)
δ_I	Immune apoptosis rate (min^{-1})	0.5 / 1440	Bergman et al. (2024)
Δt	Time step (min)	15	Bergman et al. (2024)
T	Total simulation time (min)	35×1440	Chosen
N	Grid size (μm)	35	Assumption

Table 1: Fixed parameters for the agent-based model

3.2 Random Move

We begin by characterizing the intrinsic stochastic dynamics of the immune compartment in the absence of any directed targeting: immune cells execute unbiased random walks and only eliminate tumor cells upon chance spatial overlap.

In our first experiment, we deliver periodic injections of 3000 immune cells, which produces the saw-tooth-like pattern indicating the exact size oscillation of the immune cell count (**Fig.2**). The transient rises of immune cell counts would increase the possibility of immune cancer contact and thus reduce the cancer population. Therefore, there is a

small attenuation of cancer cells after each injection. We observe that shortly after each decrease, the cancer rebounds to the equilibrium at the carrying capacity. The narrow 95% confidence interval around the cancer trajectory reflects very low inter-run variability in tumor growth. These phenomena indicate that only higher dosing would enhance the immune efficacy.

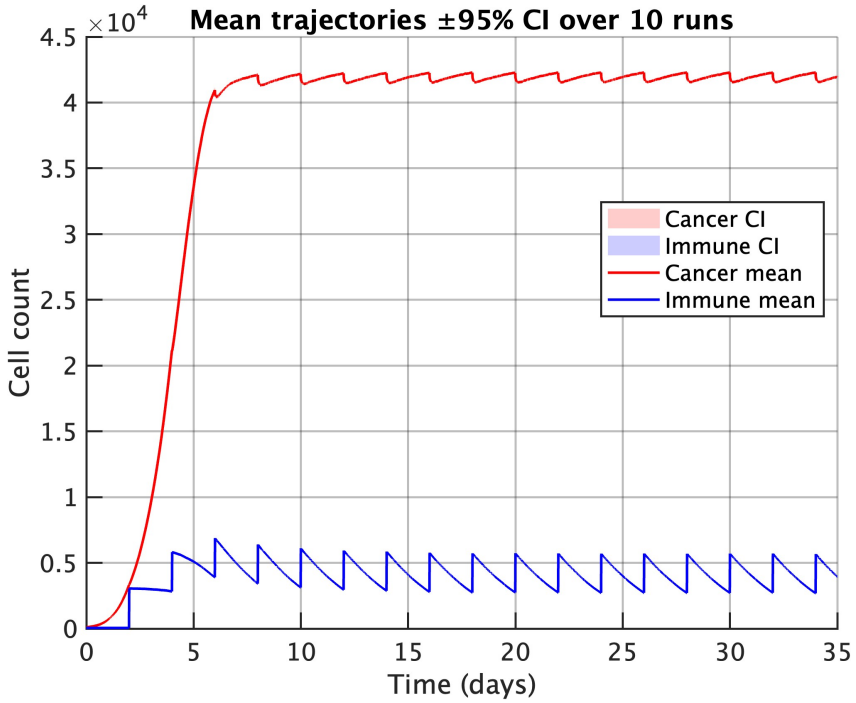

Figure 2: Periodic injections of 3000 immune cells every two days produce a saw-tooth pattern in the immune cell population, yet tumor cells continue to dominate, indicating treatment failure.

In the next set of simulations, we increased the periodic dosing to 5000 immune cells. As shown in **Fig.3**, this higher injection dramatically altered the outcome: after the fifth injection, the cancer started to decline rapidly and was completely eradicated within 3 weeks. We could notice that the interval between the second and third week displayed a wider 95% confidence interval, indicating the randomness of immune movement leads to varied cancer clearance times. Despite this heterogeneity, the overall trajectory retains the same saw-tooth structure. We could conclude that with a sufficiently large dose, the cancer treatment would work.

A remaining concern is that, under our current protocol, immune-cell injections continue even after the tumor has been cleared—an excess that could be deleterious in vivo. To mitigate this risk, we introduce a proportional dosing therapy in which each injection size is dynamically adjusted to the prevailing tumor population.

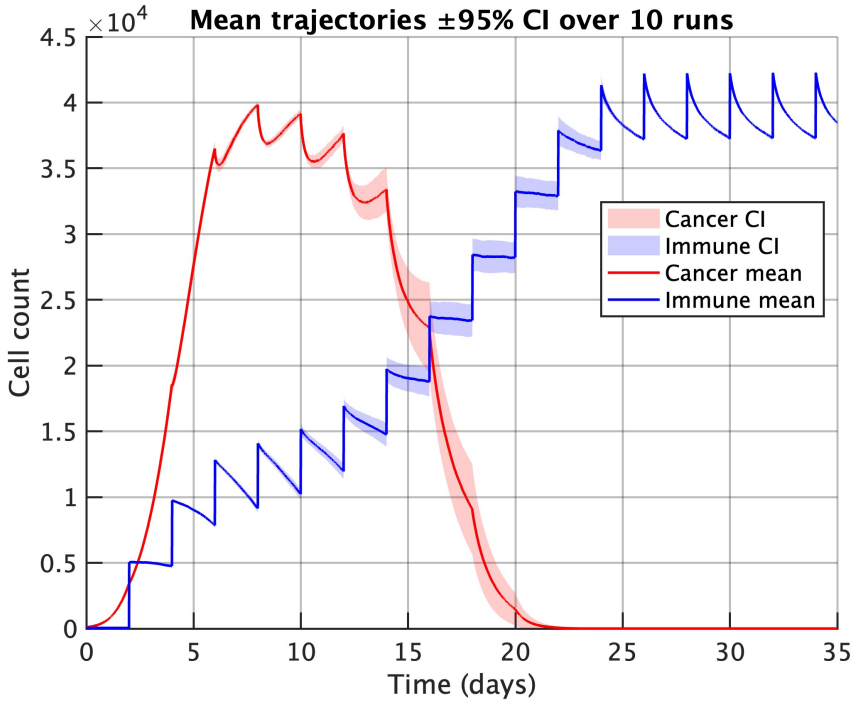

Figure 3: Periodic injections of 5000 immune cells every two days boosted the immune cell population and achieved complete tumor elimination by the third week.

We then implement a proportional-dosing, administering immune cells equal to 10% of the prevailing tumor burden every two days. The resulting trajectories still exhibit the saw-tooth-like oscillations in both immune and cancer cell counts (**Fig.4**). Crucially, after each injection, the tumor never rebounds fully to the carrying capacity; instead, it declines incrementally with each dose. This adaptive strategy dynamically scales the infusion to tumor size and produces a cumulative reduction in cancer burden, demonstrating clear advantages over fixed-dose regimens.

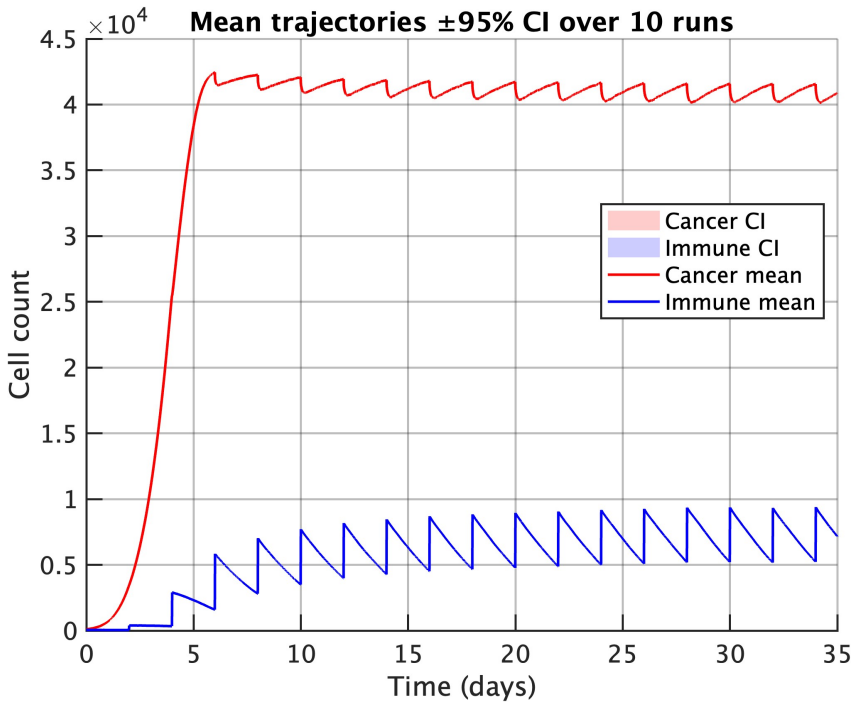

Figure 4: Proportional dosing at a rate of 0.1 produces high-amplitude saw-tooth oscillations in the immune cell population, while the tumor remains near its carrying capacity.

We further raised the dosing rate to 15%. As a result, the saw-tooth profile in the immune trajectory (blue) now features even higher peaks and steadily elevated troughs (**Fig.5**). From the second week, the cancer population decreased drastically until extinction, meaning the treatment has exhibited a robust outcome. The transient widening of the 95% confidence band during this collapse reflects stochastic variability in the exact clearance time, whereas the narrow CI before and after indicates highly reproducible dynamics. These results demonstrate that a proportional rate of 0.15 not only avoids post-clearance overshoot but also reliably tips cancer–immune balance toward extinction under purely random immune motility.

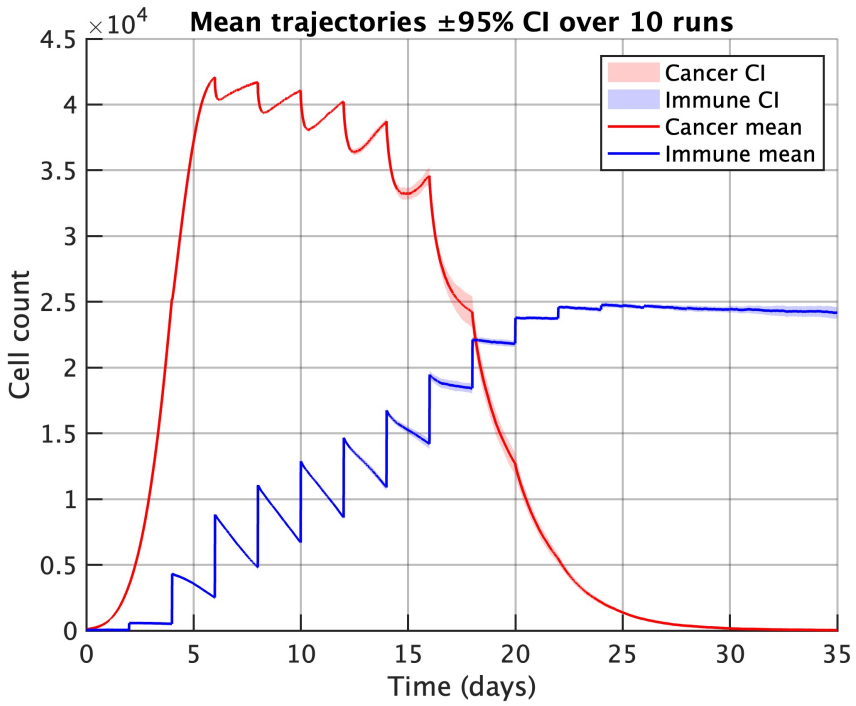

Figure 5: Proportional dosing at a rate of 0.15 achieved complete tumor elimination by 30 days while avoiding immune overshoot.

3.3 Random Hunt

We now discuss therapy under Moore's neighbor hunting combined with random movements. We still assess the same per-injection dosing of 3000 immune therapy. As shown in **Fig.6**, the cancer-immune dynamics showed identical behavior as results from random movement cases. Although Moore's neighbor hunting yields slightly larger post-injection peaks in the immune population (blue) compared with pure random motility, it does not materially improve tumor control at this dose.

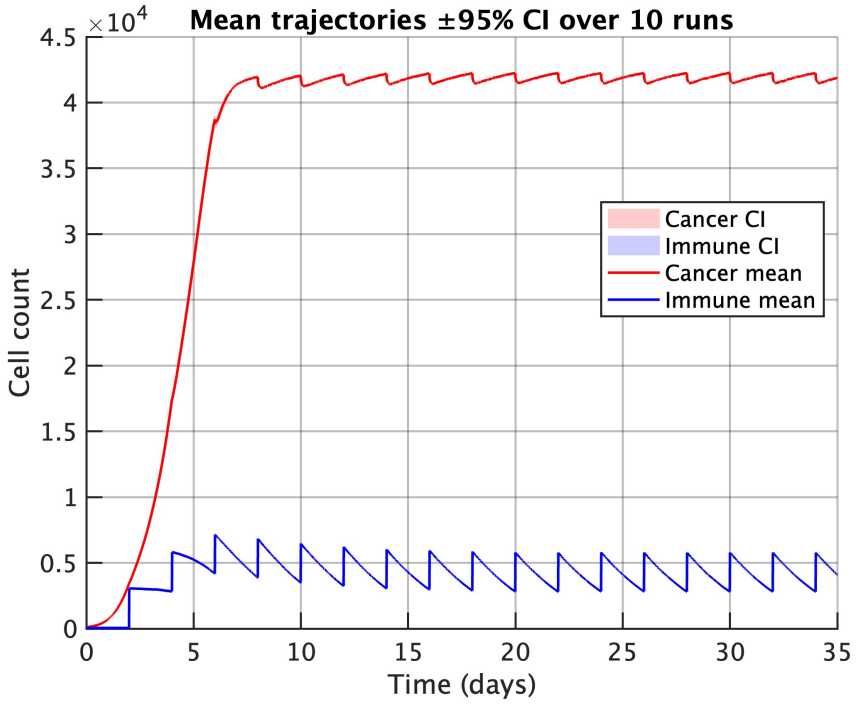

Figure 6: Periodic injections of 3000 immune cells still produce a saw-tooth pattern in the immune cell population under hunting rules, yet tumor cells continue to dominate.

We then increased the per-injection dosing to 4000 immune cells, and effective cancer eradication was achieved between the third and fourth week (**Fig.7**). It is noticeable that the decreasing cancer trajectory displayed a wide 95% confidence band, indicating a high uncertainty of the timing of effective treatment.

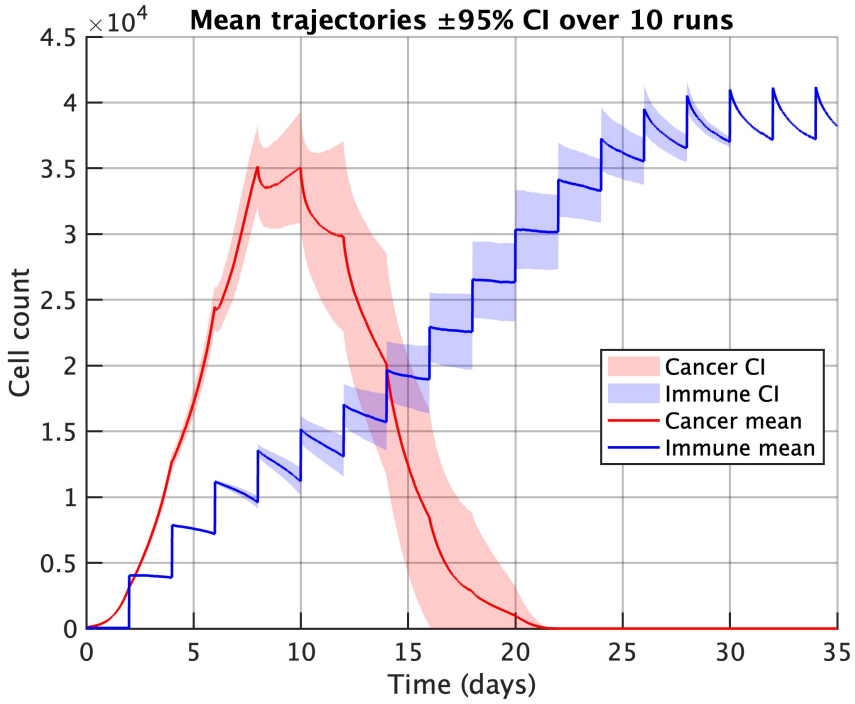

Figure 7: Utilizing the random move and hunting algorithm, dosing of $4,000$ immune cells achieves comparable tumor clearance to the non-hunting one with reduced total dosing, resulting in complete elimination within three weeks.

The proportional-dosing that administers immune cells equal to 10% of the prevailing tumor burden in a hunting situation (**Fig.8**) didn't show a significant difference from the previous results in **Fig.4**. This indicates that for cancer dominance scenarios, different killing efficiencies don't make a great difference. The dosing would need to be increased for effective treatment.

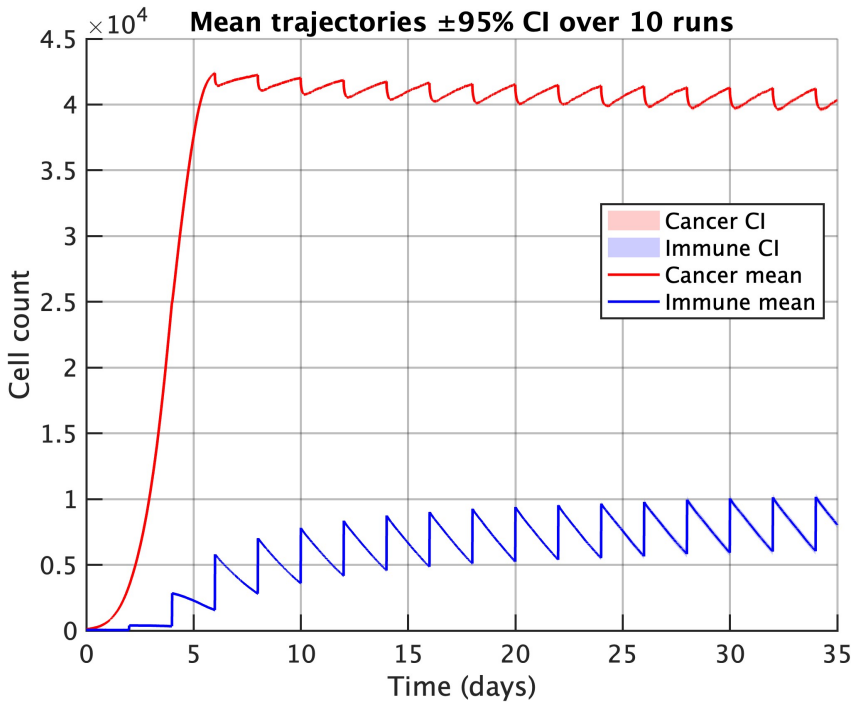

Figure 8: Proportional dosing at a rate of 0.1 produces high-amplitude saw-tooth oscillations in the immune cell population under the hunting algorithm, while the tumor remains near its carrying capacity.

We further raised the dosing rate to 15%. As expected, the therapy effectively killed cancer within the third week (**Fig.9**). It is notable that, compared to the simple random movement case in **Fig.5**, there doesn't exist any significant upward turning point after two weeks. The treatment worked efficiently with a monotonic decrease in the cancer population. The transient widening of the 95% confidence band during this clearance window reflects variability in the exact eradication time, whereas the narrow 95% confidence band before and after indicates the reproducibility of both the growth phase and the post-clearance immune stage. These findings demonstrate that, under a hunting strategy, a 0.15 proportional rate reliably tips the balance toward tumor extinction without overshooting, thus avoiding unnecessary post-clearance dosing.

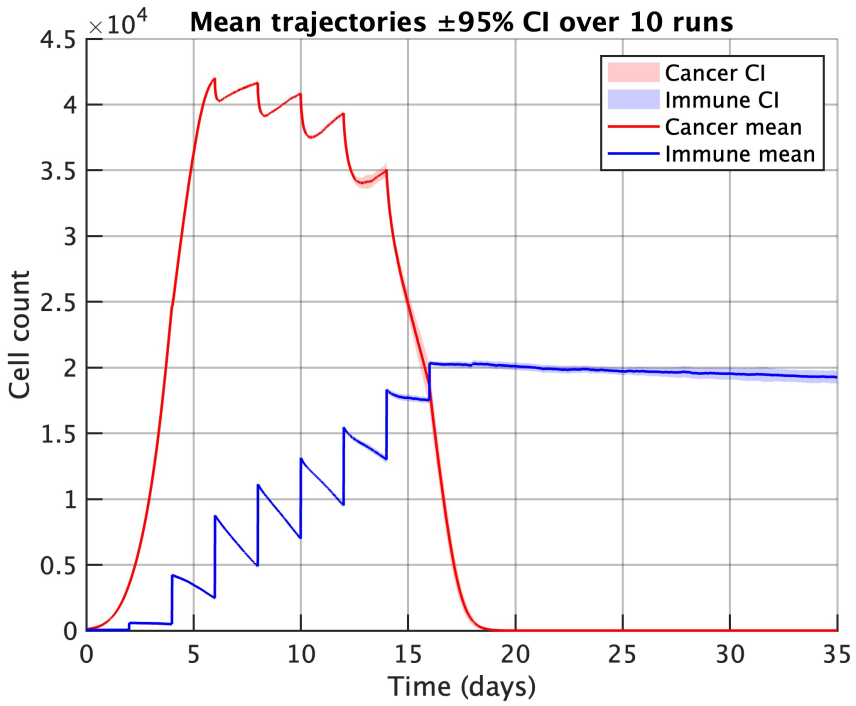

Figure 9: Proportional dosing at a rate of 0.15 achieves complete tumor clearance by 18 days while preventing immune overshoot under random movement and hunting algorithms, outperforming the pure random-movement function.

3.4 Search Hunt

In this final set of simulations, we endowed immune cells with an autonomous search and hunt phenotype. At each time step, immune cells search for the closest cancer cell rather than executing an unbiased random walk, and then hunt in Moore's neighborhood. This directed-search behavior markedly increases cancer-immune encounter rates, allowing us to achieve effective tumor control with a relatively smaller dose. Accordingly, we first tested a fixed-dose therapy of 1000 cells every two days.

As illustrated in **Fig.10**, this 1000 cells therapy with directed movement significantly delays tumor progression. Prior simulations saw the cancer population reach its carrying capacity within one week; here, it was only attained after approximately two weeks. The immune trajectory retains its characteristic saw-tooth-like shape, and the cancer curve exhibited only a very slight downward drift at late times. The wide 95% confidence interval around the tumor trajectory reflects high stochastic variability in both growth and incidental killing events at this low dose. However, there do exist some trials that the immune system kills cancer cells over time, causing this wide 95% confidence interval. But these sporadic successes don't prove that the 1000 immune cells per dose treatment are reliable. A higher per-dose injection is therefore required to ensure consistent tumor eradication.

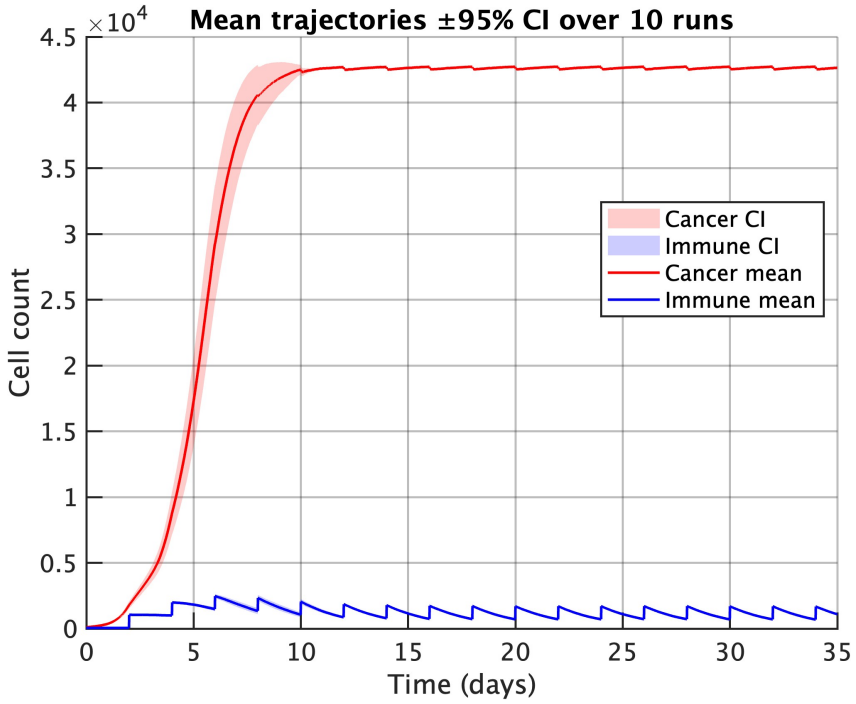

Figure 10: Periodic injections of 1000 immune cells using the search-and-hunt algorithm delay tumor progression toward carrying capacity but fail to eliminate the tumor.

We doubled the dose to 2000 immune cells per injection. As in **Fig.11**, cancer rises to a peak of around 3000 cells but then undergoes a rapid, monotonic decline until a thorough clean-out. Notice the continued dosing post-clearance leads to substantial immune overshoot.

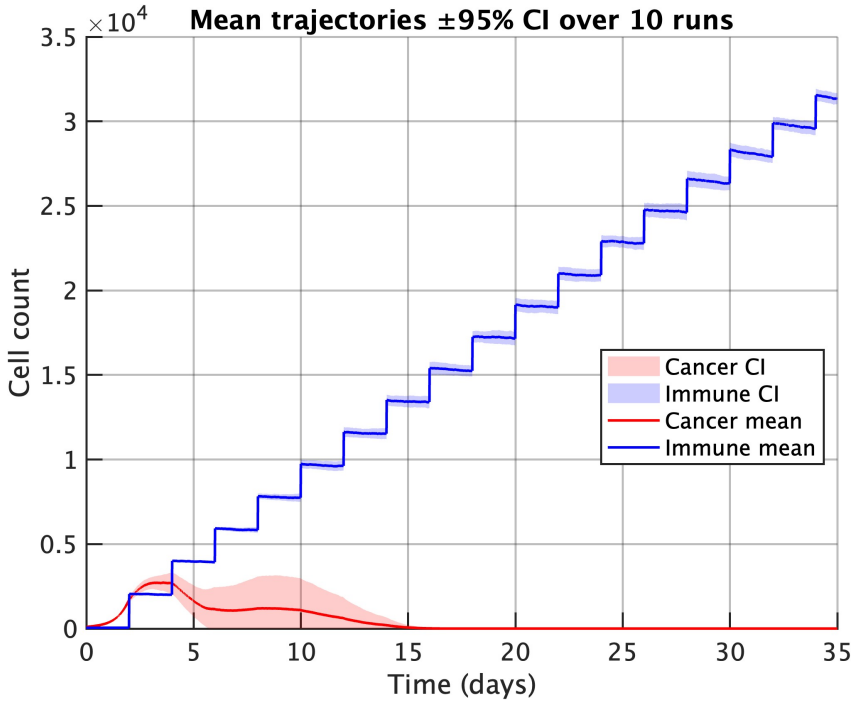

Figure 11: Periodic injections of 2000 immune cells using the search-and-hunt algorithm kept the tumor population below 3000 cells and achieved full clearance within two weeks, but caused significant immune overshoot.

We next evaluated the 10% proportional-dosing under the search-and-hunt phenotype. The immune-cell trajectory retained its upward-going saw-tooth-like shape (**Fig.12**). The cancer population followed exponential growth to carrying capacity within the first week, then entered a sustained decline during weeks 2–3. Although it did not drive the tumor to extinction, the therapy produced a clear and positive outcome starting from the third week. This therapy could eliminate cancer if the therapy were continued for another few weeks.

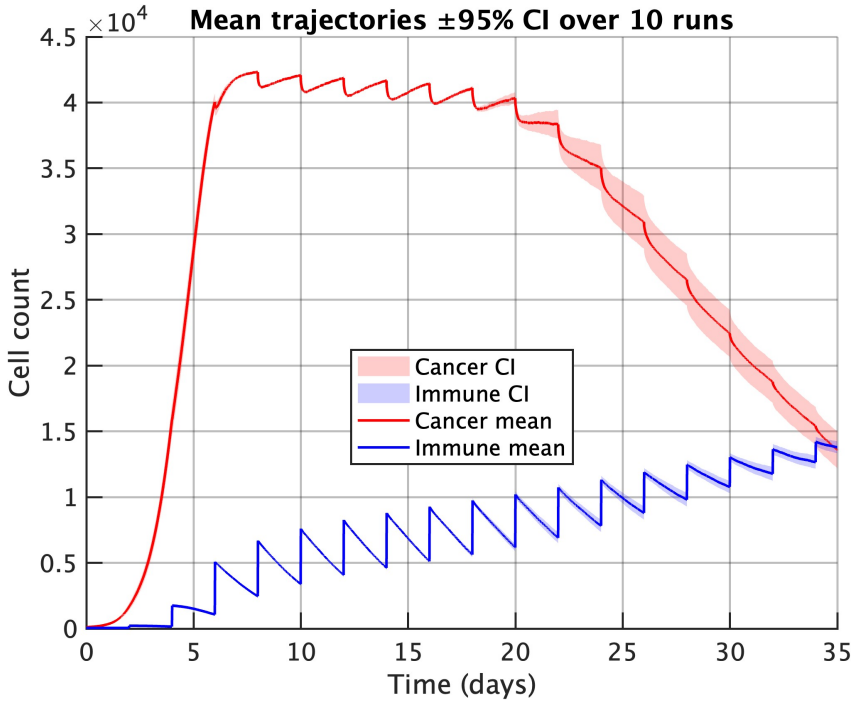

Figure 12: Proportional dosing at a rate of 0.1 with the search-and-hunt algorithm gradually reduces the tumor after three weeks and shows a trend toward complete clearance.

To improve therapeutic efficacy, we increased the dosing rate up to 0.15%. A notable difference, as shown in **Fig.13**, as soon as the cancer population reached the carrying capacity, the curve displayed rapid, monotonic decline until complete eradication by the fourth week.

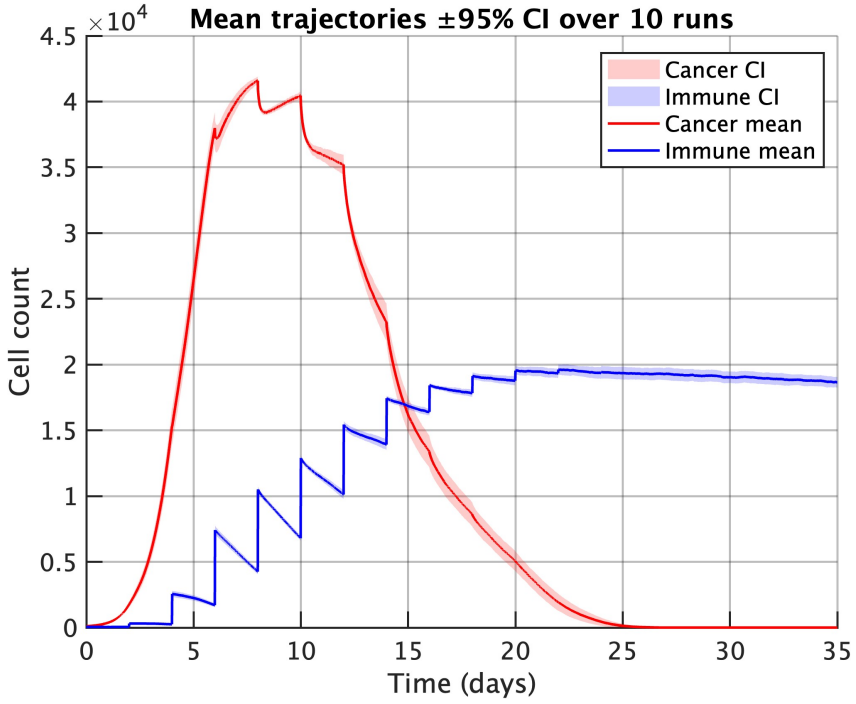

Figure 13: Proportional dosing at a rate of 0.15 with the search-and-hunt algorithm achieves complete tumor elimination within 25 days while avoiding immune overshoot.

4 Discussion

In this report, we systematically evaluated how three key factors: immune-cell motility, treatment strategy, and dosing magnitude shape the spatial tumor–immune dynamics model. By isolating each parameter and examining its impact, we identified clear principles for designing more effective immune therapies.

First, we found that the immune-cell movement phenotype critically determines search efficiency. As shown in the figures and analysis in **Part. 3**, we can understand that random movement, random movement with hunting, and search movement with hunting displayed a gradual increase in efficiency.

The pure random movement was only designed to study how stochastic behaviors are implemented in agent-based models. Under an unbiased random walk, immune cells locate tumor cells only by chance, requiring either very high doses or prolonged treatment to achieve any tumor suppression. Introducing a hunt behavior that each immune cell thoroughly checks its Moore neighborhood before moving on could substantially improve local clearance and resolve the inefficiency, reducing the fixed-dose threshold from 5000 to 4000 cells. The autonomous search-and-hunt strategy managed to let the immune cells work most efficiently: always locate the nearest cancer cell. This is more realistic as the receptor of T cells in an actual biological context would search for the closest cells and try to kill them. The search and hunt movement yields the greatest efficiency, eradicating the malignancy within two weeks with just 2000 cells per dose and inducing a net tumor decline even at a 10% proportional rate over longer cycles.

Next, we compared two different treatment strategies: fixed-dose versus tumor-adaptive dosing. We have experimented that both could work with enough dosing amount, but they showed very distinct features.

Fixed-dose therapy produces predictable and reassuring dynamics that can cure if the per-dose count exceeds a sharp threshold. But these results exhibit wide confidence interval variability, indicating that the result of treatment would vary from time, and different trials could even end up differently. This increases the uncertainty of the therapy. Therefore, it will not be a good idea to stop the dose too early, which leads to another remaining concern: continued dosing post-clearance leads to substantial immune overshoot, generating a potential threat to health.

In contrast, proportional dosing scales each injection to tumor burden, yielding tight confidence bands and avoiding unnecessary overshoot, indicating a very assured and replicable treatment. However, because the injected fraction is small during early exponential growth, the tumor still reaches carrying capacity within the first week. Allowing cancer to grow to carrying capacity would be an undesirable scenario in vivo, as this implies expanded tumor into the whole body in real-life cases.

Finally, the dosing magnitude is the most dominant factor in the research across all scenarios, as it directly affects the equilibration of the tumor–immune system balance. Only by exceeding a critical per-dose cell number or proportional rate—roughly 5000 cells (random walk), 4000 cells (hunt), or 2000 cells (search-and-hunt), or a 15% rate—can immune pressure consistently drive tumor extinction. While motility enhancements and adaptive scheduling refine treatment safety and efficiency, they cannot replace the need for sufficient dosing to achieve a reliable cure.

5 Conclusion

We introduce a novel agent-based model of the cancer–immune dynamics, parameterized to capture diverse cell phenotypes and therapeutic schemes. By systematically varying and studying immune-cell motility, treatment timing, and dosing magnitude, we demonstrate that enhanced search efficiency lowers the minimum effective therapy threshold and accelerates tumor clearance; that distinct dosing protocols yield markedly different dynamical responses; and that only when the per-dose cell number or proportional infusion rate exceeds a critical point can immune pressure successfully suppress tumor growth.

These insights pave the way for a modular agent-based modeling framework that enables interactive tuning of cellular behaviors and microenvironmental conditions and offers quantitative guidance for immunotherapy design. By comparing with existing models in the field, we must, however, transparently acknowledge our model's limitations.

As noted in Section 1, cancer is inherently a multi-scale disease. The future of computational oncology lies in hybrid modeling that integrates discrete and continuous frameworks with physics-based, data-driven, and optimization methods [16]. Our model can integrate features from diverse models to present hybrid phenotypes, thereby improving predictive accuracy and capturing cellular mobility.

While lattice-based models offer straightforward implementation, they are susceptible to grid artifacts [23] and may inadequately capture biomechanical realism [18]. Off-lattice simulations could serve as complementary approaches to validate model fidelity and assess computational performance.

Our present model emphasizes general cellular and therapeutic behaviors; however, to advance toward personalized medicine, it must be adapted to specific cancer types and account for tumor heterogeneity [24]. Rigorous integration of lab data and quantitative methods will be required for clinical translation [13].

Despite these limitations, our agent-based framework demonstrates significant potential for addressing the above challenges. Leveraging the inherent strengths of agent-based approaches, this platform represents a powerful and robust tool for oncology research and the development of novel therapeutic strategies.

References

[1] G. B. of Disease Cancer Collaboration, "Global, regional, and national cancer incidence, mortality, years of life lost, years lived with disability, and disability-adjusted life-years for 29 cancer groups, 1990 to 2017: A systematic analysis for the global burden of disease study," *JAMA Oncology*, vol. 5, pp. 1749–1768, 12 2019.

[2] T. P. Hanna, W. D. King, S. Thibodeau, M. Jalink, G. A. Paulin, E. Harvey-Jones, D. E. O'Sullivan, C. M. Booth, R. Sullivan, and A. Aggarwal, "Mortality due to cancer treatment delay: systematic review and meta-analysis," *BMJ*, vol. 371, 2020.

[3] S. Farkona, E. P. Diamandis, and I. M. Blasutig, "Cancer immunotherapy: the beginning of the end of cancer?," *BMC medicine*, vol. 14, no. 1, pp. 73–73, 2016.

[4] P. Sharma and J. P. Allison, "The future of immune checkpoint therapy," *Science (American Association for the Advancement of Science)*, vol. 348, no. 6230, pp. 56–61, 2015.

[5] A. Ribas and J. D. Wolchok, "Cancer immunotherapy using checkpoint blockade," *Science (American Association for the Advancement of Science)*, vol. 359, no. 6382, pp. 1350–1355, 2018.

[6] R. Baskar, K. A. Lee, R. Yeo, and K. W. Yeoh, "Cancer and Radiation Therapy: Current Advances and Future Directions," *Int J Med Sci*, vol. 9, no. 3, pp. 193–199, 2012.

[7] D. Hanahan and R. A. Weinberg, "Hallmarks of Cancer: The Next Generation," *Cell*, vol. 144, no. 5, pp. 646–674, 2011.

[8] S. Stephan, S. Galland, O. L. Narsis, K. Shoji, S. Vachenc, S. Gerart, and C. Nicolle, "Agent-based approaches for biological modeling in oncology: A literature review," *Artificial Intelligence in Medicine*, vol. 152, p. 102884, 2024.

[9] D. Quail and J. Joyce, "Microenvironmental regulation of tumor progression and metastasis," *Nat Med*, vol. 19, pp. 1423–1437, 2013.

[10] S. Benzekry, C. Lamont, A. Beheshti, A. Tracz, J. M. L. Ebos, L. Hlatky, and P. Hahnfeldt, "Classical Mathematical Models for Description and Prediction of Experimental Tumor Growth," *PLoS Comput Biol*, vol. 10, no. 8, p. e1003800, 2014.

[11] C. A. Valentim, J. A. Rabi, and S. A. David, "Cellular-automaton model for tumor growth dynamics: Virtualization of different scenarios," *Comput Biol Med*, vol. 153, p. 106481, 2023.

[12] K. A. Rejniak and A. R. A. Anderson, "Hybrid models of tumor growth," *WIREs Syst Biol Med*, vol. 3, no. 1, pp. 115–125, 2011.

[13] D. R. Bergman, Y. Wang, E. Trujillo, A. A. Fernald, L. Li, A. T. Pearson, R. F. Sweis, and T. L. Jackson, "Dysregulated FGFR3 signaling alters the immune landscape in bladder cancer and presents therapeutic possibilities in an agent-based model," *Frontiers in Immunology*, vol. 15, p. 1358019, 2024.

[14] M. N. G. van Genderen, J. Kneppers, A. Zaalberg, E. M. Bekers, A. M. Bergman, W. Zwart, and F. Eduati, "Agent-based modeling of the prostate tumor microenvironment uncovers spatial tumor growth constraints and immunomodulatory properties," *npj Systems Biology and Applications*, vol. 10, p. 20, 2024.

[15] H. V. Jain, K.-A. Norton, B. B. Prado, and T. L. Jackson, "SMoRe ParS: A novel methodology for bridging modeling modalities and experimental data applied to 3D vascular tumor growth," *Front. Mol. Biosci.*, vol. 9, p. 1056461, 2022.

[16] I. M. Chamseddine and K. A. Rejniak, "Hybrid modeling frameworks of tumor development and treatment," *WIREs Syst Biol Med*, vol. 12, no. 1, p. e1461, 2020.

[17] A. Holzinger, M. Barden, and H. Abken, "The growing world of CAR T cell trials: a systematic review," *Cancer Immunol Immunother*, vol. 65, no. 12, pp. 1433–1450, 2016.

[18] J. Metzcar, Y. Wang, R. Heiland, and P. Macklin, "A Review of Cell-Based Computational Modeling in Cancer Biology," *JCO Clin Cancer Inform*, vol. 3, pp. 1–13, 2019.

[19] A. Ghaffarizadeh, R. Heiland, S. H. Friedman, S. M. Mumenthaler, and P. Macklin, "PhysiCell: An open source physics-based cell simulator for 3-D multicellular systems," *PLoS Comput Biol*, vol. 14, no. 2, p. e1005991, 2018.

[20] D. T. Gillespie, "A general method for numerically simulating the stochastic time evolution of coupled chemical reactions," *Journal of Computational Physics*, vol. 22, no. 4, pp. 403–434, 1976.

[21] B. Colyer, M. Bak, D. Basanta, and R. Noble, "A seven-step guide to spatial, agent-based modelling of tumour evolution," *Evolutionary Applications*, vol. 17, no. 5, p. e13687, 2024.

[22] R. Bhat and C. Watzl, "Serial killing of tumor cells by human natural killer cells – enhancement by therapeutic antibodies," *PLOS ONE*, vol. 2, pp. 1–7, 03 2007.

[23] P. Macklin, M. E. Edgerton, V. Cristini, and J. Lowengrub, "Discrete cell modeling," in *Multiscale Modeling of Cancer: An Integrated Experimental and Mathematical Modeling Approach*, pp. 88–122, Cambridge: Cambridge University Press, 2010.

[24] M. J. Kim, R. J. Gillies, and K. A. Rejniak, "Current advances in mathematical modeling of anti-cancer drug penetration into tumor tissues," *Front. Oncol.*, vol. 3, p. 278, 2013.