BioSystems 199 (2021) 104290

Contents lists available at ScienceDirect

BioSystems

journal homepage: www.elsevier.com/locate/biosystems

Harnessing adaptive novelty for automated generation of cancer treatments ☆

Igor Balaz [a,*], Tara Petrić [a], Marina Kovacevic [b], Michail-Antisthenis Tsompanas [c], Namid Stillman [d]

[a] Laboratory of Meteorology, Biophysics and Physics, Faculty of Agriculture, University of Novi Sad, Serbia
[b] Department of Chemistry, Biochemistry, and Environmental Protection, Faculty of Sciences, University of Novi Sad, Serbia
[c] Unconventional Computing Laboratory, University of the West of England, UK
[d] University of Bristol, UK

ARTICLE INFO

Keywords:
Evolutionary innovations
Nanomedicine
Cancer
Agent-based

ABSTRACT

Nanoparticles have the potential to modulate both the pharmacokinetic and pharmacodynamic profiles of drugs, thereby enhancing their therapeutic effect. The versatility of nanoparticles allows for a wide range of customization possibilities. However, it also leads to a rich design space which is difficult to investigate and optimize. An additional problem emerges when they are applied to cancer treatment. A heterogeneous and highly adaptable tumour can quickly become resistant to primary therapy, making it inefficient. To automate the design of potential therapies for such complex cases, we propose a computational model for fast, novelty-based machine learning exploration of the nanoparticle design space. In this paper, we present an evolvable, open-ended agent-based model, where the exploration of an initially small portion of the given state space can be expanded by an ongoing generation of adaptive novelties, whenever the simulated tumour makes an adaptive leap. We demonstrate that the nano-agents can continuously reshape themselves and create a heterogeneous population of specialized groups of individuals optimized for tracking and killing different phenotypes of cancer cells. In the conclusion, we outline further development steps so this model could be used in real-world research and clinical practice.

1. Introduction

Malignant tumours are heterogeneous structures that can rapidly acquire drug resistance (Holohan et al., 2013). Therapies can lead to substantial regressions but the effect is often short-lived and cancer evolves to become resistant within a few months (Shaffer et al., 2017; Bozic et al., 2013). Resistance can arise via two main mechanisms: (i) selection of pre-existing resistant cells and (ii) the emergence of resistant cancer cells that continue to evolve under selective pressure (Hu et al., 2017). To overcome therapy resistance, clinical practices are using the recently emerging approach of combinatorial therapy — combining multiple drugs in order to synergistically eliminate the various clones that emerge in a tumour (Chen and Lahav, 2016). However, the problem remains how to properly design combinatorial therapy: which drugs to choose, whether they should be applied sequentially or simultaneously, appropriate time intervals between treatments, and the correct concentration of each drug. Choosing the appropriate target is an additional issue. Recent studies have shown that phenomena such as drug resistance and metastasis are sustained due to so-called cancer stem cells (CSCs) (Pattabiraman and Weinberg, 2014). These cells represent a small subpopulation of cancer cells that have an unlimited capacity for self-renewal, differentiation, and tumourigenesis. Moreover, they are insensitive to most of the conventional therapies, and the percentage of CSCs within a tumour often increases after chemotherapy (Kurtova et al., 2014). Finally, in designing optimal therapies, the drug's side effects should also be considered (Fanciullino et al., 2013).

The use of nanoparticles (NP) as drug carriers can overcome the limitations associated with traditional drug therapy. Due to the versatility of nanoparticles and possibilities of customization, they can improve site-specific targeting of drugs, increase *in vivo* stability, extend the drug's blood circulation time, and allow for controlled drug release (Maeda et al., 2013; Blanco et al., 2015; Wicki et al., 2015). However, despite potentially high benefits, state-of-the-art NP-based cancer treatments have not shown the expected efficacy, in part due to the lack of detailed knowledge of the behaviour of NP-drug complexes (Wilhelm et al., 2016). Therefore, there has been a lot of focus on the use of computational models for improving NP design (Rockne et al., 2019). However, simulations have mostly considered only certain tumour scenarios and have disregard other crucial aspects of treatment

☆ This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 800983.
* Corresponding author.
E-mail address: igor.balaz@df.uns.ac.rs (I. Balaz).

https://doi.org/10.1016/j.biosystems.2020.104290
Received 30 August 2020; Received in revised form 12 November 2020; Accepted 12 November 2020
Available online 17 November 2020

strategies. For example, although tumour heterogeneity (i.e. existence of various cell subpopulations: sensitive and resistant cells, proliferating, quiescent cells and CSC) is taken into consideration when modelling tumour growth, most models do not take this into account when modelling tumour treatments. Various cells can respond differently to therapy and, hence, can influence nanoparticle design (Karev et al., 2006). The accumulation of nanoparticles at effective doses in cancer cells is also not considered (Hauert et al., 2013).

In summary, in order to design an efficient treatment tailored to a specific tumour, we are facing an immense state-space generated by: (i) various drugs for targeting specific cancer cells, (ii) NPs that can modify physicochemical properties of those drugs, (iii) heterogeneous and highly adaptable tumour that can quickly become resistant to the primary therapy, and (iv) treatment specific details, such as the concentration and dose schedule. While models can certainly help us make informed decisions, computationally testing all possibilities is intractable. It has been demonstrated that both Gaussian process models and multi-layer perceptron neural network models are able to effectively explore the parameter-space of bio-physical tumour properties in order to find the optimal therapeutic strategy (Preen et al., 2019). However, that work has investigated the effect of NPs on homogeneous, static tumours that cannot evolve. To automate the creation of therapies for more realistic cases, we need a more open-ended approach where the exploration of an initially small portion of the given state space can be expanded by an ongoing generation of adaptive novelties, whenever the simulated tumour makes the adaptive leap (see the Model architecture for further details). An initial attempt towards the implementation of open-ended approaches was made by utilizing the novelty search method (Tsompanas et al., 2020), however, once again a homogeneous, static type of tumour was studied.

To demonstrate the basic implementation of, what Taylor (2019) would classify as an exploratory open-ended evolutionary model, we created a simple engine for the virtual evolution of combinatorial oncological treatments. It is a system of non-interacting evolvable nano-agents in the environment of cancer- and healthy-cells. Since the tumour can adapt and, thus, become resistant to nano-agents, the adaptive landscape of nano-agents is under constant modification so the context in which their evolution occurs is continuously changing. As a result, the population of nano-agents reshapes itself and creates a heterogeneous population of specialized groups of individuals optimized for tracking and killing different phenotypes of cancer cells.

2. Model architecture and simulation settings

We use Mesa, an agent-based modelling platform for Python (Masad and Kazil, 2015). Complete source code of the model is available at https://github.com/yotf/CancerModel/. The world is non-toroidal 2D grid populated with three types of cell-agents (cancer cells — CC; cancer stem cells — CSC; healthy cells — HC) and nano-agents — NA. Multiple agents can occupy the same location in the grid.

Cell-agents: All cell-agents are unmovable. CC and CSC can mutate with some small probability (see Simulation Setting). By mutating they change their "visible" properties through which nano-agents can recognize them. The tumour can grow through the division of CC and CSC. In scenarios with the growing tumour, after each time step, tumour cells divide with a certain probability (termed "growth rate" in the following) if there are unoccupied grid locations in their immediate neighbourhood. Healthy cells (HC) can be killed, but do not mutate or grow.

Nano-agents: The nano-agents can move, attack, observe and memorize their environment. They can have two kinds of behaviour. They can either try to kill cell-agents or inhibit their division and the quantity of each agent type is determined at the beginning of the simulation (see Table 1). At the beginning of the simulation, NA's knowledge of the environment is blank so they do not recognize any type of cell-agents. *Movement*: At each time step each nano-agent moves randomly into its Moore neighbourhood. The speed parameter s determines how far it can go. *Observing & learning*: After moving, the nano-agent observes its environment, which, in this version, is restricted only to the grid cell that the nano-agent landed on. By observing, the nano-agent checks the visible properties of the cell-agent with whom it shares the grid location and compares them with those stored in the nano-agent's memory. If it encounters an unknown cell-agent, it memorizes its properties. The memory is finite and its size is randomly chosen in the range 0−3 (e.g. if memory size is 3 it allows storage of up to three visible properties). If memory is full, new information pushes out the oldest existing one.

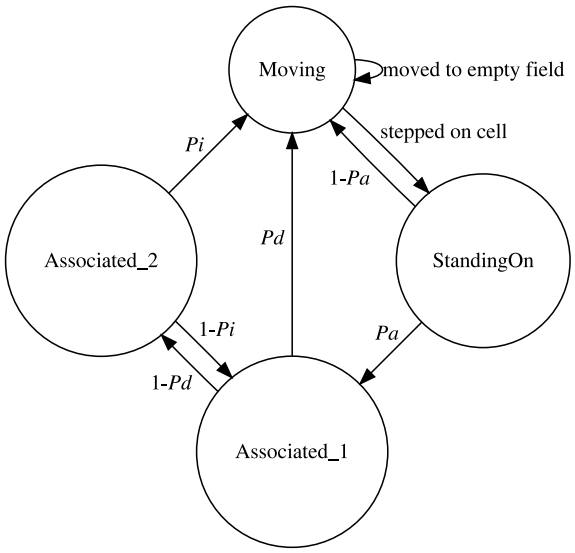

Fig. 1. State diagram describing behaviour of a nano-agent after recognizing a cancer cell. p_a: probability of binding to a cell; p_d: probability of disassociation from a cell; p_i: probability of internalization into a cell.

Attacking tumour and healthy cells: When a NA is in the same location, it goes through a series of states described in Fig. 1. The NA will attempt to kill or inhibit the division of a cell only if it gets internalized into that cell (for the probabilities of association (p_a), disassociation (p_d) and internalization (p_i) see Table 1). Since the goal is to virtually evolve potential treatment strategies, after trying to attack the cell, NAs do not perish but continue with movement in the next time step. To discriminate relative importance of cells, NA gets +1 point for killing or stopping the division of CC, −1 point for HC and +5 for CSC. If the NA encounter already memorized cell-agent, the probability of binding is defined by p_a. Otherwise, the probability of binding is reduced by multiplying p_a with the factor c. Through the factor c, we can fine-tune the "curiosity" of NAs in exploring the environment.

Tumour resistance: Randomly chosen subpopulation of tumour cells is partially resistant to NAs, by adding resistance modifiers to them. The default size of the resistant sub-population is 10% of the total number of tumour cells. Upon interacting with NAs, resistant cell can modify one of the following NA's properties: binding (p_a), disassociation (p_d), internalization (p_i), killing (p_k) and inhibition of division (p_{sd}). The exact property to which each cell-agent is resistant is randomly determined at the beginning of the simulation. Resistance strength is randomly chosen in the interval 30 – 80%. This means that upon the interaction between a resistant cell and a nano-agent, a corresponding property of a nano-agent is reduced by the chosen percentage. If cells divide, resistance is passed on to the next generation. All CSC are partially resistant to killing and inhibition of division, by randomly choosing resistance strength in the interval 50 – 80%, at the beginning of the simulation.

Evolution of nano-agents: When a NA kills a cell-agent it is assigned positive or negative points (described above). The performance of each

Table 1
Simulation parameters. p_a - probability of association; p_d - probability of disassociation; p_i - probability of internalization; p_k - probability of killing; p_{sd} - probability of stopping cell division.

Description	Value
Grid size	35*x*35
Initial number of cancer cells	297
Initial number of cancer stem cells	3
Initial number of healthy cells	925
Size of nano-agent population	500
CC and CSC mutation probability	0.1
CC and CSC growth rates	0.01, 0.005, 0.001
Curiosity (c)	0.5
Detachment (d)	10%
Number of injection sites (n)	5
Speed range	1 – 5
p_a, p_d, p_i, p_k, p_{sd} of NA^G that attack CC	0.7, 0.5, 0.7, 0.7, 0.7
p_a, p_d, p_i, p_k, p_{sd} of NA^G that attack HC	0.5, 0.5, 0.5, 0.5, 0.5

NA is measured after 10 time steps by comparing collected points. Five percent of top performers are duplicated and the same percent of worst performers are eliminated; hence, the population size of NAs remains constant. Middle five percent of remaining NAs are mutated. Mutable properties are speed, memory size, probabilities of association, disassociation, and internalization, as well as behaviour (inhibit cell division or kill a cell). To compare different simulation runs we defined the overall fitness function as $f(NA) = \sum HC/(r\sum CSC + \sum CC)$ where $r = 5$ is the relative importance, that gives advantage to NAs who killed CSCs.

Metastasis: Since CSC drive the metastasis (Gener et al., 2016) we implemented a simplified measure of metastasis as an additional numeric property the "detachment" d of a CSC agent. At each step, the CSC will leave the tumour with 10% probability. If the CSC leaves, the metastasis score is incremented by 1.

Simulation Settings: Since the two main goals of experiments in this manuscript are to test (i) whether, and to what extent, NAs are capable of dealing with an evolvable tumour, and (ii) whether spontaneously evolved combinatorial therapy can be more efficient than generalized therapy, we created three classes of NAs:

- NA^0 agents that cannot learn. They can only recognize and attack initial (non-mutated) tumour cells. If the tumour cells mutate they become invisible to NA^0. A real-life equivalent is treatment with single anti-cancer drugs that are specialized for a certain tumour type;
- NA^E fully evolvable agents;
- NA^G agents that attack all cells, both healthy and tumour, with a slight bias towards tumour cells (See Table 1). A real-life equivalent of this is treatment with cytostatic drugs;

To represent a release of nanoparticles through the bloodstream, NAs are injected at n randomly chosen locations in the grid. After injection agents disperse through the grid space, at the speed determined by s.

In order to determine how many experimental runs (denoted as e^*) are necessary to mitigate the uncertainty that stems from model stochasticity, we first performed a Consistency Analysis. Detailed methodology is described elsewhere (Alden et al., 2013; Read et al., 2012). In short, we analysed sample sizes of $1, 5, 50, 100, 150, 200$ and 300 *in silico* runs. For each sample size we created 20 groups of distributions, all generated using the same fixed set of parameter values and containing identical numbers of simulation samples. So, for example, sample size of 100, each of the 20 groups of distributions contains Fitness function results of 100 runs (which gives a total of 2000 runs). Then, to compare distributions within one group, we employ a non-parametric effect magnitude test based on the Vargha–Delaney Â test (Vargha and Delaney, 2000). The Â test compares two population distributions and returns a value in the range [0.0,1.0] that represents the probability that a randomly chosen sample taken from population A is larger than a randomly chosen sample from population B. A value of 0.5 indicates no difference at all, whereas values above 0.71 indicate a "large" difference in the distributions. As an acceptable level for "small" difference, we used 0.56. The minimal sample size where maxÂ value is below 0.56 for all tested run durations (up to 1000 steps) and for both scenarios (static and growing tumour) is $e^* = 200$ (Fig. 2). Therefore, all results shown below represent an average of 200 runs.

3. Simulation results and discussion

Non-evolvable, homogeneous tumour.

To test the basic dynamics of nano-agents, we first simulated their behaviour in the most simple tumour scenario: non-evolvable homogeneous tumour. Here, cancer cells cannot mutate and do not have resistance modifiers. In combination with NA^0 agents this constitutes the minimal setting of our model. In such scenario, fully evolvable NA^E agents and non-evolvable NA^0 agents behave identically (Fig. 3a). They eliminate the entire tumour and reach the fitness peak in the first 500 time steps. NA^G agents are the worst-performing. Even though they manage to kill the tumour, in the process they also kill all the healthy cells. Given the high incentive to eliminate CSC (defined by the fitness function) and the low CSC number, all three types of nano-agents eliminate the entire CSC population within the first 40–60 generations, so the metastatic score remains low.

Although the results seem promising, this tumour scenario is highly unrealistic. Real tumours are not homogeneous masses of identical cells. In fact, the main limiting factor in treating cancer patients is drug resistance, that stems from tumour heterogeneity (Vasan et al., 2019). Tumour resistance can be divided into two categories: the intrinsic resistance (as a result of pre-existing resistant cells in a tumour, before drug treatment) and the acquired resistance (developed during or after therapy either by genetic changes in a tumour microenvironment). Therefore, in the following section, we investigate to what extent the open-ended evolution of nano-agents can deal with both categories of resistance.

Tumour drug resistance and exploratory open-endedness of NAs.
Here, we simulated the evolution of all three classes of nano-agents in three tumour scenarios: (1) Non-evolvable heterogeneous tumour (cancer cells cannot mutate and 10% of cancer cells are partially resistant), (2) Evolvable homogeneous tumour (cancer cells can mutate and do not have resistance modifiers) and (3) Evolvable heterogeneous tumour (cancer cells can mutate and 10% of the cells are partially resistant). In the model, mutations of cancer cells change their phenotype, so they can evade recognition by NA. Therefore, the first scenario represents a tumour with only intrinsic resistance, the second scenario is a homogeneous tumour with only acquired resistance, while the third scenario represents a more realistic scenario of a tumour with both intrinsic and acquired resistance.

In all of these more challenging scenarios, fully evolvable NA^E agents show similar efficacy. Within the first 400 – 500 generations they learn to recognize and kill all cancer cells (Fig. 3b,c,d) while keeping the healthy cells mostly intact (as shown in Fig. 4).

In the non-evolvable scenario, NA^0 agents behave similarly to NA^E agents. The only difference is that NA^0 agents reach the fitness peak slightly earlier than NA^E agents, regardless of the initial percentage of resistant tumour cells (Fig. 5). The reason for that is the functioning of the learning algorithm in NA^E agents. NA^0 agents do not learn during the simulation, so they immediately proceed with killing cancer cells. In contrast, evolvable NA^E agents have limited memory (maximum 3) and they can "forget" old data. Therefore, in order to kill all cancer cells, some of NA^E agents need to re-learn to recognize diverse cancer cells. This process of forgetting and re-learning generates a time lag, that disappears when we increase maximum memory capacity to 10 (data not shown).

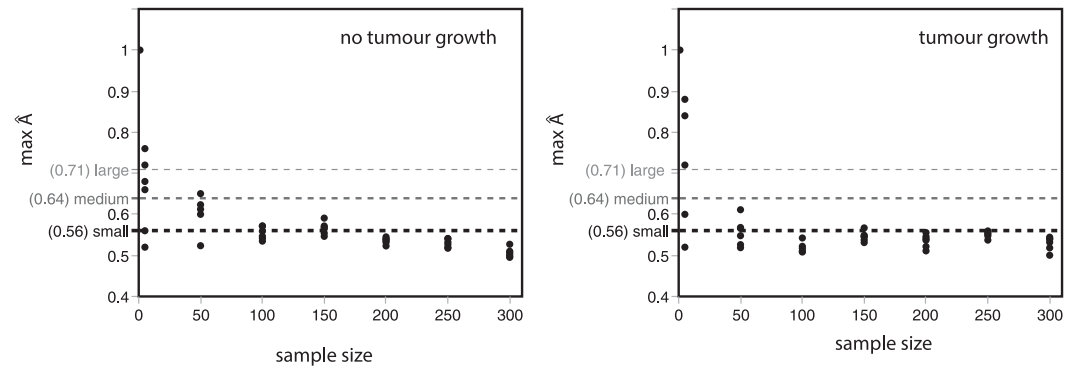

Fig. 2. Â test results for up to 1000-step simulation runs. For each sample size we calculated Â test after 50, 100, 250, 300, 500 and 1000 time steps (black dots).

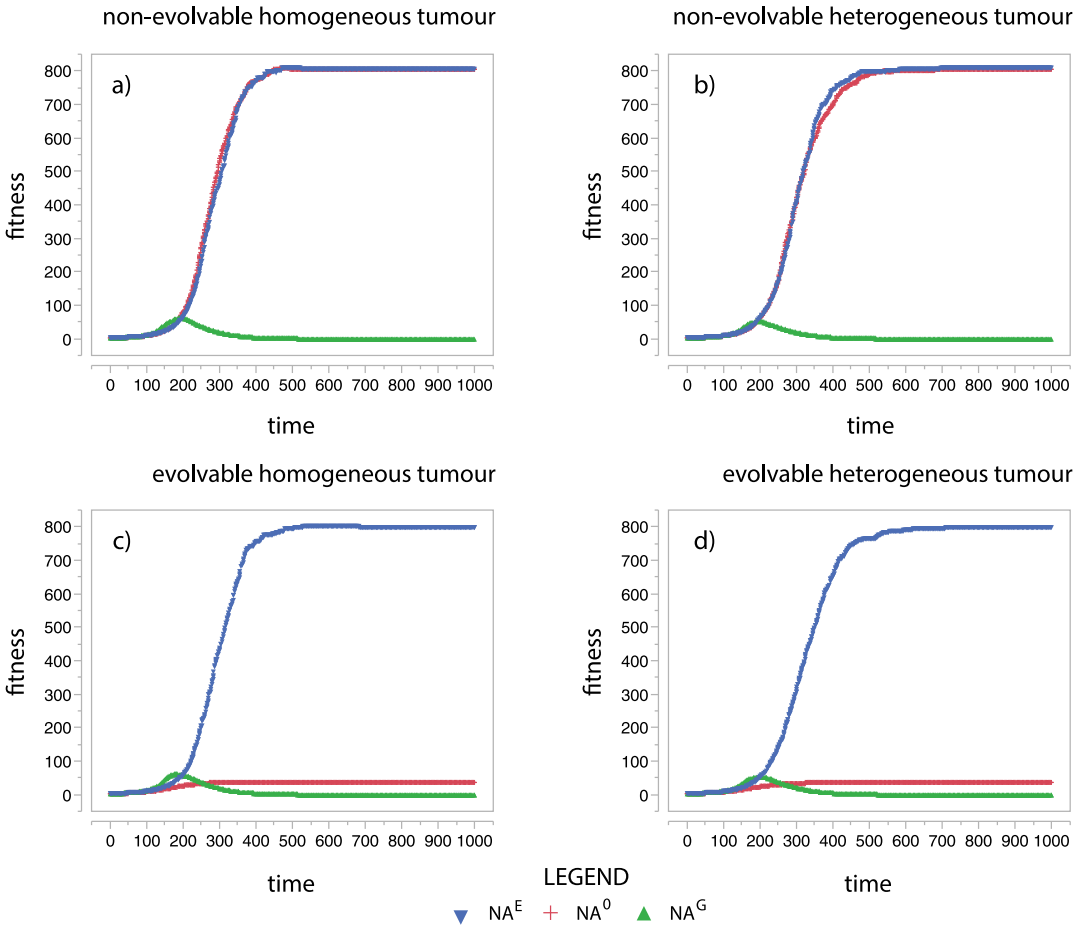

Fig. 3. Fitness evolution of NA^0, NA^G and NA^E agents for different tumour scenarios. Heterogeneous tumours have 10% of resistant cells.

In evolvable homogeneous and heterogeneous tumour scenarios, NA^E agents show their full potential compared to NA^0 and NA^G agents. In both evolvable scenarios, the tumour managed to escape NA^0 agents, through evolution, and become composed entirely of mutated cells (for heterogeneous example see Fig. 6a). When treated with NA^E and NA^G agents, this is not the case. The percent of mutated tumour cells starts to rise at the beginning of the simulation but then rapidly declines. NA^G agents can, right away, without having to memorize and recognize them, eliminate all cancer cells, so the maximum reached percentage of mutated cells is quite low. However, when both (homogeneous and heterogeneous) tumours are treated with NA^E agents, the number of mutated cells reaches 40% and their peak is at about 100 time steps later than in the case of NA^G agents. As with the lag in reaching fitness peak in a non-evolvable tumour scenario, the reason lies in the learning process. In contrast to NA^G agents, in order to eliminate emerging mutated cancer cells, NA^E agents have to continuously learn, memorize, and eventually re-learn the environment. Therefore, in this scenario, NA^E agents are somewhat slower in eliminating tumour cells than NA^G agents.

If in the evolvable heterogeneous tumour we observe intrinsic tumour resistance, represented by heterogeneity in the resistance of the cells, we can see a similar trend for NA^E and NA^G agents (Fig. 6b). The percentage of resistant cells first grows, because of the selective pressure: NAs can more easily kill non-resistant cells, so they are the first to be eliminated. This leaves the empty space for the more resistant cells to replicate. With NA^E and NA^G treatment, the percentage of resistant tumour cells goes up to approximately 40%, whereas the peak and duration of tumour resistance are much smaller for NA^G treatment.

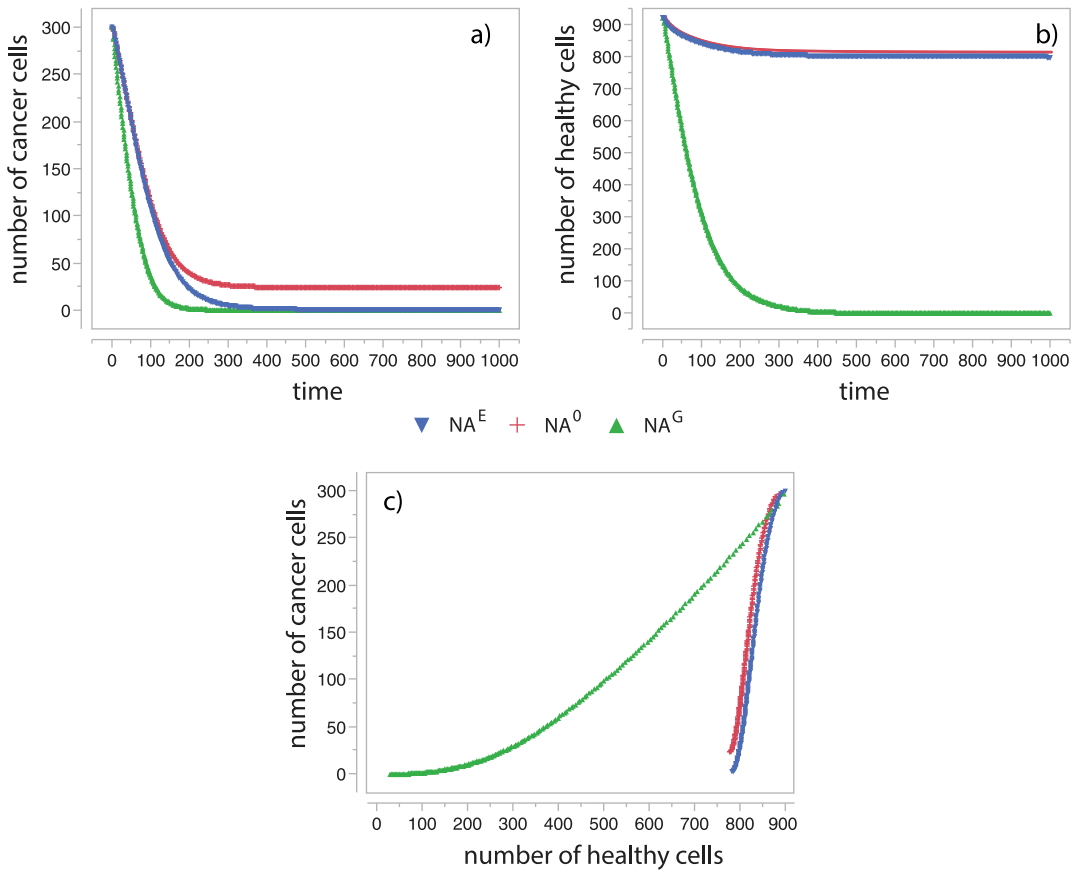

Fig. 4. Efficacy of all three classes of nano-agents in evolvable heterogeneous tumour scenario. (a) number of cancer cells as a function of time. (b) number of healthy cells as a function of time. (c) relation between the number of healthy cells and the number of cancer cells.

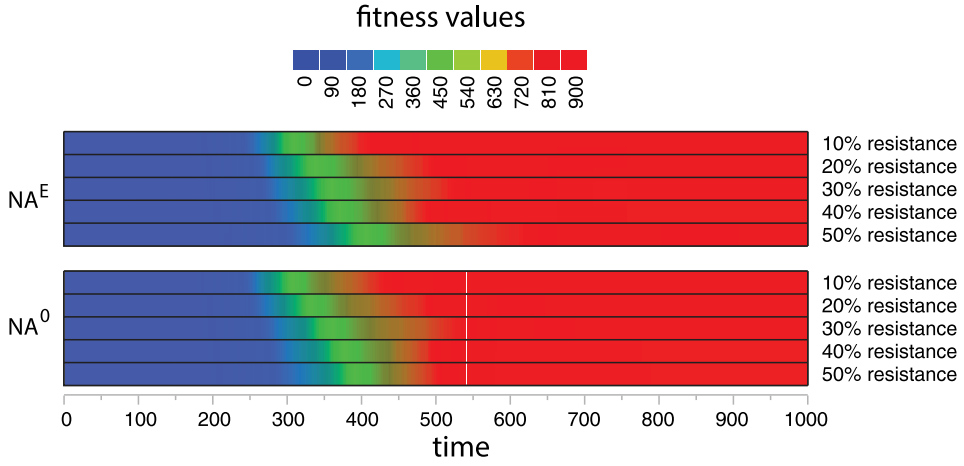

Fig. 5. Fitness evolution of NA^E and NA^0 agents for non-evolvable heterogeneous tumour. Percent of resistance indicate initial number of cancer cells with resistance modifiers.

Again, NA^G agents do not need to spend time learning to adapt to newly emerging resistant cancer cells, as is the case for NA^E agents. Nevertheless, NA^E agents are very efficient in learning to recognize new cancer cell phenotypes (Fig. 6c) needing less than 100 generations to be able to recognize almost all emerging phenotypes. In the case of NA^0 treatment, only a fraction of initially recognized cancer cells will be eliminated while the rest of the tumour will escape recognition via phenotype mutations. Non-resistant cells have a higher probability to be among firstly eliminated cells so the percentage of resistant cells only slightly increases and then remains constant.

To conclude this section, we would like to emphasize that even though NA^G agents are fastest to eliminate the entire tumour (Fig. 4a) and to deal with mutated and resistant cancer cells, their efficacy is consistently the worst, primarily because they end up killing both tumour and healthy cells (Fig. 4b). We tested several scenarios of NA^G's bias towards killing tumour cells by increasing the probability of attacking tumour cells and decreasing the probability of attacking healthy cells (Fig. 7). As expected, higher bias towards cancer cells leads to a sharper increase in fitness, but eventually, due to the killing of healthy cells, fitness converges towards zero in all tested cases.

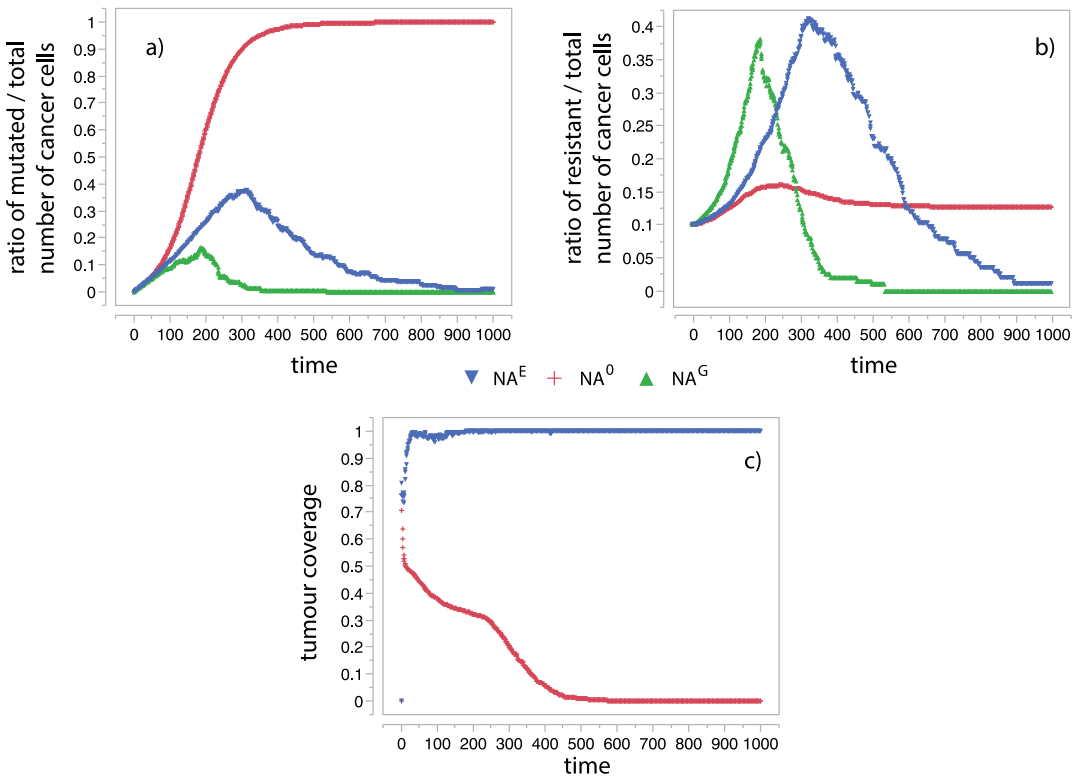

Fig. 6. Behaviour of evolvable heterogeneous tumour in different treatment regimes. (a) Temporal evolution of the ratio of phenotype mutated tumour cell and the total number of tumour cells. (b) Temporal evolution of the ratio of resistant tumour cell and the total number of tumour cells. (c) Tumour coverage is the ratio of the number of tumour cell types that the population of nano-agents can recognize and the total number of tumour cell types. NA^G agents are not shown because they can recognize all cells so their coverage would always be 1.

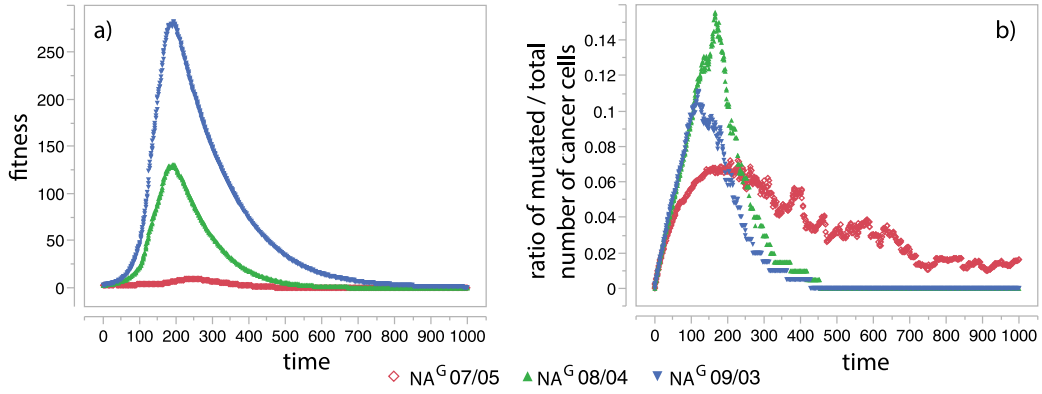

Fig. 7. (a) Fitness and (b) tumour mutation response when treated with NA^G with different biases towards cancer cells. In addition to the default one (specified in Table 1 and indicated here as NA^G 07/05) for which p_a, p_i, p_k and p_{sd} values are 0.7 (when attacking cancer cells) and 0.5 (when attacking healthy cells), we also tested NA^G 08/04 and NA^G 09/03.

Growing tumour To test whether the growing tumour will significantly change the outcome of treatments with NAs, we simulated a heterogeneous, evolvable tumour with the following growth rates: 0.01, 0.005 , and 0.001. As expected, for higher growth rates, NAs need more time to eliminate all cancer cells. For the growth rate of 0.001, NA^E agents need about 500 generations (compared to approximately 300 generations for the non-growing tumour — see Fig. 4a), while for the growth rate of 0.005 time needed is much longer (almost 5000 generations). Only in the case of the fastest tested growth rate (0.01), after the initial sharp decline, the population of remaining cancer cells starts to grow (Fig. 8a) and eventually became larger than the original tumour.

For all tested growth rates, tumour resistance continuously grows (as an example see Fig. 8b for the scenario with the growth rate = 0.01), mostly as a result of 50 fold increase of disassociation probability. This is in contrast with a non-growing tumour where probabilities of association, disassociation, and internalization remain constant during the simulation period. Such discrepancy is a result of different survival strategies. In a non-growing tumour, the number of cells is fixed, so there is no possibility to select for any property that will proliferate, evolve, and became dominant through generations. However, in a growing tumour, the fact that cells can divide gives exactly that opportunity to more resistant cells. Apparently, the property of cancer cells that is mostly favoured is the disassociation probability (probability of internalization is mostly constant, while the probability of association slightly increases).

At the same time, the percentage of mutated cancer cells shows the expected trend: in NA^0 treatment tumour population gradually escape

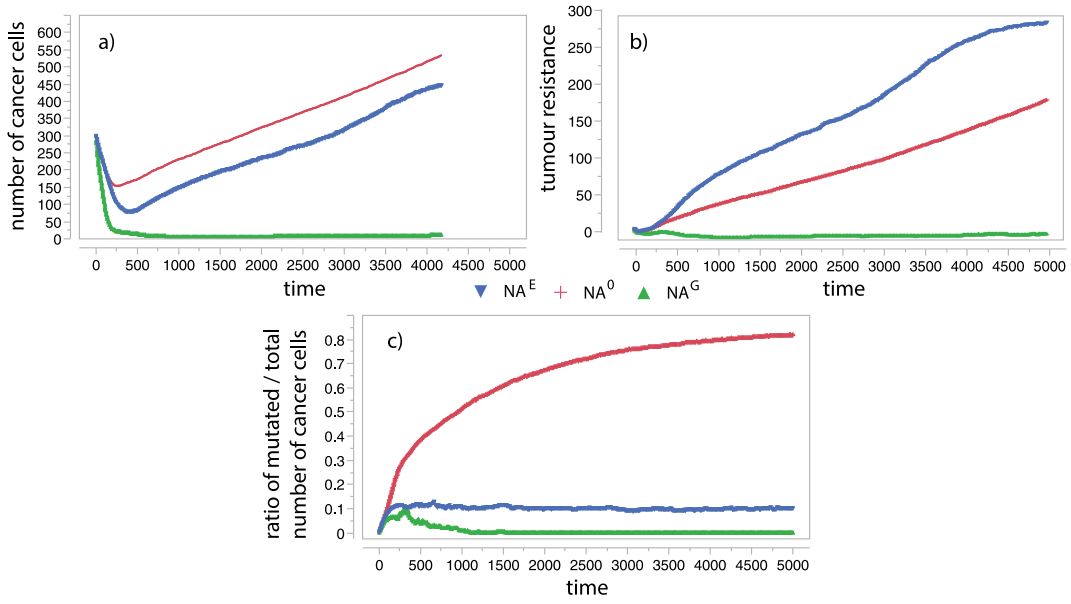

Fig. 8. Behaviour of evolvable heterogeneous growing tumour in different treatment regimes (growth rate = 0.01). (a) Temporal evolution of the total number of cancer cells. (b) Temporal evolution of tumour resistance, defined as a sum of resistance effect of all living cancer cells. (c) Temporal evolution of the ratio of phenotype mutated tumour cell and the total number of tumour cells.

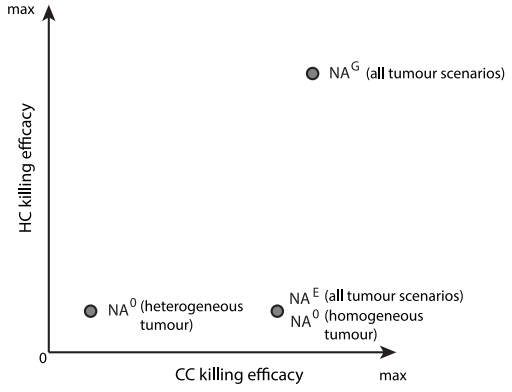

Fig. 9. Summary scheme of efficacy of all three NA classes simulated in this paper. CC — cancer cells; HC — healthy cells.

treatment by mutating, while open-ended evolution of NA^E agents can continuously track down newly emerging tumour mutations and keeps them under control (Fig. 8c).

4. Concluding remarks

In the model presented here, our primary aim was to demonstrate that open-ended evolution of NA^E agents leads to the emergence of adaptive and combinatorial therapy that can eliminate a heterogeneous, evolvable tumour, which is one of the main problems in today's oncological clinical practice (Vasan et al., 2019). Our main finding is that, despite the tumour's phenotypic plasticity, NA^E agents are flexible enough to respond to such changes and successfully track down and kill all cancer cells. Their efficacy of killing cancer cells is somewhat lower compared with NA^G agents (Fig. 9), but at the same time, NA^G agents kill most of the healthy cells making extensive damage to a patient. During simulations, in both the growing and the non-growing tumour, initially homogeneous NA^E population starts to diverge and became increasingly heterogeneous as nano-agents become specialized for recognizing each newly emerged type of cancer cells, while at the same time optimizing their own parameters needed for eliminating the tumour. For the first 300 generations Pearson's correlation between NA^E population heterogeneity and the percent of mutated cancer cells is 0.77. After that, the percentage of mutated cancer cells starts to decline, while NA^E population heterogeneity remains stable until the end of the simulation period (data not shown).

However, the model presented here lacks many of the details necessary for the computational platform that could be used in real-world research and clinical practice. In modelling the influence and impact of nanoparticles on cancer progression, multiple scales have to be considered: molecular-scale interactions of individual nanoparticles with drugs and their environment; cellular scale interactions of numerous nanoparticles with the cell environment as well as the macroscopic scale which describes details of cancer cells, their mutual interactions, and their interactions with the environment (Bellomo et al., 2003). Furthermore, in order to apply this finding in a real-world setting, both the simulated agents and the simulated tumour must be based on realistic parameters derived from the nanoparticle's physicochemical properties and tumour's physiology. Physicochemical properties of nanoparticles add complexity associated with pharmacodynamics, pharmacokinetics, tumour accumulation, and biodistribution. Also, when nanoparticles enter the bloodstream, they encounter several transport barriers such as the formation of the protein corona, vascular transport, extravasation, endocytosis, and RES uptake. Each of these barriers is a considerable individual challenge (Stillman et al., 2020). Therefore, nanoparticles must be designed in such a way to overcome several (if not all) of them to reach maximum efficacy.

As emphasized in the Introduction, the versatility of NPs gives a wide range of customization possibilities. However, it also leads to a rich design space that is difficult to optimize and is computationally expensive. In this paper, using a relatively simple model, we demonstrated that open-ended evolution coupled with multidimensional optimization can lead to the emergence of successful virtual anti-cancer therapies. However, the computational cost of simulations will become unmanageable if we try to implement all aspects mentioned above. Therefore, we intend to further develop this model as a part of a more complex pipeline that will be composed of several modules. The pipeline under development consists of three main modules: an open-ended evolution for the automated virtual design of treatments (this model), simulation of a virtual tumour (PhysiCell (Ghaffarizadeh et al., 2018)), and simulation of nanoparticle–cell

interactions (STEPS (Hepburn et al., 2012)). In our next steps we will further refine this model by adding three main features:

- realistic mapping between properties of simulated nano-agents and real chemical counterparts (e.g. setting up speeds p_a, p_d, p_i, p_k and p_{sd} based on measured chemical properties);
- automated evaluation of evolved treatment strategies;
- input/output interface between this model and tumour and cellular simulators so that evolved virtual treatments could be easily tested in a more realistic setting.

How such added complexity will influence the evolutionary emergence of combinatorial therapies will be investigated in our future work.

As a final note, we would like to emphasize that in generalized terms, the approach developed here can be applied to many problems whose structure and outcomes can change over time. Besides oncology, a similar approach could work in designing treatment for infectious diseases where bacteria or viruses can evolve, develop resistance to treatments, and form heterogeneous populations. Outside of the medical field, an example of a possible application is designing a plan of social policy actions where a given segment of society is also changeable and heterogeneous. However, to translate the approach we outlined here to any of the new problem domains would require careful consideration of constraints and structure of that domain. That is also a task for our future work.

CRediT authorship contribution statement

Igor Balaz: Conceptualization of this study, Writing, Methodology, Analysis of results, Data interpretation. **Tara Petrić:** Software development, Data curation, Data interpretation. **Marina Kovacevic:** Analysis of results, Data interpretation. **Michail-Antisthenis Tsompanas:** Analysis of results, Data interpretation. **Namid Stillman:** Analysis of results, Data interpretation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Alden, K., Read, M., Timmis, J., Andrews, P., Veiga-Fernandes, H., Coles, M., 2013. Spartan: A comprehensive tool for understanding uncertainty in simulations of biological systems. PLoS Comput. Biol. 9, e1002916.

Bellomo, N., De Angelis, E., Preziosis, L., 2003. Review article: Multiscale modeling and mathematical problems related to tumor evolution and medical therapy. J. Theor. Med. 5, 111–136.

Blanco, E., Shen, H., Ferrari, M., 2015. Principles of nanoparticle design for overcoming biological barriers to drug delivery. Nature Biotechnol. 33, 941–951.

Bozic, I., Reiter, J., Allen, B., Antal, T., et al., 2013. Evolutionary dynamics of cancer in response to targeted combination therapy. eLIFE 2, e00747.

Chen, S.h., Lahav, G., 2016. Two is better than one; toward a rational design of combinatorial therapy. Curr. Opin. Struct. Biol. 41, 145–150.

Fanciullino, R., Ciccolini, J., Milano, G., 2013. Challenges, expectations and limits for nanoparticles-based therapeutics in cancer: A focus on nano-albumin-bound drugs. Crit. Rev. Oncol./Hematol. 88, 504–513.

Gener, P., de Sous Rafael, D., Fernandez, Y., Ortega, J., Arango, D., Abasolo, I., Videira, M., Schwartz, S., 2016. Cancer stem cells and personalized cancer nanomedicine. Nanomedicine 11, 307–320.

Ghaffarizadeh, A., Heiland, R., Friedman, S.H., Mumenthaler, S.M., Macklin, P., 2018. Physicell: an open source physics-based cell simulator for 3-d multicellular systems. PLoS Comput. Biol. 14, e1005991.

Hauert, S., Berman, S., Nagpal, R., Bhatia, S., 2013. A computational framework for identifying design guidelines to increase the penetration of targeted nanoparticles into tumors. Nano Today 8, 566–576.

Hepburn, I., Chen, W., Wils, S., De Schutter, E., 2012. Steps: efficient simulation of stochastic reaction–diffusion models in realistic morphologies. BMC Syst. Biol. 6, 1–19.

Holohan, C., Van Schaeybroeck, S., Longley, D., Johnston, P., 2013. Cancer drug resistance: and evolving paradigm. Nature Rev. Cancer 13, 714–726.

Hu, Z., Sun, R., Curtis, C., 2017. A population genetics perspective on the determinants of intra-tumour heterogeneity. Biochim. Biophys. Acta 1867, 109–126.

Karev, G., Novozhilov, A., Koonin, E., 2006. Mathematical modeling of tumor therapy with oncolytic viruses: effects of parametric heterogeneity on cell dynamics. Nature Rev. Mater. 1, 30.

Kurtova, A., Xiao, J., Mo, Q., Pazhanisamy, S., Krasnow, R., Lerner, S., Chen, F., Roh, T., Lay, E., Ho, P., Chan, K., 2014. Blocking pge2-induced tumour repopulation abrogates bladder cancer chemoresistance. Nature 517, 209–213.

Maeda, H., Nakamura, H., Fang, J., 2013. The epr effect for macromolecular drug delivery to solid tumors: Improvement of tumor uptake, lowering of systemic toxicity, and distinct tumor imaging in vivo. Adv. Drug Deliv. Rev. 65, 71–79.

Masad, D., Kazil, J., 2015. Mesa: An agent-based modeling framework. In: Huff, K., Bergstra, J. (Eds.), Proceedings of the 14th Python in Science Conference. pp. 51–58. http://dx.doi.org/10.25080/Majora-7b98e3ed-009.

Pattabiraman, D., Weinberg, R., 2014. Tackling the cancer stem cells - what challenges do they pose. Nat. Rev. Drug Discov. 13, 497–512.

Preen, R., Bull, L., Adamatzky, A., 2019. Towards and evolvable cancer treatment simulator. BioSystems 182, 1–7.

Read, M., Andrews, P., Timmis, J., Kumar, V., 2012. Techniques for grounding agent-based simulations in the real domain: a case study in experimental autoimmune encephalomyelitis. Math. Comput. Model. Dyn. Syst. 18, 67–86.

Rockne, R., Hawking-Daarud, A., Swanson, K., Sluka, J., Glazier, J., Macklin, P., Hormuth II, D., Jarett, A., Lima, E., Tinsle Oden, J., Biros, G., Yankeelov, T., Curtius, K., al Bakir, I., Wodarz, D., Komarova, N., Bordyuh, L., Rabadan, R., Finley, S., Enderling, H., Caudell, J., Moros, E., Anderson, A.R.A.G., Kaznatcheev, A., Jeavons, P., Krishnan, N., Pelesko, J., Wadhwa, R., Yoon, N., Nichol, D., Marusyk, A., Hinczewski, M., Scott, J., 2019. The 2019 mathematical oncology roadmap. Phys. Biol. 16, 041005, 1–47.

Shaffer, S., Dunagin, M., Torborg, S., Torre, E., et al., 2017. Rare cell variability and drug-induced reprogramming as a mode of cancer drug resistance. Nature 546, 431–435.

Stillman, N., Kovacevic, M., Balaz, I., Hauert, S., 2020. In silico modelling of cancer nanomedicine, across scales and transport barriers. npj Comput. Mater. 6 (92).

Taylor, T., 2019. Evolutionary innovations and where to find them: Routes to open-ended evolution in natural and artificial systems. Artif. Life 25, 207–224.

Tsompanas, M.A., Bull, L., Adamatzky, A., Balaz, I., 2020. Novelty search employed into the development of cancer treatment simulations. Inform. Med. Unlocked 100347.

Vargha, A., Delaney, H., 2000. A critique and improvement of the cl common language effect size statistics of mcgraw and wong. J. Educ. Behav. Stat. 25, 101–132.

Vasan, N., Baselga, J., Hyman, D., 2019. A view on drug resistance in cancer. Nature 575, 299–309.

Wicki, A., Witzigmann, D., Balasubramanian, V., Huwyler, J., 2015. Nanomedicine in cancer therapy: Challenges, opportunities, and clinical applications. J. Control. Release 28, 138–157.

Wilhelm, S., Tavares, A., Dai, Q., Ohta, S., Audet, J., Dvorak, H., Chan, C., 2016. Analysis of nanoparticle delivery to tumours. Nature Rev. Mater. 1–12, 16014, 1–.