nature reviews bioengineering

https://doi.org/10.1038/s44222-023-00089-7

Review article

Check for updates

Computational approaches to modelling and optimizing cancer treatment

Thomas O. McDonald[1,2,3,4], Yu-Chen Cheng[1,3], Christopher Graser[1], Phillip B. Nicol[1,3], Daniel Temko[1] & Franziska Michor[1,2,3,4,5,6]

Abstract

Computational models can be applied to optimize treatment schedules and model treatment responses in cancer therapy. In this Review, we provide an overview of such computational approaches, including deterministic models, such as those based on ordinary and partial differential equations, stochastic models, spatially explicit agent-based approaches as well as control theory and machine learning methods. We discuss their advantages and current limitations in different scenarios. We outline how therapeutic decision-making can be aided by mathematical and computational approaches and how patient-specific responses can be assessed and incorporated into such methods. We also survey models that can incorporate adaptive changes throughout the course of treatment and discuss data and parameter estimation approaches. Finally, we highlight how such methods can lead to the identification of optimum treatment options for individual cancer and treatment types, and examine the challenges that remain to be addressed to enable the clinical translation of computational models in cancer therapy.

Sections

[1]Department of Data Science, Dana-Farber Cancer Institute, Boston, MA, USA. [2]Center for Cancer Evolution, Dana-Farber Cancer Institute, Boston, MA, USA. [3]Department of Biostatistics, Harvard T. H. Chan School of Public Health, Boston, MA, USA. [4]Department of Stem Cell and Regenerative Biology, Harvard University, Cambridge, MA, USA. [5]The Broad Institute of MIT and Harvard, Cambridge, MA, USA. [6]The Ludwig Center at Harvard, Boston, MA, USA. e-mail: michor@jimmy.harvard.edu

Key points

- Computational approaches can be applied to describe the response of tumour cells to cancer treatment.
- Such computational methods can be based on ordinary and partial differential equation modelling, stochastic modelling and spatially explicit agent-based models.
- Adaptive treatment schedules and incorporation of patient-specific responses allow a personalized assessment of treatment options.
- Control theory and machine learning methods, such as reinforcement learning, can be applied to design cancer treatment schedules.
- Multimodal and longitudinal data sets could be integrated into patient-specific models to identify the best therapeutic options for individual patients.

Introduction

Cancer biology and clinical care have benefited greatly from large-scale, cost-efficient and high-throughput genomics approaches at the bulk and single-cell levels, which can be applied for tens of thousands of samples and millions of single cells[1]. In addition, phenotypic data of tumour and microenvironmental cells, such as growth rates, migration, invasion and interaction kinetics, can be obtained for multiple tumour types, treatments and in vivo models[2], giving insight into tumour diversification during cancer progression and bottleneck effects of treatments. These approaches also allow the molecular profiling of a patient's tumour and, thus, the identification of the best treatment option, thereby advancing precision medicine in immunotherapy, radiation, chemotherapy and combinations thereof[3].

However, the complexity of human tumours at such high resolution requires new analyses to assess big data sets and filter important biological nodes, pathways and networks associated with treatment responses to create and choose the best therapeutic options for individual patients. In particular, statistical and bioinformatic methods can be applied to identify differentially expressed genes[4], proteins[5] and metabolites[6], correct for potential batch effects and confounders[7] and predict responders versus non-responders using multi-omics data sets[8].

In many cases, however, these methodologies are limited to the identification of associations between molecular nodes and outcomes, precluding the study of causal relationships. To make mechanistic connections between inputs and outputs in large complex systems, mathematical and computational models can be applied, which enable a systematic analysis of scenarios, hypotheses and counterfactuals, to derive a quantitative, mechanistic understanding of cancer treatment response[9,10] (Box 1).

The goal of cancer treatment is to improve patient survival. Clinical strategies are routinely identified using preclinical data sets that suggest a potentially tolerable dose-concentration profile; these profiles are then tested in a phase I/II dose-escalation approach, in which patient groups are treated using successively escalating concentrations until the maximally tolerated dose (MTD) is reached. The expansion cohort and subsequent phase III trials are then used to ascertain that this schedule is safe and efficacious. Such therapies are developed with the goal of reducing tumour burden and maximizing survival; however, resistance often arises and poses an obstacle to continued response under the recommended treatment even when therapies are initially successful.

Mathematical models can make use of available data and knowledge of the underlying biological mechanisms of cancer growth and resistance to study the dynamics of treatment response and the emergence of resistance, and to suggest more effective protocols[10]. Scenarios arising in cancer treatment, such as targeted agents with known resistance profiles, combination treatments such as chemoradiation, or immunotherapy, all require different computational modelling methods with different assumptions, input data and predictions. Therefore, models developed to investigate and propose treatment strategies in cancer vary greatly (Table 1); for example, some strategies aim to reduce tumour burden as quickly as possible before resistance can develop, generated through the acquisition of (epi)genetic changes during the continuous cell division of drug-sensitive tumour cells[11,12]. The emergence of resistance may then lead to changes in treatment schedules, treatment combinations or personalized dosages that do not fall into the 'one-size-fits-all' category of MTD-based drug dosing[13]. By contrast, in some approaches, the consideration of resistance can lead to a new goal, such as prolonged tumour control instead of tumour eradication. Cancer management strategies such as adaptive therapy[14] are subject to change as the tumour progresses or regresses during treatment. These strategies allow drug-sensitive tumour cells to survive with the aim that they will inhibit the growth of resistant subpopulations through intratumour competition to delay the outgrowth of resistant cells, known as competitive release[15].

In this Review, we discuss computational approaches to modelling cancer treatment response, studying the emergence of resistance and identifying optimum treatment strategies. We outline different methodological approaches and discuss their applications and potential for clinical translation. We conclude by summarizing the challenges and clinical opportunities in this field.

Ordinary differential equation-based models

As a tumour changes in size, composition and behaviour over time, ordinary differential equations (ODEs) are a useful tool in the investigation of cancer dynamics, in particular, the dynamics of treatment response. Using biological assumptions describing the growth dynamics and interactions of cell types (Fig. 1a), a tumour and/or its microenvironment can be mathematically described with a set of ODEs (Fig. 1b) and solved analytically, providing a trajectory of tumour growth over time given the initial condition (Fig. 1c). In addition, treatment can be integrated as a time-dependent function affecting the system (Fig. 1d). At the population scale, ODE models can be used to understand the dynamics of different cell populations and their interactions and allow researchers to explore how treatment interventions may affect the tumour (microenvironment) composition. These models can then be used to create strategies to combat the emergence of resistance and influence the extent of intra-tumour heterogeneity[16–20]. As such, ODEs have long been used for investigating cancer treatment response; for example, the effects of radiation on cell survival was originally described by the linear-quadratic ODE-based model in 1942 (ref. 21), which is still used today as the basis for modelling cell survival after radiation treatment owing to its close agreement with experimental results[19]. Similarly, ODE models are often used to understand the population dynamics of tumour growth and resistance[20,22–29].

Box 1

Computational models

Mathematical modelling: the process of using mathematical formalism and concepts to quantitatively describe a system to explain phenomena and predict behaviour.

Computational modelling: the process of using computational tools and algorithms to simulate and study a system; applicable if simple analytic solutions are not easily accessible.

Adaptive therapy: the strategy of maintaining a stable tumour burden by modelling intratumour competition between sensitive and resistant cells and altering treatments and schedules.

Ordinary differential equation: an equation that describes how a dependent variable changes with respect to another independent variable.

Partial differential equation: an equation that describes how a dependent variable changes with respect to multiple independent variables.

Stochastic differential equation: a differential equation in which some terms are random variables, introducing noise into the system.

Control theory: a field of mathematics that deals with applications of dynamical systems that are designed to control a system to achieve a desired state.

Reinforcement learning: feedback-based machine learning that includes different data-driven methods to learn policies for sequential decision-making.

Evolutionary game theory: a mathematical formulation for describing the growth and evolution of multiple populations by accounting for competition between strategies, payoffs and interactions of individuals within the populations.

Reaction–diffusion systems: mathematical models that describe the interaction of multiple species as well as their spread across space over time.

Branching process: a stochastic process that models the growth and composition of a population of dividing individuals based on the rates of division, death and other transitions of the individuals.

Pharmacokinetic and pharmacodynamic modelling: Pharmacokinetic modelling is a mathematical or computational modelling approach to describe the effect of a drug over time in terms of drug efficacy (pharmacodynamics) as well as absorption, distribution, metabolism and excretion (pharmacokinetics).

Moran process: a discrete-time stochastic process for studying evolution in fixed populations; in each step, an individual is chosen to divide with certain probability according to its fitness while a randomly chosen cell dies.

Wright–Fisher process: a discrete-time stochastic process for studying evolution in fixed populations; in each time step, the composition of the current generation is drawn independently at random from all types in the previous generation.

Agent-based modelling: computational models for individuals based on a set of rules that use simulation to describe and predict emergent phenomena of complex systems.

Logistic growth: the logistic growth model of population growth characterizes growth as a sigmoid function: at the early stage, the population grows exponentially as resources are abundant; however, as the system approaches the carrying capacity, population growth reaches a plateau owing to limited resources.

Gompertzian growth: the Gompertzian growth function is a special case of general logistic growth. The Gompertzian function approaches two asymptotes (initial value on the left side and future value on the right side) at different rates.

Equilibrium point: in the theory of dynamical systems, an equilibrium point is a constant solution, also known as a steady state. In general, there could exist multiple equilibrium points in a system. Different initial conditions lead to different corresponding equilibrium points.

Deep neural network: a type of artificial neural network consisting of multiple hidden layers between the input and the output. The network structure and algorithm are inspired by the human brain and use data as input to train themselves to learn the patterns of the data.

In addition, quantitative systems pharmacology applies ODEs to predict drug response and to model biological pathways to aid in drug development[30,31].

Treatment response modelling of drug resistance

Genetic alterations and non-genetic plasticity can be integrated into computational models of treatment responses[32,33] to shed light on therapeutic opportunities[34]. For example, ODE-based models can be applied to predict treatment responses under different resistance mechanisms, such as drug-induced resistance that occurs if the drug causes a cell transition from a sensitive to a resistant phenotype, or acquired resistance that arises owing to non-genetic plasticity or the acquisition of genetic alterations[20,22–29]. Such resistance may exist prior to treatment onset or arise stochastically after initiation of therapy. Some treatment types, particularly those that induce cell killing through DNA damage, such as radiation, can also contribute to the generation of genetic resistance mechanisms[35].

Mathematical models of tumorigenesis. Models combining sensitive and resistant subclone dynamics in the presence and absence of treatment can be used to identify treatment strategies for delaying progression. ODE models allow the modelling of interactions between different cell types and the resulting cellular dynamics by describing cell-cell interactions without accounting for individual cell fates. For example, coupled linear ODEs can be applied to model the growth and transition between drug-sensitive and drug-resistant cells, which was found to have a more important role than selection in the emergence of resistance in HL60 leukaemic cells following treatment with the

Table 1 | Computational modelling methods in cancer therapy

Model class	Use case	Advantages	Disadvantages
Ordinary differential equations	Useful if drug response depends on the aggregate behaviour of a large population of cells or molecules, and if the rules governing the time development of these populations in terms of their current states are known[21]	Mathematical results can describe the general behaviour of the system without resorting to simulations	Limited applicability owing to the modelling of general behaviours that do not account for stochasticity that may be present at early stages
Partial differential equations	Useful if drug response depends on the spatiotemporal behaviour of large populations of molecules, and if the rules governing the development of these populations in time and space in terms of their current states are known[56]	Formal mathematical results are available, but analysis can be more challenging than for ordinary differential equations	Presents a balance between mathematical tractability and expressivity if spatial effects are important
Stochastic models	Useful if drug response is inherently uncertain, often owing to the effect of stochasticity in small populations[75]	Enables a rigorous quantification of uncertainty for stochastic events	Mathematical analysis is highly specialized and can be challenging
Spatial and agent-based models	Useful if drug response depends on the spatiotemporal behaviour of individual cells[108]	Flexible framework; any rules can be incorporated as long as they can be coded in a computer	Requires computationally intensive simulations
Control theory	Mathematical analysis framework for identifying optimal treatment schedules under a model; useful if the underlying model dynamics are not too complex[127]	Optimal schedules can be obtained by mathematical analysis	The need to specify a correct model for the system limits the applicability of control theory; some systems may be too complex to specify mathematically
Reinforcement learning	Machine learning framework for identifying near-to-optimal treatment schedules; useful for data-rich settings, in which the model or real-world scenario is complex[149]	Flexible framework for obtaining near-to-optimal schedules for complex models and real-world settings	Computationally and data-intensive

chemotherapeutic agent vincristine[22]; in this work, a subpopulation of cells was observed to exhibit high levels of multidrug resistance 1 (MDR1), conferring resistance to vincristine, and the growth rate of each cell type was then estimated in the presence and absence of the drug, revealing that selection plays a minor role in the emergence of MDR1-high cells. In a model of glioblastoma growth and treatment response, continuous dosing was determined to be the best strategy for slowing tumour growth while accounting for tumour heterogeneity and drug resistance[36]. This model was parameterized using in vitro treatment response data of glioblastoma cell lines and by incorporating patient pharmacokinetics data from clinical trials. Building upon a mechanism of resistance based on switching between transcriptional states in melanoma, a three-state model of drug-sensitive, pre-resistant and irreversibly resistant cells, in which pre-resistant cells act as an intermediate state between sensitive and resistant phenotypes prior to treatment with vemurafenib[20], together with estimated parameters, demonstrated that continuous dosing is optimal across tested treatments including an alternating on/off schedule.

Adaptive therapy. An alternative approach for identifying treatment schedules is adaptive therapy, in which schedules are decided based on the current state and projected trajectory of the tumour, which may lead to a maintenance strategy exploiting intratumour competition to suppress resistance[14,37]. One model of adaptive therapy, designed to modulate treatment to maintain a tumour's size[14], assumes that competition exists between tumour cells (Fig. 1e,f) and that there is a fitness tradeoff to resistance so that resistant cells are naturally suppressed by the drug-sensitive population (Fig. 1g). Other models investigating adaptive dosing strategies assume that the drug-sensitive tumour population competes with a pre-existing resistant population such that tumour control can be attained by maintaining instead of eradicating the drug-sensitive tumour cell population so that it may outcompete resistant cell clones[23,24] (Fig. 1h). Such models typically use ODEs and integrate competition into the dynamics to describe the general behaviours of the populations when noise is negligible[23,24]. ODEs also enable the investigation of long-term behaviours, such as the stability of the solutions and how interactions between cell types affect such stability. If competition exists between cells, models need to account for cell–cell interactions and their effects on tumour dynamics (Fig. 1e–g); for example, in a model in which treatment induces resistance, an increase in drug concentration does not always lead to a better treatment response, because the drug-induced transition to resistance can dominate the killing effect of high drug concentrations[29]. Using Lotka–Volterra dynamics of competition between sensitive and resistant cells, a range of parameters can be defined, including the initial cell type distribution, growth rates and transition rates, with which to investigate predictive factors of the underlying tumour dynamics leading to changes in survival under an adaptive therapy strategy[28].

Evolutionary game theory. Evolutionary game theory (EGT) has been used to understand competition and optimize adaptive therapy strategies. EGT is a mathematical formulation describing the competition, along with strategies, interactions and payoffs (that is, fitness), among cell types in a population[38]. EGT employs ODEs to understand population dynamic trends and how perturbations (for example, treatment) and cell–cell interactions affect such solutions. In models of cancer evolution that use EGT, the various 'players' are defined as populations in competition for resources[15,37,39–42]. Alternatively, a 'game' may take place between the physician and the tumour, suggesting that optimal strategies exist. For example, a physician may guide a tumour's trajectory by anticipating and monitoring resistance mechanisms to guide treatment decisions rather than by responding to resistance after it is observed[43]. In such a scenario, cancer cells are considered to be adaptable to strategies employed by the physician, a rational player,

who is able to anticipate future events. This assumption creates a series of actions, in which the physician applies treatment and the tumour adapts in response, and knowledge of the tumour's resistance mechanisms allows the physician to consider strategies that prolong the duration of response and possibly allow for a cure. Furthermore, EGT has been used to suggest strategies for combination therapy and to determine optimal schedules for multiple drugs to sensitize a population[39] or treat different mechanisms of resistance[40,44]. Such models are also used to investigate treatment response, and, owing to the cyclical nature of adaptive therapies, the stability of such strategies can be studied to make decisions about treatment[42] (Fig. 1h).

Understanding the underlying model of tumour development amid this competition between drug-sensitive and resistant cell types allows a comparison between strategies if faced with a heterogeneous tumour cell population. For example, a system of coupled ODEs describing drug-sensitive and drug-resistant cells shows that adaptive therapy may be beneficial, but that cell turnover and the rate of growth of the resistant population may serve as primary factors influencing the benefit of adaptive therapy[27]. In certain parameter regimes, however, conventional (that is, non-adaptive) therapies lead to better outcomes.

Targeting cancer metabolism

Treatment response modelling can also address the effects of therapies on the complex network of genomic, transcriptomic and metabolomic factors[45–50]; in particular, networks that may contain feedback loops and other motifs can be modelled using ODEs. For example, ODEs can be applied to study how such feedbacks affect long-term behaviours using stability analyses. Targeting cancer metabolism is the goal of many preclinical and clinical trials[51,52], and several approaches were developed to model the treatment response of such strategies. For example, the inhibition of enzymes that catalyse certain metabolic pathways influences cell proliferation rates can be investigated by characterizing the dynamics of metabolites in the pathways of glycolysis, glutaminolysis and the tricarboxylic acid cycle to identify the optimal target of metabolic enzymes for inhibiting the growth of Kirsten rat sarcoma viral oncogene homologue (KRAS)-mediated pancreatic cancer cells[45]. Using this model, the relation of target enzyme knockdown and pancreatic cancer cell growth was studied, which may also be combined with the inhibition of glutamate oxaloacetate transaminase 1 (GOT1), a key enzyme for the regulation of glutaminolysis.

Furthermore, the steady-state dynamics of these models can be exploited to analyse the levels of the gene-regulatory factors AMP-activated protein kinase (AMPK), hypoxia-inducible factor 1 (HIF1) and reactive oxygen species (ROS) in response to a treatment that inhibits the glycolytic pathway[46,47]. In particular, an ODE-based model[47] has predicted the existence of a stable, hybrid metabolic phenotype, in which both glycolysis and oxidative phosphorylation can be employed simultaneously in triple-negative breast cancer (TNBC) cells. The model established that the optimal therapeutic strategy for decreasing the aggressiveness of TNBC targets the pathways of glycolysis and oxidative phosphorylation simultaneously. A multi-scale ODE model incorporating metabolite fluxes, gene regulatory networks and tumour growth dynamics can predict tumour treatment responses by modulation of the Warburg effect, reverse Warburg effect and glutamine addiction in solid tumours[48], revealing that population-scale growth of solid tumours can be inhibited by targeting the Warburg effect, that initial growth of tumour cells may be reduced by targeting the reverse Warburg effect, and that there is no obvious response to inhibition of glutamine uptake via the metabolic pathways included in the model.

A comprehensive perturbation analysis of a dynamical systems model of tumour growth in response to treatment was used to predict the most effective combination therapy for cancer cell types using oxidative phosphorylation versus glycolysis[49]. If tumour cells obtain energy from oxidative phosphorylation, the most effective therapy promotes glucose transporter 1 and inhibits pyruvate kinase isozymes M2. However, simultaneously promoting mammalian target of rapamycin (mTOR) and NADPH oxidase is most effective in highly glycolytic cancers. An ODE-based model was also used to characterize how the metabolic traits of cancer cells depend on the availability of glucose and glutamine in their environment[50]. This model predicted the changes in the genetic and metabolic profiles of melanoma cells in response to treatment with a proto-oncogene B-RAF (BRAF) inhibitor that suppresses glutamine uptake.

Partial differential equations

In contrast to ODE models that focus on temporal dynamics, partial differential equations (PDEs) capture the spatial and temporal dynamics of tumour progression in response to cancer therapy[53–58]. A system of PDEs (Fig. 2a) can be used to predict treatment responses based on a spatio-temporal model of regulatory factors of tumour invasion, spatial competition between drug-sensitive and drug-resistant cells, and interactions between tumour cells and their microenvironment[53–55] (Fig. 2b). For example, a system of PDEs has been used to simulate the dynamics of intratumoural drug concentrations to model the treatment response of a vascular tumour consisting of a mixture of cells with high and low drug susceptibility[54]. Comparing chemotherapeutic administration by a bolus injection to continuous infusion of the same amount of drug showed that the treatment schedule by bolus is slightly better at inhibiting tumour growth and leads to a notably higher level of drug plasma concentration (related to drug toxicity) than the continuous infusion schedules. Alternatively, reaction-diffusion equations have been applied to simulate the spatio-temporal dynamics of drugs delivered to a tumour to predict the development of pre-existing and drug-induced resistant tumour subpopulations under treatment schedules with low or high drug dosages[55]; here, drug resistance of the whole population is amplified by high drug dosages through selection if resistance is pre-existing or independent of the drug; by contrast, low and medium dosages lead to higher levels of drug-induced resistance than high doses.

Angiogenesis, which involves interactions between tumour and endothelial cells, is an important factor in tumorigenesis[59,60]. Angiogenesis has been described using PDE-based models to predict responses to antiangiogenic therapy[56–58]; for example, the transport of endothelial growth factors throughout the interstitium and vessel wall was modelled to predict antiangiogenic therapy effects on tumour vessel permeability and interstitial fluid pressure[56]. Such a model predicted that treatment inducing a decreasing level of convection of endothelial growth factors reduces peritumour hyperplasia and angiogenesis, eventually reducing the occurrence of lymphatic metastasis. As an alternative to angiogenesis, tumours can also progress through vessel cooption using pre-existing vessel networks. Tumour vessel cooption and angiogenesis were modelled by describing the crosstalk between tumour and endothelial cells as well as the concentrations of endothelial growth factors and oxygen transport[58]. This model was used to investigate the effects of anti-angiogenesis and anti-cooption agents, showing that tumour progression is most effectively attenuated by a combination of both types of agent. Additionally, a sequential therapeutic strategy is more effective than a simultaneous administration strategy.

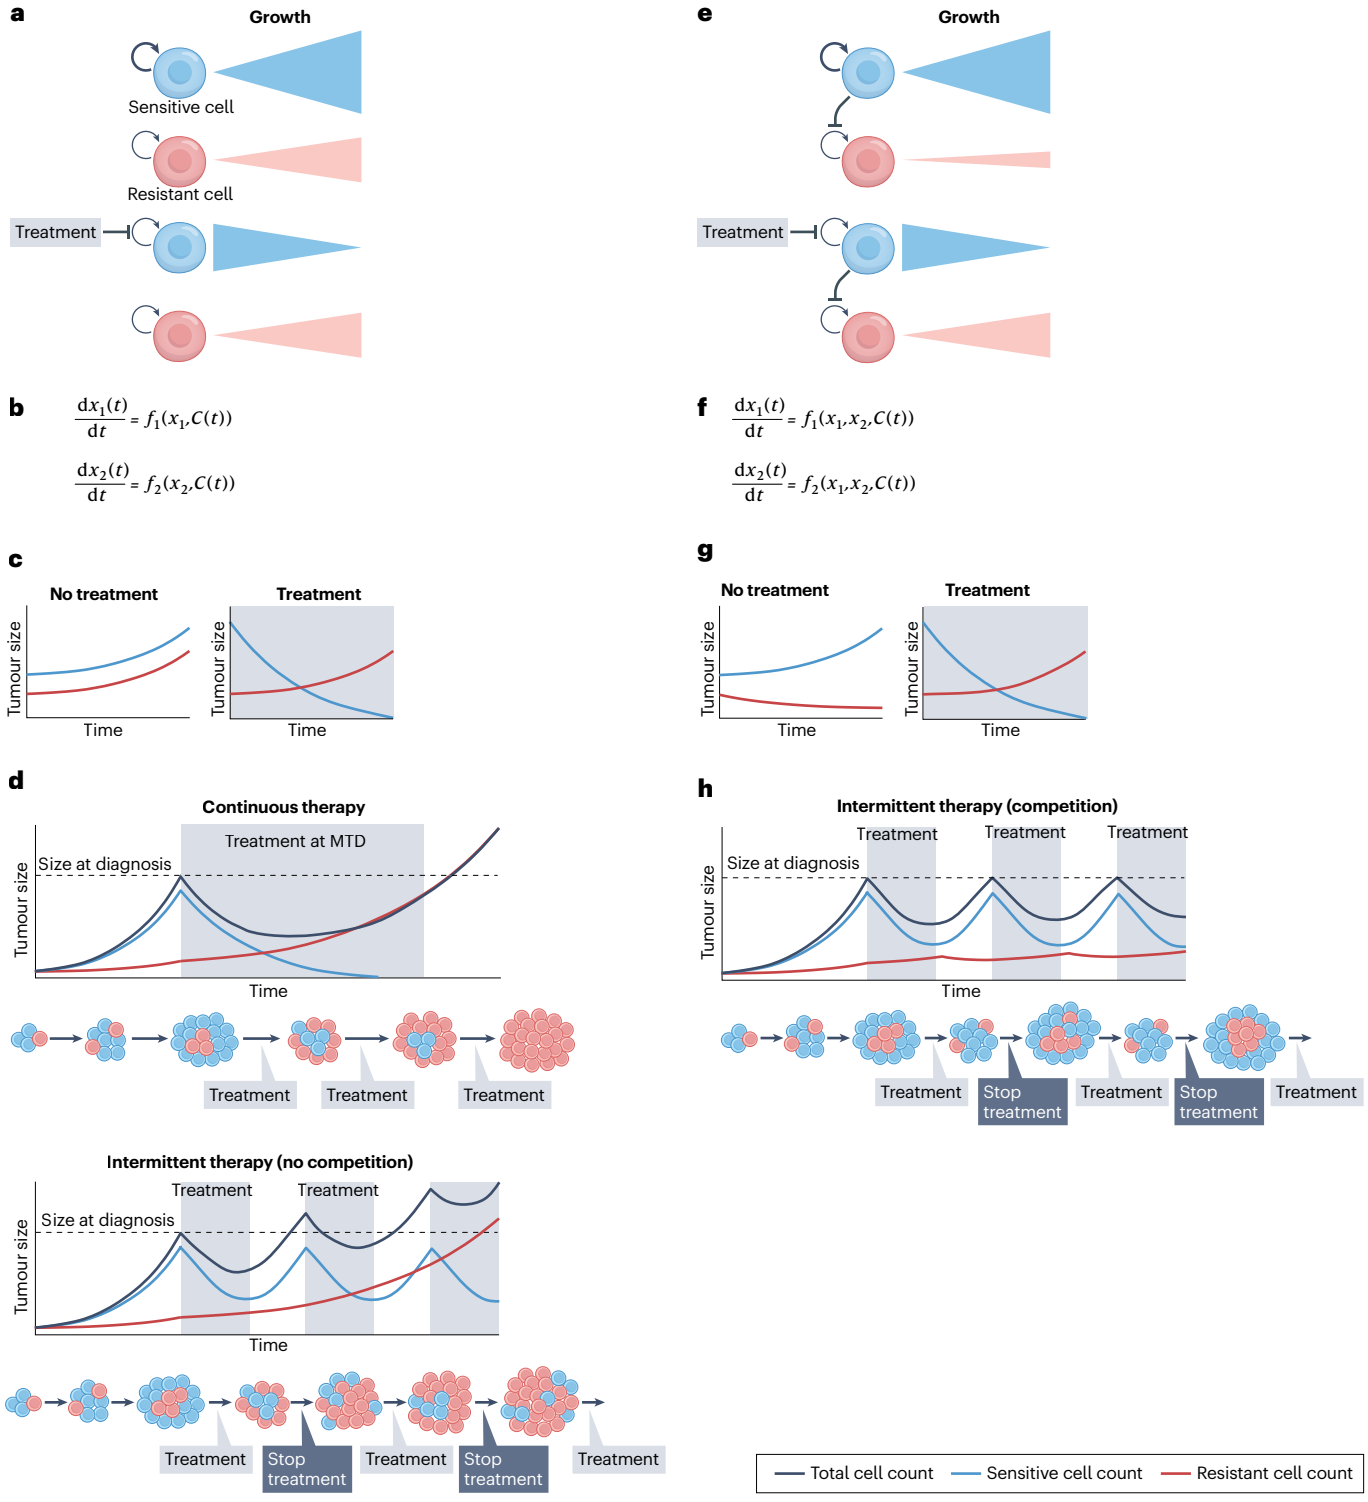

PDE-based models can also be used in modelling drug delivery to predict the treatment response of solid tumours and tumour metastases[61–63]; for example, convection–diffusion equations were applied to model drug transport in a remodelled microvascular network, induced by a solid tumour, to predict the dependence of drug delivery on different vascular networks conditions[62]. This model showed that an avascular network predicts more uniform and higher drug concentrations to reduce tumour growth than either a static or dynamic vascular network. In addition, a dynamic network predicts a more heterogeneous distribution of drug than a tumour without a

Fig. 1 | ODE-based models. a, Deterministic modelling begins by specifying the biological processes and assumptions of the proposed model. A model of sensitive and resistant cells without competition leads to undisturbed growth of the resistant population when treatment is applied. **b**, Biological models are written as a system of ordinary differential equations (ODEs) describing the dynamics of the sensitive (x_1) and resistant (x_2) populations modelled. **c**, Under different conditions, such as without or with treatment, analysis of the system of ODEs provides a trajectory over time given the initial condition. **d**, Knowledge of the dynamics under the initial conditions and environment allows the construction of different scenarios, in which the therapy can be changed over time to propose strategies for treatment. Continuous and intermittent therapies can lead to different outcomes and resistant outgrowth. **e**, Adaptive therapy incorporates competition between sensitive and resistant cells into the underlying model. **f**, ODEs for competition models depend on the state of other species in the population. **g**, In tumours without treatment, one population may shrink that would otherwise grow if treatment is applied owing to the suppression of the fitter population. **h**, These dynamics can be used to formulate strategies for maintenance of a tumour if resistance would be the likely outcome of treatment. MTD, maximally tolerated dose.

dynamic network or having no vasculature owing to increased interstitial fluid pressure and irregular capillary networks generated. Convection–diffusion equations have also been used to simulate the pharmacokinetics of doxorubicin and ado-trastuzumab, which are used to treat breast cancer brain metastases, by modelling drug delivery across the blood–brain and blood–tumour barriers using microbubbles[63]. Microbubble-based drug delivery is based on a protein or lipid shell that is used in conjunction with ultrasound to disrupt the blood–brain barrier for chemotherapy delivery[64]. This model predicted that the application of focused ultrasound in combination with microbubbles enhances the penetration of those two anti-cancer agents and improves drug delivery to the brain.

Stochastic models

ODE and PDE models represent deterministic descriptions of the average or aggregate behaviour of a biological system. They rely on the assumption that any stochasticity in the behaviour of the individual components does not relevantly affect their aggregated dynamics. However, in many situations arising in treatment response modelling, the individual stochastic components of the system (for example, cells that stochastically either self-renew, differentiate or die) are not numerous enough or sufficiently homogeneous to justify this assumption. For example, targeted cell populations might be small, as is often the case with (cancer) stem cell populations[65] or when targeting minimal residual disease[66], or when the tumour cell number substantially decreases during treatment. Moreover, tumours often contain small heterogeneous subclones that vary in sensitivity to treatment[67]. This phenomenon is especially relevant when considering the emergence of treatment resistance, which often starts with an initially rare resistant subclone[68]. In those settings, the dynamics of interest are more adequately described by a stochastic model.

Stochastic models yield a probability distribution over a set of possible outcomes (Fig. 3a). With small or heterogeneous populations, optimal treatment regimens can vary among this set of outcomes. Owing to the greater role of stochasticity, even subpopulations with comparatively higher proliferation rates, which are on average expected to increase in frequency, have non-negligible extinction probability. Thus, such stochastic fluctuations can mark the difference between resistant clones becoming predominant or being eradicated.

Treatment administration schedules often need to strike a balance between targeting different clones that respond best to different therapeutic agents. In the short term, targeting the most frequent and proliferative clone yields the largest reduction in tumour growth. However, continued treatment of a less abundant clone, until it is fully eradicated, prevents this clone from resurging when treatment is interrupted or if the dose is altered, which might result in more successful treatment outcomes in the long term. To optimize such schedules, it is important to create an accurate description of the stochastic dynamics of proliferation and possible extinction of individual clones.

Stochastic models for population dynamics

Branching processes. Branching processes can be applied to modelling the division and mutation-accumulation dynamics of cell populations[11,69–71]. In these models, individual cells are categorized by type (or state). For each type, the corresponding rates, or probabilities, of death, division or transitioning to another type are specified, and the model tracks the number of cells of different types (Fig. 3b). An advantage of this representation is that the branching behaviours of a cell and its descendants follow the same rules, so that recursive formulations can be used to analytically describe aggregate properties of these processes. Furthermore, explicit formulas for the distribution of single-type branching processes exist[69], as does a central limit theorem for multi-type branching processes[72], allowing the use of approximations, which are faster than simulation and account for variability when estimating model parameters[73].

Branching processes can be used to model the dynamics of treatment response and the development of resistance if treatment-sensitive cells may become quiescent[74], for example, to investigate the dynamics of acquired resistance owing to the T790M mutation in epidermal growth factor receptor (EGFR)-mutant non-small-cell lung cancer treated with erlotinib[75]. Here, resistant cells were identified to have a selective disadvantage in the absence of treatment, and stochastic modelling was used to identify optimum erlotinib administration strategies, which were later tested in a clinical trial[76]. A branching process model of primary and metastasized pancreatic tumour cells, parameterized with human clinical data, revealed that delay in treatment initiation can have negative consequences for survival outcomes[77]. Branching process models have further been applied for treatment optimization[12,78], for example, in the treatment of chronic myeloid leukaemia and acute myeloid leukaemia[12,78], investigating different drug cycles to improve outcomes, such as delayed emergence of resistance as compared to single-drug treatments. Similarly, such models have been used to study the dynamics of resistance to specific drugs or drug combinations and the efficacy of treatment regimens that combine several drugs[70]. Moreover, branching process models have found application in the optimization of screening strategies[79,80].

Branching processes and pharmacokinetics. Branching processes can also be combined with pharmacokinetics models to capture the variability in how administered dosages translate to effective concentrations in the target tissue, depending on patient-specific characteristics (Fig. 3c). For example, a model of cell type-specific combination responses to osimertinib and selumetinib in EGFR-mutant non-small-cell lung cancer was combined with a two-parameter exponential decay model to describe the change in drug concentration over time

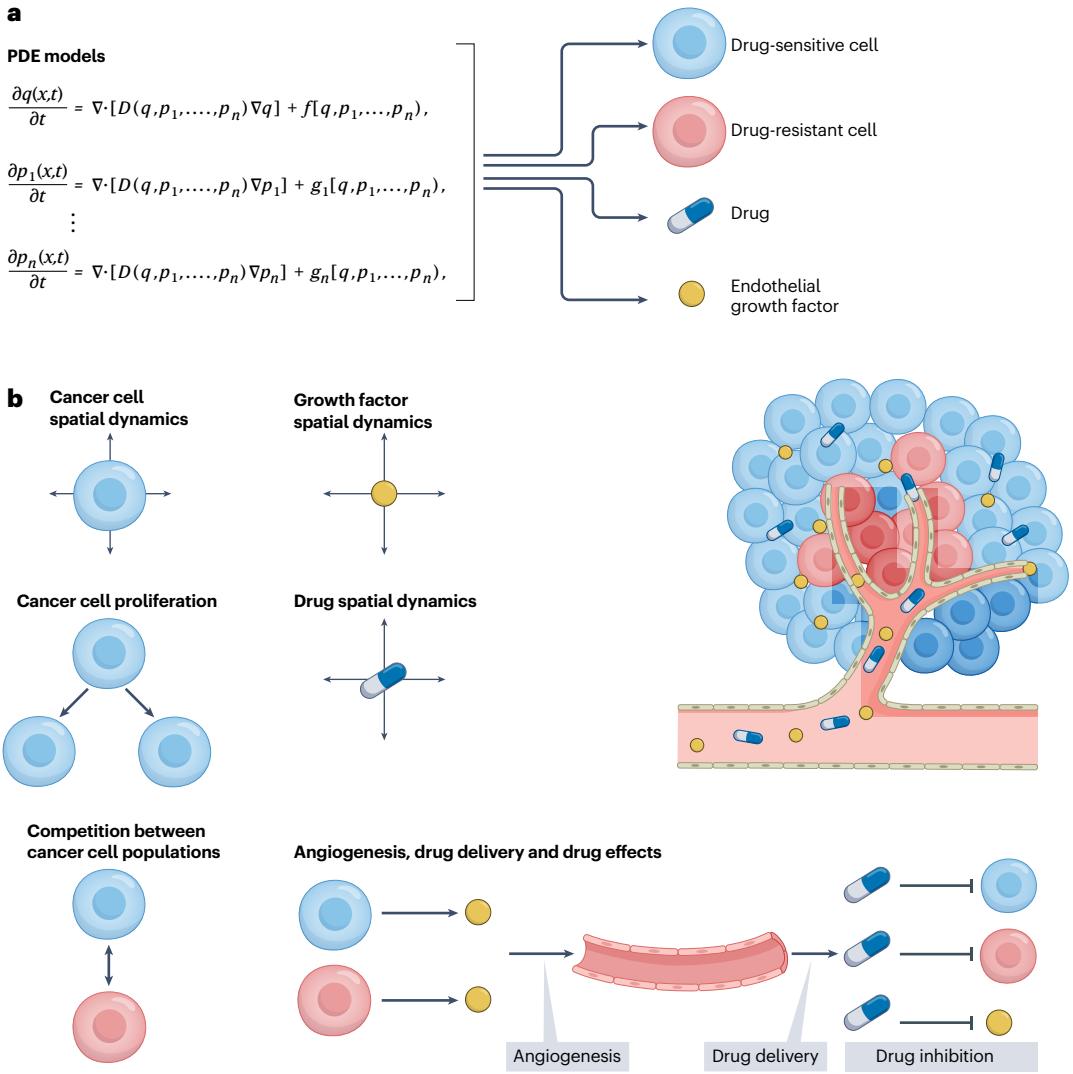

Fig. 2 | PDE models. a, Partial differential equation (PDE) models incorporate dynamics in time and space by accounting for cells, drugs and growth factors in the system along with their spatial distribution and movement over time. **b**, A PDE model of angiogenesis can model the delivery of drugs through the vasculature and extravasation into the tumour as well as diffusion in the tumour. The growth dynamics (proliferation and death) are incorporated, along with the movement and interaction with other cell populations. D, a function representing the diffusion coefficient; $f, g_1, \ldots, g_n$; functions representing the respective reaction terms; $p_1, \ldots, p_n$, density functions of different types of cells and drugs; q, a density function of endothelial growth factor; t, time; x, location.

after administration[81]; here, conditioning treatment regimens on these two patient-specific parameters notably improved treatment efficacy and/or reduced toxicity. Similarly, combination schedules of dacomitinib and osimertinib were investigated in EGFR-positive non-small-cell lung cancer[13] to analyse different dosing regimens by integrating a branching process model of the population dynamics of drug-sensitive cells and different types of resistant cells with pharmacodynamics models; here, optimal dosing schedules consist of low, frequent doses of osimertinib combined with fewer, high doses of dacomitinib, a suggestion that led to the implementation of a phase I clinical trial[82] (NCT03810807).

Branching models considering interdependence. In situations in which the type space becomes too complex or the behaviour of different cell types too interdependent, branching processes are less analytically tractable. There are, nevertheless, successful applications of branching process models in scenarios with interdependent rates, for example, to model population dynamics during immunotherapy; in this model, death rates of melanoma cells depend on the prevalence of T cells[83]. Another important case of interdependence between cells is given by capacity constraints, which cap the expansion of a cell population at some limit, and which consequently bring about a negative interdependence between the expansion of different subclones. Interdependence between lineages has also been described using logistic branching process models[77], for example, to investigate capacity-constrained expansion of pancreatic tumours and their response to treatment with the chemotherapeutic agents folfirinox and gemcitabine[84,85].

Moran and Wright–Fisher processes. The Moran process[86,87] and the Wright–Fisher process can model competition and replacement between cells and allow some analytical exploration[88] (Fig. 3b). In the Moran process, the population is fixed at a certain size and each dividing cell consequently replaces another randomly chosen cell. Division likelihoods at each time step are set to be proportional to a cell-type-specific fitness value. If fitness is constant across types, this process is entirely governed by genetic drift, whereas increasing differences in the fitness values between types result in a higher selection relative to drift. The fitness may further depend on the administered drug dose. For example, a model in which drug-sensitive cells have a higher fitness than drug-resistant cells in the absence of a drug[89] enables the investigation of the efficacy of adaptive treatment schedules, which periodically allow drug-sensitive cells to proliferate to antagonize resistant cells.

In the Wright–Fisher process (Fig. 3b), the population also remains at a fixed size, but the evolution of the cell population is modelled as a

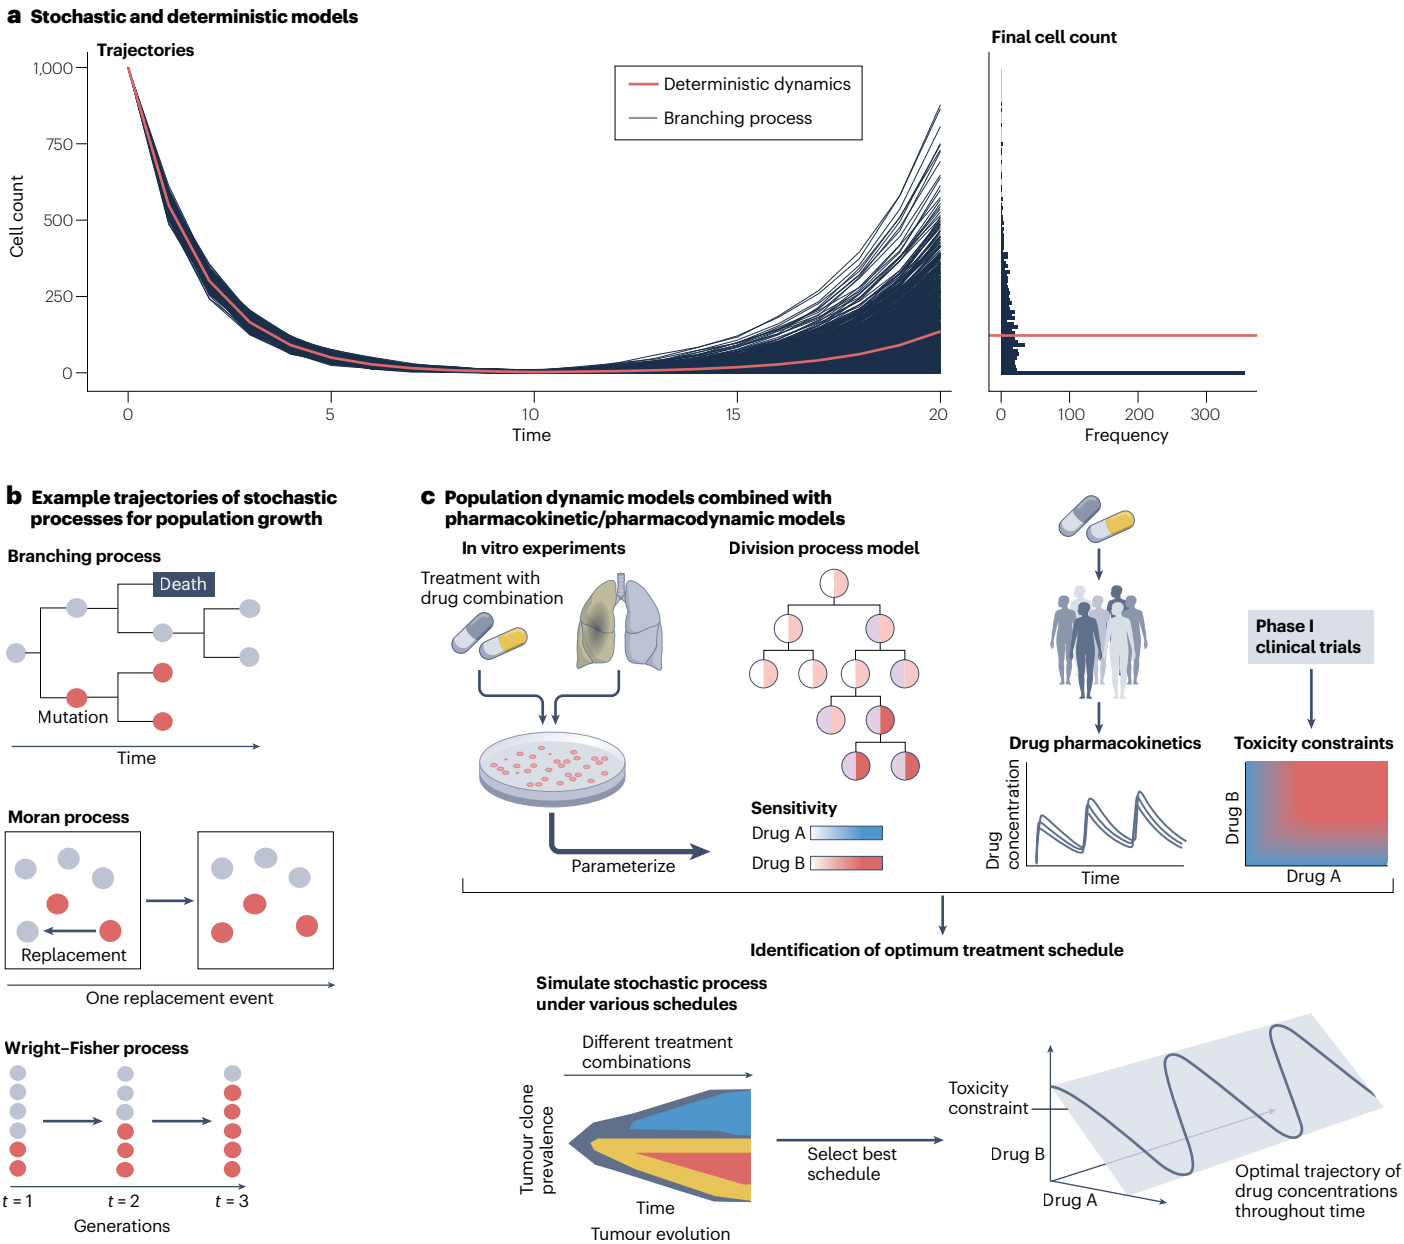

Fig. 3 | Stochastic models. a, Comparison of stochastic and deterministic models of tumour growth with 1,000 simulation runs of a branching process (black) and the deterministic dynamics (red) illustrates the effect of variability in either scenario. In the first half of the simulated time, the death rate is higher than the birth rate whereas, in the second half, the death rate is lower than the birth rate, reminiscent of a drug treatment and a withdrawal phase. The histogram shows the distribution of observed cell counts at the end of the simulations with the deterministic result (red). **b**, The branching process, Moran process and Wright–Fisher process are stochastic processes for modelling division and death. These models have different rules for division and generation time. **c**, Treatment schedule optimization workflow in frameworks that combine stochastic population dynamics models with pharmacokinetics and pharmacodynamics models.

sequence of disjoint generations. At any time step, the type composition of the current cell generation is drawn according to a distribution wherein each type has a weight proportional to a type-dependent fitness multiplied with the number of cells of this type in the previous generation. In the context of treatment response modelling, this stochastic process has been applied to demonstrate the value of monitoring in adaptive therapies[90].

Hybrid models. Hybrid models comprising systems of differential equations often remain tractable even if some variables are assumed to be stochastic. For example, a stochastic differential equation model has been used to describe intermittent androgen deprivation therapy in prostate cancer with stochasticity in growth rates and antigenicity of tumour cells[91]. Similarly, a mechanistic model that includes a noise component in the system of differential equations describing the growth dynamics of cancer cells and cytotoxic T lymphocytes allows for investigation of how the efficacy of oncolytic virus therapy targeting tumour cells depends dynamically on the prevalence of virus-specific cytotoxic T lymphocytes[92]. This framework was used to predict the likelihood of treatment success given various parameters governing cell growth and the transmission of the virus. Furthermore, an ODE-based model of glioblastoma treatment response to radiation was designed by extending the linear quadratic model such that tumour cells asynchronously exit quiescence after administration of radiation[19]. This model was applied to establish an alternative radiation administration schedule that nearly doubled the efficacy of each Gray of radiation administered compared to the clinical standard-of-care schedule, a finding that was validated in a mouse model and tested for feasibility in a pilot clinical trial (NCT03557372)[93]. Stochastic differential equation models have also been used to model the response of prostate cancer cells to intermittent hormone therapy by accounting for noise in both growth response and treatment monitoring[94]. In addition, a model of metastases and resistance to targeted therapies allows for assessment of how combinations of different targeted therapies can prolong progression-free survival[53].

Stochastic drift and diffusion models have found application in modelling tumour growth and cell migratory behaviour. For example, stochastic versions of the Gompertz diffusion process enable estimation of the effect of treatment on the growth rates of cells, which can be parameterized for uveal melanoma patient-derived xenografts[95]. Moreover, stochastic differential equations can describe the interdependent migration of epithelial and mesenchymal cells into the surrounding tumour and microenvironmental tissue[96].

Spatial and agent-based models

Agent-based models (ABMs) use simulation to investigate cancer treatment dynamics at single-cell resolution by simulating every cell's potential fates, such as movement, division and death, at each time step (Fig. 4a). Simple spatial ABMs assume that cells occupy sites on a grid (or lattice) and formulate a set of rules for how each cell interacts with neighbouring sites in the tumour microenvironment at each time step[97]. Particularly important is the ability of ABMs to capture spatial heterogeneity in treatment efficacy; for example, the tumour cell's proximity to a blood vessel could affect the effectiveness of chemotherapy[98]. ABMs can be used to address this scenario by explicitly modelling angiogenesis and the tumour vasculature[99]. ABMs are also capable of accurately modelling the spatial distribution of tumour subclones, showing that tumour evolution is dependent on its spatial structure and architecture[100]. Simpler mathematical models often assume that the tumour population is well mixed (that is, that subclones are uniformly distributed within the tumour)[101]. In solid tumours, this assumption is violated, because the process of somatic evolution implies that nearby cells are more likely to have a recent common ancestor[102]. As a result, treatment-resistant subclones are expected to occupy spatially localized regions of the tumour. Importantly, spatial diversification of cells within tumours is related to outcomes, reinforcing the notion that space is important in accurately modelling a treatment response[103].

The high resolution and flexibility of ABMs come at a price, as they are rarely mathematically tractable and must use simulation to study population behaviours[104]. However, the dynamics of some ABMs may be studied through ODEs and other methods[105,106]. As a result, efficient algorithms and high-performance computing are often necessary to simulate large tumour populations; for example, a recent ABM simulation[107] required 36 million CPU hours. Parameter tuning and optimization are also challenging, because complex ABMs can contain hundreds of parameters, many of which might be difficult to measure in patients or realistic model systems.

On-lattice versus off-lattice models

ABMs are classified as on-lattice if they assume that cells occupy sites on a grid (Fig. 4b). On-lattice extensions of birth–death processes can be used to model spatial intratumour heterogeneity and treatment resistance[108,109]. In these models, cells only replicate into unoccupied adjacent sites. These models are examples of stochastic cellular automata if time is discrete or interacting particle systems if time is continuous[110]. The spatial constraints on division restrict the growth rate to approximately t^3 (in the three-dimensional case) as opposed to approximately $\exp(t)$ in its non-spatial counterpart. On-lattice models have the advantage that they are typically more computationally efficient than off-lattice models. Several software packages can be used to simulate tumours with billions of cells; for example, Tumour Generator[108] (C++) and SITH[111] (R) implement a basic three-dimensional on-lattice model and offer the possibility of modelling treatment-resistant clones and targeted therapy. CHESS[109] (R/C++) uses a slightly different model, in which tumour cells push others out of the way during division. J-Space[112] (Julia) implements a variety of interacting particle system models and allows the generation of synthetic next-generation sequencing data from the simulated tumour. Finally, CancerSim[113] provides a basic two-dimensional on-lattice simulator for Python users.

Off-lattice models increase flexibility by allowing cells to be arranged into an arbitrary configuration (Fig. 4c). A subclass of off-lattice models, known as deformable cell models, account for cell plasticity by explicitly modelling cell shape and boundaries (using polygons, for example)[114]. The software package Cancer, Heart and Soft Tissue Environment (CHASTE)[97] simulates off- and on-lattice models under a variety of different assumptions about cell states and interactions; for example, CHASTE can be used to investigate colorectal cancer initiation, because it accurately models the structure of crypts while considering heterogeneous cell types (for example, stem cells versus differentiated cells)[115]. The software package PhysiCell[116] has also been applied to simulate interactions between immune and tumour cells, providing a useful tool for investigating response to immunotherapy.

Spatial competition limits treatment-resistant subclones

An on-lattice model of tumour growth and targeted therapy suggests that a high rate of long-range migration results in faster tumour recurrence[108]; here, increased cell motility leads to a more spatially

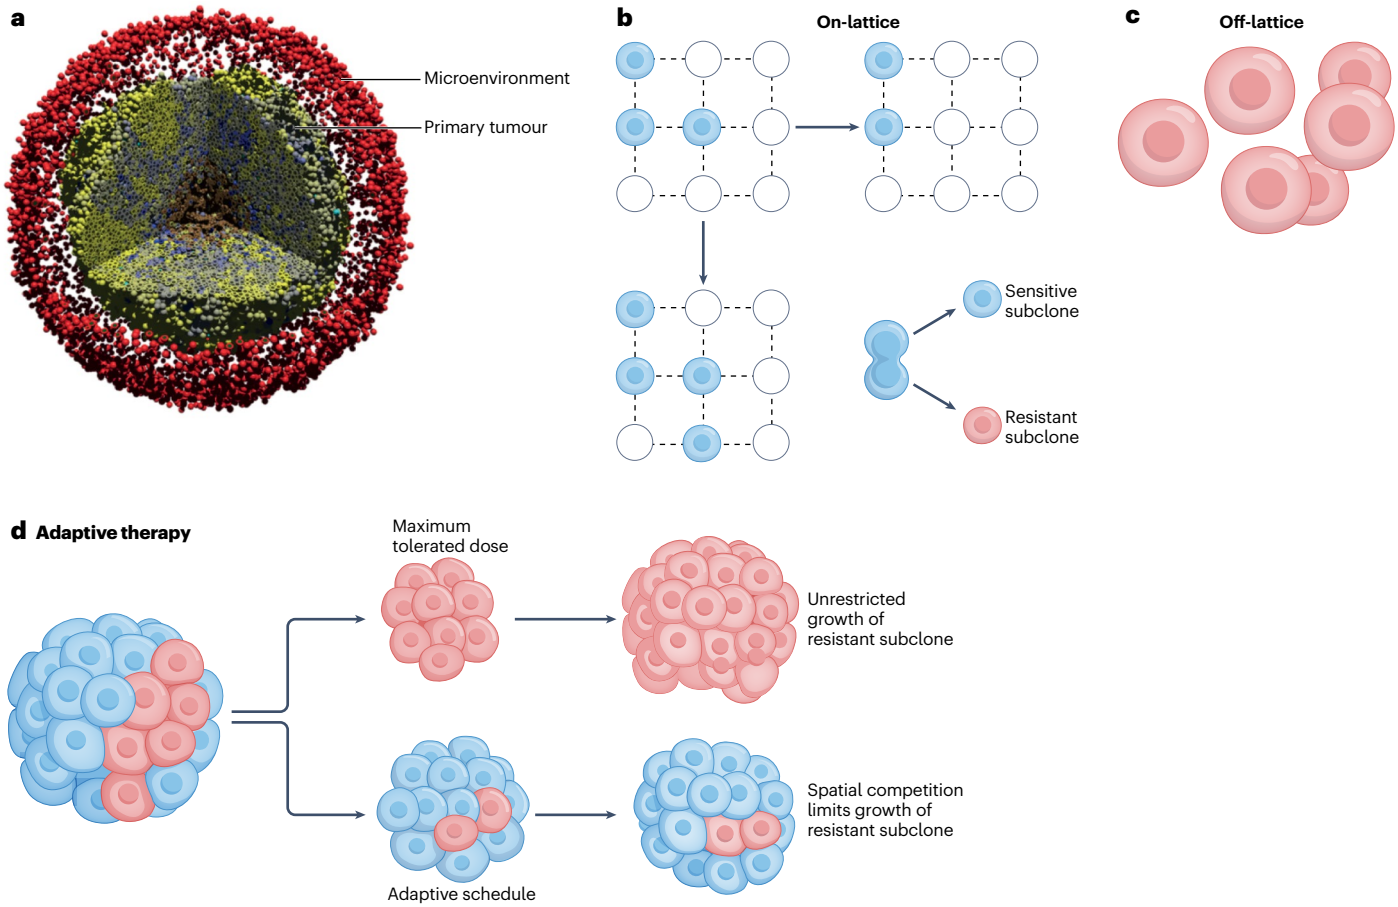

Fig. 4 | ABMs. a, An agent-based model (ABM) can account for the primary tumour (yellow cells) and the microenvironment (red cells) by modelling each cell as an individual unit. **b**, On-lattice ABMs require cells to occupy predefined spatial positions on a lattice. Movement and division occur only on the lattice. **c**, Off-lattice models do not have spatial or size constraints and allow cells more flexibility in how they move, with the added cost of computational complexity. **d**, Treatment schedules, such as adaptive therapy, can be tested by in silico experiments using ABMs that consider the spatial effects of tumours and accessibility to the drug, to model treatment effects at a local scale. Panel **a** is adapted from ref. 116, CC BY 4.0 (https://creativecommons.org/licenses/by/4.0/). Panels **b** and **d** adapted from Harvard Library, ©2021 President and Fellows of Harvard College, licensed under a Creative Commons Attribution 4.0 International License.

dispersed, resistant subclone, which results in less spatial competition (that is, faster growth rate) if targeted therapy kills the non-resistant subclone. Indeed, the non-resistant subclone could constrain the growth of the resistant subclone[14,117] (Fig. 1b), as shown using an off-lattice model comparing the conventional treatment schedule of applying the MTD to an adaptive schedule with dose modulation and treatment vacations. In heterogeneous tumours, MTD consistently led to unconstrained growth of the resistant subclone. Under the adaptive schedule, spatial competition kept the resistant subclone in check, but ultimately, no single strategy was consistently better. The finding was validated in vitro by showing that chemotherapy-sensitive MCF7 cells outcompete resistant cells in co-culture[117]; however, the presence and extent of competition in human tumours remain to be validated.

Agent-based models can also be used to test adaptive therapy strategies of multiple drugs with multiple resistance mechanisms; for example, one study using in silico experiments suggests that dose modulation tends to fare better than conventional fixed dosages across a range of parameters not meant to represent a particular drug or cancer type[118] (Fig. 4d). ABMs have also been applied to investigate the effects of space on the tumour to identify the primary features that influence competition and to determine treatment response kinetics[119]. Exploiting intratumour competition to decide on a strategy of containment or maximal cell kill, ABMs can also account for patient variability, as each patient's tumour spatial heterogeneity is likely to be unique and thus the benefits of adaptive therapy or curative therapy may also vary[28].

Hybrid models combine ABMs (addressing position of a cell in space) with ODEs or PDEs (to model continuous features of the tumour microenvironment) to consider diffusion of chemokines, drugs or oxygen[120]; for example, in a hybrid model, cells were assumed to occupy lattice sites and oxygen and drug concentrations were modelled with diffusion equations[121]. In this model, the relative benefit of an adaptive schedule (as opposed to MTD) was quantified as a function of the fitness of resistant cells. An extension to the PhysiCell ABM framework[116], called PhysiPKPD, introduced pharmacodynamics and pharmocokinetics modelling of drug delivery using ODEs in addition to ABMs that model tumour cells and local diffusion of the drug after extravasation

into the tumour[122]. This framework allows in silico experiments of different drug schedules and accounts for the administration and distribution of the drug in the body.

Modelling the tumour microenvironment during treatment

The flexibility of ABMs allows them to capture the idiosyncrasies of tumour microenvironments that cannot be considered in standard mathematical models. For example, glioblastomas have a distinct spatial structure (called a perivascular niche), in which a blood vessel is surrounded by resistant stem-like cells and sensitive differentiated tumour cells (Fig. 1c). After radiation, cells can transition between the differentiated and de-differentiated states. An ABM was formulated that captures the structure of the perivascular niche, accounting for cell-state switching during treatment as well as the pharmacokinetics of the chemotherapeutic agent temozolomide[107]. This ABM was used to find the optimal schedule of combining temozolomide with radiation, which was validated in a mouse model, leading to notably improved survival times as compared to a scramble control schedule.

ABMs can reduce modelling assumptions and increase robustness by using genomic data as a baseline; for example, an agent-based model parameterized by data from fluorescence in situ hybridization (FISH) experiments was applied to study genetic heterogeneity and phenotypic diversity at the single cell level in a cohort of pre- and post-treatment breast tumours[123], showing that cells post-treatment strongly cluster by phenotype (as opposed to genotype). This spatial clustering can in principle be due to phenotypic switching, cell motility or both, and delineation of the contribution of both factors is of clinical interest. An ABM based on CHASTE that does not include motility or phenotypic switching was found to generate substantially less spatial clustering compared to the observed data[123], suggesting that switching and motility take place in this cancer type and differ between subtypes of the disease. Importantly, the ABM uses the FISH results from the baseline samples to specify the initial coordinates and phenotypes of cells, highlighting the ability of ABMs to use real data to bypass additional modelling assumptions. These examples demonstrate the usefulness of ABM for studying cancer treatment response, in particular, if spatial considerations are essential for accurately describing a cell's response to therapy.

Control theory and machine learning approaches

A longstanding aim of computational approaches in cancer research is the optimal scheduling of therapies to maximize overall patient benefit. Optimal control theory (OCT) is a branch of mathematics concerned with identifying the interventions in dynamical systems that produce the best possible outcomes, providing another framework for modelling. For example, OCT can be applied to balance treatment toxicity with efficacy, and to assess tumour–immune dynamics, tumour–virus interactions, and the development of therapy resistance under treatment. Similarly, reinforcement learning can be used to improve cancer therapy scheduling. Reinforcement learning and OCT both attempt to automate learning and decisions via a controller (intervention) that optimizes a system's behaviour over time. OCT can also be characterized as a subtype of reinforcement learning, given that reinforcement learning may be more broadly defined as a class of methods that attempt to understand and automate decisions under a framework, in which the decision-maker (or agent) directly interacts with its environments, actions and rewards to maximize an objective[124]. Here, we distinguish between the two by focusing on the differences in their approaches to the same class of problem.

Optimal control theory

In OCT, externally manipulable variables (controls) are set such that some cost is minimized with respect to observed variables[125]. In a typical cancer therapy setting, the observed variables might describe time-varying tumour population sizes, whereas the controls might describe time-varying drug concentration, and the cost is the final number of tumour cells at the end of treatment and the total drug doses applied over the course of treatment[126] (Fig. 5a). For example, OCT has been used to model the evolving total number of tumour cells in multiple myeloma under treatment with the chemotherapeutic agents melphalan, cyclophosphamide and prednisone[127]; here, tumour cells are assumed to follow carrying capacity-limited growth, and a saturating effect of drug treatment at high doses is considered to investigate the optimal dose schedule and minimize the total administered dose under an equality constraint on the final tumour size. This approach showed that the optimal dose decreases over the course of the treatment.

Because of these considerations, OCT allows the optimal scheduling of cancer therapies[23,24,107,126,128–138]; for example, by generating models of tumour response to immunotherapies[129–131] or viral therapies[132–134], and of the dynamic outgrowth of therapy-resistant clones under treatment[23,24,42,128,135,139]. Optimal antigen receptor therapy dosing schedules have been defined by modelling the interplay between tumour cells, osteoblasts and osteoclasts in the bone metastatic niche, accounting for life stage and changes in bone remodelling dynamics with age[129]. Here, the optimal schedules for different life stages and parameter settings vary in the total amount of administered drug and the timing of change points; however, each schedule consistently contains an initial fixed dose phase followed by a period of dose decrease.

OCT approaches have also been used to optimize treatments when considering situations in which resistance may arise. Here, the goal is to maximize time to progression with minimal drug administered (Fig. 5a); for example, optimal chemotherapy schedules were calculated that maximize the time until tumours reach a threshold total size, if tumours comprise a mixed population of therapy-sensitive and fully resistant cells under three scenarios of population growth dynamics: unconstrained exponential growth, carrying-capacity-limited logistic and Gompertzian growth[139]. The optimal approach under exponential growth involves eradicating the entire sensitive tumour population with treatment, whereas the optimal strategy for Gompertzian growth maintains the tumour at a large size. Under logistic growth, the choice of whether to maintain the tumour at a large or small size has limited effect on the success of the therapy. In addition, OCT was applied to study an evolutionary game theory model of castration-resistant prostate cancer dynamics[15,135]. By finding a stable, equilibrium point between cell types, the hormone therapy abiraterone was administered to steer newly diagnosed tumours towards this point, and a simple titration protocol was defined that gradually increases the dose before reaching a plateau of the average optimized approach across varying initial conditions. This strategy was predicted to outperform an aggressive treatment and an adaptive therapy approach. Furthermore, a model in which treatment can cause sensitive tumour cells to become resistant in a dose-dependent manner[128] applies OCT to calculate an optimal schedule that minimizes the probability of resistance emerging from an initially fully sensitive population, which is close to a constant dose.

Reinforcement learning

Reinforcement learning is a subfield of machine learning that encompasses data-driven methods to learn policies for sequential

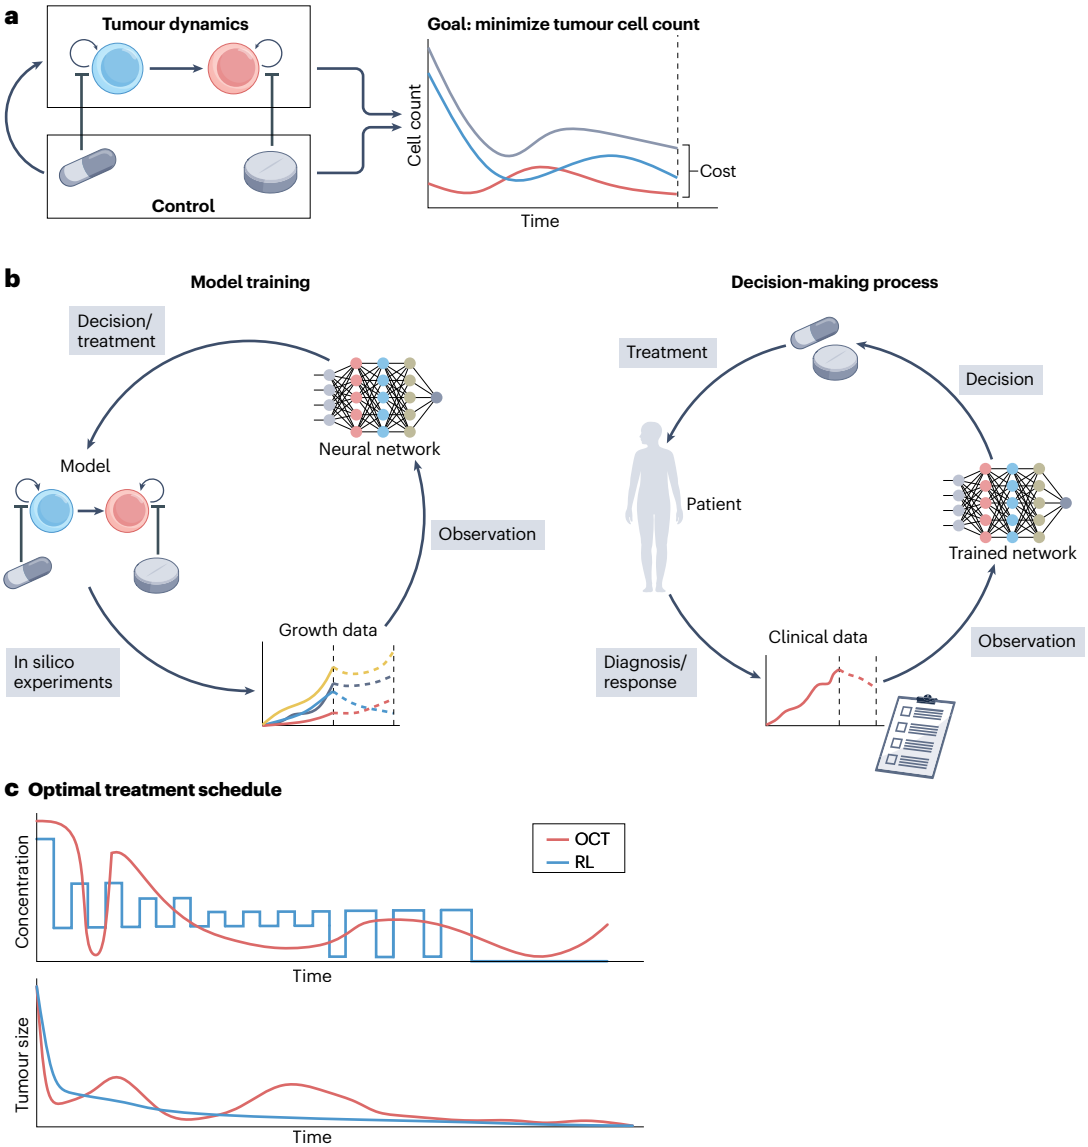

Fig. 5 | OCT. a, Optimal control theory (OCT) approaches seek to determine a function (treatment strategy) that affects the growth dynamics of the tumour in a way that minimizes the selected cost function (for example, tumour cell count). The growth dynamics are specified as part of the model that should be optimized. **b**, Reinforcement learning (RL) for tumour dynamics can use in silico experiments to train a neural network to make a decision and to optimize the decision based on some measure of cost or reward. The trained model can then be used to predict the optimal decision and treatment for a patient by incorporating clinical data and observed trajectories, leading to updated decisions. **c**, The resulting optimal treatment schedule under RL may be directly comparable to an OCT schedule without requiring the specified model, using only observation to steer the neural network outputs toward a locally optimal solution.

decision-making[124,140]. There are, broadly, two related ways in which methods from reinforcement learning are becoming relevant for computational approaches to scheduling cancer therapy.

First, reinforcement learning methods provide a flexible alternative to studying therapy scheduling under mathematical models of cancer therapy in situations in which features of the model make traditional OCT challenging (Fig. 5b). In these cases, reinforcement learning methods have been used in conjunction with extensive model simulation to learn near-to-optimal policies under the model in a data-driven way[141–147]. In addition, stochastic optimal control methods exist[128]; however, high-dimensional model state spaces and nonlinear dynamics pose challenges for traditional control methods[148]. Reinforcement learning methods have been used in conjunction with extensive model simulation to learn policies that perform well under the model[141–147]. For example, a deep neural network reinforcement learning model of cancer evolution under chemotherapy treatment was trained on in silico patients over the course of treatment[141]. In this approach, the scheduling algorithm can control tumour cell proliferation in the context of emerging resistance. Similarly, a reinforcement learning model was parameterized by a deep neural network and trained using simulated

tumour growth data to suggest chemotherapy dosing schedules[142]; here, the trained algorithm performs competitively compared to an approach based on a traditional OCT analysis of the model (Fig. 5c).

Second, these methods may also be used in cancer therapy scheduling by learning successful policies directly from high-resolution patient data[149–151]. For example, a neural network-based reinforcement learning framework based on patient-specific cytokine, imaging and dosage features was used to suggest dose adaptations for patients with non-small-cell lung cancer who undergo radiation[149]. The framework was trained on data from 114 patients, and the resulting treatment recommendations show common features with decisions recommended by the treating physicians, motivating further testing of the performance on independent data sets. Similarly, a neural network-based reinforcement learning framework could make recommendations for a sequence of three clinical decisions in the care of patients with oropharyngeal squamous carcinoma[150]. The decision-making system was trained on patient history features using retrospective treatment response data from 402 patients and tested using held-out data from 134 patients. Using a predictive model trained on the same data set, it was predicted that the recommended decisions would have led to a small improvement in survival and dysphagia rate compared to standard treatments.

Outlook

Computational approaches can address a variety of questions in cancer treatment response dynamics and optimization, including the description of interactions between cell types and prediction of treatment outcomes to identify best treatment strategies, some of which have been tested in clinical trials (NCT03557372 (ref. 93), NCT03810807 (ref. 82), NCT02415621 (ref. 152), NCT03511196 (ref. 153), NCT03543969 (ref. 154), NCT03630120 (ref. 155)) (Table 2). However, challenges remain to be addressed to enable the broad clinical application of computational modelling in cancer treatment (Box 2). Importantly, the connection between data and model should be strengthened, both in the design of the model based on data-driven insights in cancer biology and treatment resistance mechanisms, and when measuring tumour burden and the pre-existence and/or generation and outgrowth of resistant clones over time. For example, methods are available for monitoring and quantifying the degree of resistance present at the start of treatment[14,27,117]. Similarly, tumour burden over time can be measured, for example, by assessing prostate-specific antigen (PSA) levels in prostate cancer, lactate dehydrogenase (LDH) in melanoma[28] and circulating tumour DNA (ctDNA) in metastatic colorectal cancer[50]. Moreover, single-cell mRNA expression data can be used to identify and determine distinct transcriptional states with respect to BRAF inhibitor response in melanoma[156]. To make clinically relevant predictions and recommendations, the model design must reflect the true behaviour of the modelled system. Model fitting can, however, be more challenging in certain situations; for example, model parameterization is more complex for stochastic models than for deterministic models, as the variance of a process is itself an important component of stochastic models; however, this variance might be obfuscated by further sources of variability in the measurement or by experimental processes. Model fitting in deterministic models consists of eliciting a mean behaviour of a system and removing noise in the system; by contrast, fitting stochastic models must additionally master the challenge of disentangling the intrinsic variance of a process from external sources of variance. Similarly, the design parameterization of large complex computational models such as ABMs relies on data input from multiple sources so that cell behaviour, localization and interactions can be adequately described. Furthermore, modelling approaches may rely on assumptions that might not be accurate for specific situations, which can lead to overfitted or oversimplified models. The next generation of models should take advantage of the large amount of available patient data to produce more accurate and better parameterized descriptions of the biological system of interest.

Approaches could also be designed that learn successful treatment policies directly from patient data; for example, historic patient time-course data has been used to train a reinforcement learning model to generate treatment suggestions for adult patients with sepsis in the intensive care unit[157]; here, in an independent cohort, mortality was lowest for patients whose actual received treatment matched what the system would have recommended, and mortality increased in a dose-dependent manner for patients whose treatment deviated from the recommendations. As we take advantage of the era of big data in cancer research, it is becoming increasingly possible to collect data with high temporal, spatial and molecular resolution[136]. In the future, it may be possible to use approaches, such as those outlined in this Review article, to train directly on patient treatment data to obtain improvements in patient care.

Table 2 | Clinical trials of computational models in cancer treatment

Clinical trial identified	Disease	Intervention	Status	Refs.
NCT03557372	Recurrent glioblastoma (stage IV)	Radiation schedule developed with a mathematical model	Phase I completed	19,158
NCT03810807	EGFR-mutant lung cancer	Model-optimized combination of dacomitinib and osimertinib	Active, phase 1	13
NCT02415621	Prostate cancer	Adaptive dosing of abiraterone based on an evolutionary game theoretic model	Active, early phase 1	15
NCT03511196	Prostate cancer	Adaptive androgen deprivation therapy in combination with abiraterone and prednisone	Active, early phase 1	159
NCT03543969	BRAF-mutant metastatic melanoma	Adaptive dosing of vemurafenib and cobimetinib	Active, recruiting, early phase 1	–
NCT03630120	Thyroid cancer	Adaptive dosing of lenvatinib, sorafenib, cabozantinib and vandetanib (depending on subtype)	Phase 2 terminated owing to lack of efficacy	–
NCT01967095	EGFR-mutant lung cancer	Erlotinib dosing schedule developed using a stochastic birth–death process	Phase I completed	75,76

BRAF, proto-oncogene B-RAF; EGFR, epidermal growth factor receptor.

Box 2

Translational considerations

The translation of a predicted optimum treatment schedule based on mathematical modelling to a clinical trial is the ultimate test of the accuracy of the prediction. However, the clinical translation of computational models in cancer treatment faces multiple challenges. First, the predicted schedule must compete with other clinical trials. There might be a small patient pool available that is eligible to enroll, and competition between a modelling-suggested trial and other trials must be overcome. Second, as in trials designed without the use of modelling approaches, preclinical evidence must be carefully investigated to conclude whether a clinical benefit is expected. Statistical methodology[158] to compare the expected results of mouse and human trials might be helpful to increase the confidence that positive results might be achieved. Third, the high degree of variability of human patients in terms of inter- and intra-tumour heterogeneity as well as variability in drug absorption and metabolism kinetics might require a large sample size to observe efficacy; an underestimation of the extent of this variability might lead to smaller than needed sample sizes and inconclusive results. Although this issue is not specific for modelling-suggested translation, it can be addressed using computational models, providing an opportunity for overcoming the challenges that all clinical trials face.

Citation diversity statement

The authors acknowledge that papers authored by scholars from historically excluded groups are systematically under-cited. Here, the authors have made every attempt to reference relevant papers in a manner that is equitable in terms of racial, ethnic, gender and geographical representation.

Published online: 19 July 2023

References

1. Jiang, P. et al. Big data in basic and translational cancer research. *Nat. Rev. Cancer* **22**, 625–639 (2022).
2. Basu, A. et al. An interactive resource to identify cancer genetic and lineage dependencies targeted by small molecules. *Cell* **154**, 1151–1161 (2013).
3. Van Allen, E. M. et al. Whole-exome sequencing and clinical interpretation of formalin-fixed, paraffin-embedded tumor samples to guide precision cancer medicine. *Nat. Med.* **20**, 682–688 (2014).
4. Xue, J.-M., Liu, Y., Wan, L.-H. & Zhu, Y.-X. Comprehensive analysis of differential gene expression to identify common gene signatures in multiple cancers. *Med. Sci. Monit.* **26**, e919953 (2020).
5. Song, Q. et al. Proteomic analysis reveals key differences between squamous cell carcinomas and adenocarcinomas across multiple tissues. *Nat. Commun.* **13**, 4167 (2022).
6. Reznik, E. et al. A landscape of metabolic variation across tumor types. *Cell Syst.* **6**, 301–313.e3 (2018).
7. Büttner, M., Miao, Z., Wolf, F. A., Teichmann, S. A. & Theis, F. J. A test metric for assessing single-cell RNA-seq batch correction. *Nat. Methods* **16**, 43–49 (2019).
8. Sammut, S.-J. et al. Multi-omic machine learning predictor of breast cancer therapy response. *Nature* **601**, 623–629 (2022).
9. Gallasch, R., Efremova, M., Charoentong, P., Hackl, H. & Trajanoski, Z. Mathematical models for translational and clinical oncology. *J. Clin. Bioinform.* **3**, 23 (2013).
10. Altrock, P. M., Liu, L. L. & Michor, F. The mathematics of cancer: integrating quantitative models. *Nat. Rev. Cancer* **15**, 730–745 (2015).
11. Foo, J. & Michor, F. Evolution of resistance to targeted anti-cancer therapies during continuous and pulsed administration strategies. *PLoS Comput. Biol.* **5**, e1000557 (2009).
12. Yang, J., Lindström, H. J. G. & Friedman, R. Combating drug resistance in acute myeloid leukaemia by drug rotations: the effects of quizartinib and pexidartinib. *Cancer Cell Int.* **21**, 198 (2021).
13. Poels, K. E. et al. Identification of optimal dosing schedules of dacomitinib and osimertinib for a phase I/II trial in advanced EGFR-mutant non-small cell lung cancer. *Nat. Commun.* **12**, 3697 (2021).
14. Gatenby, R. A., Silva, A. S., Gillies, R. J. & Frieden, B. R. Adaptive therapy. *Cancer Res.* **69**, 4894–4903 (2009).
15. Zhang, J., Cunningham, J. J., Brown, J. S. & Gatenby, R. A. Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. *Nat. Commun.* **8**, 1816 (2017).
16. Luria, S. E. & Delbrück, M. Mutations of bacteria from virus sensitivity to virus resistance. *Genetics* **28**, 491–511 (1943).
17. Gardner, S. N. Modeling multi-drug chemotherapy: tailoring treatment to individuals. *J. Theor. Biol.* **214**, 181–207 (2002).
18. Michelson, S. & Leith, J. T. Effects of differential cell kill on the dynamic composition of heterogeneous tumors. *Comput. Math. Appl.* **20**, 149–159 (1990).
19. Leder, K. et al. Mathematical modeling of PDGF-driven glioblastoma reveals optimized radiation dosing schedules. *Cell* **156**, 603–616 (2014).
20. Paryad-Zanjani, S., Saint-Antoine, M. M. & Singh, A. Optimal scheduling of therapy to delay cancer drug resistance. *IFAC-Pap.* **54**, 239–244 (2021).
21. Lea, D. E. & Catcheside, D. G. The mechanism of the induction by radiation of chromosome aberrations in *Tradescantia*. *J. Genet.* **44**, 216–245 (1942).
22. Pisco, A. O. et al. Non-Darwinian dynamics in therapy-induced cancer drug resistance. *Nat. Commun.* **4**, 2467 (2013).
23. Greene, J. M., Gevertz, J. L. & Sontag, E. D. Mathematical approach to differentiate spontaneous and induced evolution to drug resistance during cancer treatment. *JCO Clin. Cancer Inform.* **3**, 1–20 (2019).
24. Greene, J. M., Sanchez-Tapia, C. & Sontag, E. D. Mathematical details on a cancer resistance model. *Front. Bioeng. Biotechnol.* **8**, 501 (2020).
25. Johnson, K. E. et al. Integrating transcriptomics and bulk time course data into a mathematical framework to describe and predict therapeutic resistance in cancer. *Phys. Biol.* **18**, 016001 (2020).
26. Owolabi, K. M. & Shikongo, A. Mathematical modelling of multi-mutation and drug resistance model with fractional derivative. *Alex. Eng. J.* **59**, 2291–2304 (2020).
27. Strobl, M. A. R. et al. Turnover modulates the need for a cost of resistance in adaptive therapy. *Cancer Res.* **81**, 1135–1147 (2021).
28. Kim, E., Brown, J. S., Eroglu, Z. & Anderson, A. R. A. Adaptive therapy for metastatic melanoma: predictions from patient calibrated mathematical models. *Cancers* **13**, 823 (2021).
29. Angelini, E., Wang, Y., Zhou, J. X., Qian, H. & Huang, S. A model for the intrinsic limit of cancer therapy: duality of treatment-induced cell death and treatment-induced stemness. *PLoS Comput. Biol.* **18**, e1010319 (2022).
30. Fröhlich, F. et al. Efficient parameter estimation enables the prediction of drug response using a mechanistic pan-cancer pathway model. *Cell Syst.* **7**, 567–579.e6 (2018).
31. Aghamiri, S. S., Amin, R. & Helikar, T. Recent applications of quantitative systems pharmacology and machine learning models across diseases. *J. Pharmacokinet. Pharmacodyn.* **49**, 19–37 (2022).
32. Neftel, C. et al. An integrative model of cellular states, plasticity, and genetics for glioblastoma. *Cell* **178**, 835–849.e21 (2019).
33. Shaffer, S. M. et al. Memory sequencing reveals heritable single-cell gene expression programs associated with distinct cellular behaviors. *Cell* **182**, 947–959.e17 (2020).
34. Rambow, F., Marine, J.-C. & Goding, C. R. Melanoma plasticity and phenotypic diversity: therapeutic barriers and opportunities. *Genes Dev.* **33**, 1295–1318 (2019).
35. Fukui, R. et al. Tumor radioresistance caused by radiation-induced changes of stem-like cell content and sub-lethal damage repair capability. *Sci. Rep.* **12**, 1056 (2022).
36. Stein, S., Zhao, R., Haeno, H., Vivanco, I. & Michor, F. Mathematical modeling identifies optimum lapatinib dosing schedules for the treatment of glioblastoma patients. *PLoS Comput. Biol.* **14**, e1005924 (2018).
37. Gillies, R. J., Verduzco, D. & Gatenby, R. A. Evolutionary dynamics of carcinogenesis and why targeted therapy does not work. *Nat. Rev. Cancer* **12**, 487–493 (2012).
38. Smith, J. M. & Price, G. R. The logic of animal conflict. *Nature* **246**, 15–18 (1973).
39. Basanta, D., Gatenby, R. A. & Anderson, A. R. A. Exploiting evolution to treat drug resistance: combination therapy and the double bind. *Mol. Pharm.* **9**, 914–921 (2012).
40. Orlando, P. A., Gatenby, R. A. & Brown, J. S. Cancer treatment as a game: integrating evolutionary game theory into the optimal control of chemotherapy. *Phys. Biol.* **9**, 065007 (2012).
41. West, J. B. et al. Multidrug cancer therapy in metastatic castrate-resistant prostate cancer: an evolution-based strategy. *Clin. Cancer Res.* **25**, 4413–4421 (2019).
42. Gluzman, M., Scott, J. G. & Vladimirsky, A. Optimizing adaptive cancer therapy: dynamic programming and evolutionary game theory. *Proc. Biol. Sci.* **287**, 20192454 (2020).
43. Stanková, K., Brown, J. S., Dalton, W. S. & Gatenby, R. A. Optimizing cancer treatment using game theory: a review. *JAMA Oncol.* **5**, 96–103 (2019).
44. West, J. et al. Towards multidrug adaptive therapy. *Cancer Res.* **80**, 1578–1589 (2020).
45. Roy, M. & Finley, S. D. Computational model predicts the effects of targeting cellular metabolism in pancreatic cancer. *Front. Physiol.* **8**, 217 (2017).

46. Yu, L. et al. Modeling the genetic regulation of cancer metabolism: interplay between glycolysis and oxidative phosphorylation. *Cancer Res.* **77**, 1564–1574 (2017).
47. Jia, D. et al. Elucidating cancer metabolic plasticity by coupling gene regulation with metabolic pathways. *Proc. Natl Acad. Sci. USA* **116**, 3909–3918 (2019).
48. Shan, M., Dai, D., Vudem, A., Varner, J. D. & Stroock, A. D. Multi-scale computational study of the Warburg effect, reverse Warburg effect and glutamine addiction in solid tumors. *PLoS Comput. Biol.* **14**, e1006584 (2018).
49. Li, W. & Wang, J. Uncovering the underlying mechanisms of cancer metabolism through the landscapes and probability flux quantifications. *iScience* **23**, 101002 (2020).
50. Tripathi, S. et al. A mechanistic modeling framework reveals the key principles underlying tumor metabolism. *PLoS Comput. Biol.* **18**, e1009841 (2022).
51. Vernieri, C. et al. Targeting cancer metabolism: dietary and pharmacologic interventions. *Cancer Discov.* **6**, 1315–1333 (2016).
52. Stine, Z. E., Schug, Z. T., Salvino, J. M. & Dang, C. V. Targeting cancer metabolism in the era of precision oncology. *Nat. Rev. Drug Discov.* **21**, 141–162 (2022).
53. Sun, X., Bao, J. & Shao, Y. Mathematical modeling of therapy-induced cancer drug resistance: connecting cancer mechanisms to population survival rates. *Sci. Rep.* **6**, 22498 (2016).
54. Jackson, T. L. & Byrne, H. M. A mathematical model to study the effects of drug resistance and vasculature on the response of solid tumors to chemotherapy. *Math. Biosci.* **164**, 17–38 (2000).
55. Hamis, S., Nithiarasu, P. & Powathil, G. G. What does not kill a tumour may make it stronger: in silico insights into chemotherapeutic drug resistance. *J. Theor. Biol.* **454**, 253–267 (2018).
56. Jain, R. K., Tong, R. T. & Munn, L. L. Effect of vascular normalization by antiangiogenic therapy on interstitial hypertension, peritumor edema, and lymphatic metastasis: insights from a mathematical model. *Cancer Res.* **67**, 2729–2735 (2007).
57. Zheng, X. et al. A continuous model of angiogenesis: initiation, extension, and maturation of new blood vessels modulated by vascular endothelial growth factor, angiopoietins, platelet-derived growth factor-B, and pericytes. *Discret. Contin. Dyn. Syst. B* **18**, 1109–1154 (2013).
58. Voutouri, C. et al. Experimental and computational analyses reveal dynamics of tumor vessel cooption and optimal treatment strategies. *Proc. Natl Acad. Sci. USA* **116**, 2662–2671 (2019).
59. De Palma, M., Biziato, D. & Petrova, T. V. Microenvironmental regulation of tumour angiogenesis. *Nat. Rev. Cancer* **17**, 457–474 (2017).
60. Lugano, R., Ramachandran, M. & Dimberg, A. Tumor angiogenesis: causes, consequences, challenges and opportunities. *Cell. Mol. Life Sci.* **77**, 1745–1770 (2020).
61. Stylianopoulos, T. The solid mechanics of cancer and strategies for improved therapy. *J. Biomech. Eng.* **139**, 4034991 (2017).
62. Sefidgar, M. et al. Numerical modeling of drug delivery in a dynamic solid tumor microvasculature. *Microvasc. Res.* **99**, 43–56 (2015).
63. Arvanitis, C. D. et al. Mechanisms of enhanced drug delivery in brain metastases with focused ultrasound-induced blood–tumor barrier disruption. *Proc. Natl Acad. Sci. USA* **115**, E8717–E8726 (2018).
64. Mainprize, T. et al. Blood–brain barrier opening in primary brain tumors with non-invasive MR-guided focused ultrasound: a clinical safety and feasibility study. *Sci. Rep.* **9**, 321 (2019).
65. Ischenko, I., Seeliger, H., Schaffer, M., Jauch, K.-W. & Bruns, C. J. Cancer stem cells: how can we target them? *Curr. Med. Chem.* **15**, 3171–3184 (2008).
66. Anderson, K. C. et al. The role of minimal residual disease testing in myeloma treatment selection and drug development: current value and future applications. *Clin. Cancer Res.* **23**, 3980–3993 (2017).
67. da Silva-Diz, V., Lorenzo-Sanz, L., Bernat-Peguera, A., Lopez-Cerda, M. & Muñoz, P. Cancer cell plasticity: impact on tumor progression and therapy response. *Semin. Cancer Biol.* **53**, 48–58 (2018).
68. Hinohara, K. et al. KDM5 histone demethylase activity links cellular transcriptomic heterogeneity to therapeutic resistance. *Cancer Cell* **34**, 939–953.e9 (2018).
69. Kimmel, M. & Axelrod, D. E. *Branching Processes in Biology* (Springer, 2015).
70. Bozic, I. et al. Accumulation of driver and passenger mutations during tumor progression. *Proc. Natl Acad. Sci. USA* **107**, 18545–18550 (2010).
71. Bauer, B., Siebert, R. & Traulsen, A. Cancer initiation with epistatic interactions between driver and passenger mutations. *J. Theor. Biol.* **358**, 52–60 (2014).
72. Yakovlev, A. Y. & Yanev, N. M. Relative frequencies in multitype branching processes. *Ann. Appl. Probab.* **19**, 1–14 (2009).
73. Roney, J. P., Ferlic, J., Michor, F. & McDonald, T. O. ESTIpop: a computational tool to simulate and estimate parameters for continuous-time Markov branching processes. *Bioinformatics* **36**, 4372–4373 (2020).
74. Komarova, N. L. & Wodarz, D. Stochastic modeling of cellular colonies with quiescence: an application to drug resistance in cancer. *Theor. Popul. Biol.* **72**, 523–538 (2007).
75. Chmielecki, J. et al. Optimization of dosing for EGFR-mutant non-small cell lung cancer with evolutionary cancer modeling. *Sci. Transl Med.* **3**, 90ra59 (2011).
76. Yu, H. A. et al. Phase 1 study of twice weekly pulse dose and daily low-dose erlotinib as initial treatment for patients with EGFR-mutant lung cancers. *Ann. Oncol.* **28**, 278–284 (2017).
77. Haeno, H. et al. Computational modeling of pancreatic cancer reveals kinetics of metastasis suggesting optimum treatment strategies. *Cell* **148**, 362–375 (2012).
78. Lindström, H. J. G., de Wijn, A. S. & Friedman, R. Stochastic modelling of tyrosine kinase inhibitor rotation therapy in chronic myeloid leukaemia. *BMC Cancer* **19**, 508 (2019).
79. Danesh, K., Durrett, R., Havrilesky, L. J. & Myers, E. A branching process model of ovarian cancer. *J. Theor. Biol.* **314**, 10–15 (2012).
80. Dean, J., Goldberg, E. & Michor, F. Designing optimal allocations for cancer screening using queuing network models. *PLoS Comput. Biol.* **18**, e1010179 (2022).
81. Chakrabarti, S. & Michor, F. Pharmacokinetics and drug interactions determine optimum combination strategies in computational models of cancer evolution. *Cancer Res.* **77**, 3908–3921 (2017).
82. US National Library of Medicine. *ClinicalTrials.gov* https://clinicaltrials.gov/ct2/show/NCT03810807 (2023).
83. Baar, M. et al. A stochastic model for immunotherapy of cancer. *Sci. Rep.* **6**, 24169 (2016).
84. Yamamoto, K. N. et al. Computational modeling of pancreatic cancer patients receiving FOLFIRINOX and gemcitabine-based therapies identifies optimum intervention strategies. *PLoS ONE* **14**, e0215409 (2019).
85. Yamamoto, K. N., Liu, L. L., Nakamura, A., Haeno, H. & Michor, F. Stochastic evolution of pancreatic cancer metastases during logistic clonal expansion. *JCO Clin. Cancer Inf.* **3**, 1–11 (2019).
86. Moran, P. A. P. *The Statistical Processes Of Evolutionary Theory* (Clarendon Press, 1962).
87. Michor, F., Iwasa, Y. & Nowak, M. A. Dynamics of cancer progression. *Nat. Rev. Cancer* **4**, 197–205 (2004).
88. Beerenwinkel, N. et al. Genetic progression and the waiting time to cancer. *PLoS Comput. Biol.* **3**, e225 (2007).
89. Park, J. & Newton, P. K. Stochastic competitive release and adaptive chemotherapy. Preprint at *bioRxiv* https://doi.org/10.1101/2022.06.17.496594 (2022).
90. Fischer, A., Vázquez-García, I. & Mustonen, V. The value of monitoring to control evolving populations. *Proc. Natl Acad. Sci. USA* **112**, 1007–1012 (2015).
91. Chen, L., Yang, J., Tan, Y., Liu, Z. & Cheke, R. A. Threshold dynamics of a stochastic model of intermittent androgen deprivation therapy for prostate cancer. *Commun. Nonlinear Sci. Numer. Simul.* **100**, 105856 (2021).
92. Camara, B. I., Mokrani, H., Diouf, A., Sané, I. & Diallo, A. S. Stochastic model analysis of cancer oncolytic virus therapy: estimation of the extinction mean times and their probabilities. *Nonlinear Dyn.* **107**, 2819–2846 (2022).
93. US National Library of Medicine. *ClinicalTrials.gov* https://clinicaltrials.gov/ct2/show/NCT03557372 (2021).
94. Tanaka, G., Hirata, Y., Goldenberg, S. L., Bruchovsky, N. & Aihara, K. Mathematical modelling of prostate cancer growth and its application to hormone therapy. *Phil. Trans. R. Soc. A* **368**, 5029–5044 (2010).
95. Albano, G., Giorno, V., Román-Román, P., Román-Román, S. & Torres-Ruiz, F. Estimating and determining the effect of a therapy on tumor dynamics by means of a modified Gompertz diffusion process. *J. Theor. Biol.* **364**, 206–219 (2015).
96. Sfakianakis, N., Madzvamuse, A. & Chaplain, M. A. J. A hybrid multiscale model for cancer invasion of the extracellular matrix. *Multiscale Model. Simul.* **18**, 824–850 (2020).
97. Mirams, G. R. et al. Chaste: an open source C++ library for computational physiology and biology. *PLoS Comput. Biol.* **9**, e1002970 (2013).
98. Chamseddine, I. M. & Rejniak, K. A. Hybrid modeling frameworks of tumor development and treatment. *Wiley Interdisc. Rev. Syst. Biol. Med.* **12**, e1461 (2020).
99. Olsen, M. M. & Siegelmann, H. T. Multiscale agent-based model of tumor angiogenesis. *Proc. Comput. Sci.* **18**, 1016–1025 (2013).
100. Noble, R. et al. Spatial structure governs the mode of tumour evolution. *Nat. Ecol. Evol.* **6**, 207–217 (2022).
101. Beerenwinkel, N., Schwarz, R. F., Gerstung, M. & Markowetz, F. Cancer evolution: mathematical models and computational inference. *Syst. Biol.* **64**, e1–e25 (2015).
102. Swanton, C. Intratumor heterogeneity: evolution through space and time. *Cancer Res.* **72**, 4875–4882 (2012).
103. Wu, H.-J. et al. Spatial intra-tumor heterogeneity is associated with survival of lung adenocarcinoma patients. *Cell Genom.* **2**, 100165 (2022).
104. Zhang, L., Wang, Z., Sagotsky, J. A. & Deisboeck, T. S. Multiscale agent-based cancer modeling. *J. Math. Biol.* **58**, 545–559 (2009).
105. Fusco, D., Gralka, M., Kayser, J., Anderson, A. & Hallatschek, O. Excess of mutational jackpot events in expanding populations revealed by spatial Luria–Delbrück experiments. *Nat. Commun.* **7**, 12760 (2016).
106. Paterson, C., Nowak, M. A. & Waclaw, B. An exactly solvable, spatial model of mutation accumulation in cancer. *Sci. Rep.* **6**, 39511 (2016).
107. Randles, A. et al. Computational modelling of perivascular-niche dynamics for the optimization of treatment schedules for glioblastoma. *Nat. Biomed. Eng.* **5**, 346–359 (2021).
108. Waclaw, B. et al. A spatial model predicts that dispersal and cell turnover limit intratumour heterogeneity. *Nature* **525**, 261–264 (2015).
109. Chkhaidze, K. et al. Spatially constrained tumour growth affects the patterns of clonal selection and neutral drift in cancer genomic data. *PLoS Comput. Biol.* **15**, e1007243 (2019).
110. Liggett, T. M. *Interacting Particle Systems* (Springer, 2005).
111. Nicol, P. B., Barabási, D. L., Coombes, K. R. & Asiaee, A. SITH: an R package for visualizing and analyzing a spatial model of intratumor heterogeneity. *Comput. Syst. Oncol.* **2**, e1033 (2022).
112. Angaroni, F. et al. J-SPACE: a Julia package for the simulation of spatial models of cancer evolution and of sequencing experiments. *BMC Bioinformatics* **23**, 269 (2022).
113. Opasic, L., Scott, J., Traulsen, A. & Fortmann-Grote, C. CancerSim: a cancer simulation package for Python 3. *J. Open Source Softw.* **5**, 2436 (2020).

114. Van Liedekerke, P., Palm, M. M., Jagiella, N. & Drasdo, D. Simulating tissue mechanics with agent-based models: concepts, perspectives and some novel results. *Comp. Part. Mech.* **2**, 401–444 (2015).
115. van Leeuwen, I. M. M. et al. An integrative computational model for intestinal tissue renewal. *Cell Prolif.* **42**, 617–636 (2009).
116. Ghaffarizadeh, A., Heiland, R., Friedman, S. H., Mumenthaler, S. M. & Macklin, P. PhysiCell: an open source physics-based cell simulator for 3-D multicellular systems. *PLoS Comput. Biol.* **14**, e1005991 (2018).
117. Gallaher, J. A., Enriquez-Navas, P. M., Luddy, K. A., Gatenby, R. A. & Anderson, A. R. A. Spatial heterogeneity and evolutionary dynamics modulate time to recurrence in continuous and adaptive cancer therapies. *Cancer Res.* **78**, 2127–2139 (2018).
118. Thomas, D. S., Cisneros, L. H., Anderson, A. R. A. & Maley, C. C. In silico investigations of multi-drug adaptive therapy protocols. *Cancers* **14**, 2699 (2022).
119. Strobl, M. A. R. et al. Spatial structure impacts adaptive therapy by shaping intra-tumoral competition. *Commun. Med.* **2**, 46 (2022).
120. Rejniak, K. A. & Anderson, A. R. A. Hybrid models of tumor growth. *Wiley Interdisc. Rev. Syst. Biol. Med.* **3**, 115–125 (2011).
121. Bacevic, K. et al. Spatial competition constrains resistance to targeted cancer therapy. *Nat. Commun.* **8**, 1995 (2017).
122. Bergman, D. et al. PhysiPKPD: a pharmacokinetics and pharmacodynamics module for PhysiCell. *Gigabyte* **2022**, gigabyte72 (2022).
123. Almendro, V. et al. Inference of tumor evolution during chemotherapy by computational modeling and in situ analysis of genetic and phenotypic cellular diversity. *Cell Rep.* **6**, 514–527 (2014).
124. Sutton, R. S. & Barto, A. G. *Reinforcement Learning: An Introduction* (MIT Press, 2018).
125. Schättler, H. M. & Ledzewicz, U. *Optimal Control for Mathematical Models of Cancer Therapies: An Application of Geometric Methods* (Springer, 2015).
126. Jarrett, A. M. et al. Optimal control theory for personalized therapeutic regimens in oncology: background, history, challenges, and opportunities. *J. Clin. Med.* **9**, 1314 (2020).
127. Swan, G. W. & Vincent, T. L. Optimal control analysis in the chemotherapy of IgG multiple myeloma. *Bull. Math. Biol.* **39**, 317–337 (1977).
128. Kuosmanen, T. et al. Drug-induced resistance evolution necessitates less aggressive treatment. *PLoS Comput. Biol.* **17**, e1009418 (2021).
129. Jerez, S., Pliego, E., Solis, F. J. & Miller, A. K. Antigen receptor therapy in bone metastasis via optimal control for different human life stages. *J. Math. Biol.* **83**, 44 (2021).
130. Hu, X., Ke, G. & Jang, S. R.-J. Modeling pancreatic cancer dynamics with immunotherapy. *Bull. Math. Biol.* **81**, 1885–1915 (2019).
131. de Los Reyes, A. A. & Kim, Y. Optimal regulation of tumour-associated neutrophils in cancer progression. *R. Soc. Open. Sci.* **9**, 210705 (2022).
132. Lee, T., Jenner, A. L., Kim, P. S. & Lee, J. Application of control theory in a delayed-infection and immune-evading oncolytic virotherapy. *Math. Biosci. Eng.* **17**, 2361–2383 (2020).
133. Aspirin, A. P., de Los Reyes V, A. A. & Kim, Y. Polytherapeutic strategies with oncolytic virus-bortezomib and adjuvant NK cells in cancer treatment. *J. R. Soc. Interf.* **18**, 20200669 (2021).
134. Anelone, A. J. N., Villa-Tamayo, M. F. & Rivadeneira, P. S. Oncolytic virus therapy benefits from control theory. *R. Soc. Open. Sci.* **7**, 200473 (2020).
135. Cunningham, J. et al. Optimal control to reach eco-evolutionary stability in metastatic castrate-resistant prostate cancer. *PLoS ONE* **15**, e0243386 (2020).
136. Wu, C. et al. Towards patient-specific optimization of neoadjuvant treatment protocols for breast cancer based on image-guided fluid dynamics. *IEEE Trans. Biomed. Eng.* **69**, 3334–3344 (2022).
137. Angaroni, F. et al. An optimal control framework for the automated design of personalized cancer treatments. *Front. Bioeng. Biotechnol.* **8**, 523 (2020).
138. Lee, J., Lee, D. & Kim, Y. Mathematical model of STAT signalling pathways in cancer development and optimal control approaches. *R. Soc. Open Sci.* **8**, 210594 (2021).
139. Martin, R. B., Fisher, M. E., Minchin, R. F. & Teo, K. L. Optimal control of tumor size used to maximize survival time when cells are resistant to chemotherapy. *Math. Biosci.* **110**, 201–219 (1992).
140. François-Lavet, V., Henderson, P., Islam, R., Bellemare, M. G. & Pineau, J. An introduction to deep reinforcement learning. *Found. Trends Mach. Learn.* **11**, 219–354 (2018).
141. Engelhardt, D. Dynamic control of stochastic evolution: a deep reinforcement learning approach to adaptively targeting emergent drug resistance. *J. Mach. Learn. Res.* **21**, 1–30 (2020).
142. Eastman, B., Przedborski, M. & Kohandel, M. Reinforcement learning derived chemotherapeutic schedules for robust patient-specific therapy. *Sci. Rep.* **11**, 17882 (2021).
143. Ebrahimi Zade, A., Shahabi Haghighi, S. & Soltani, M. Deep neural networks for neuro-oncology: towards patient individualized design of chemo-radiation therapy for glioblastoma patients. *J. Biomed. Inf.* **127**, 104006 (2022).
144. Moreau, G., François-Lavet, V., Desbordes, P. & Macq, B. Reinforcement learning for radiotherapy dose fractioning automation. *Biomedicines* **9**, 214 (2021).
145. Tortora, M. et al. Deep reinforcement learning for fractionated radiotherapy in non-small cell lung carcinoma. *Artif. Intell. Med.* **119**, 102137 (2021).
146. Yauney, G. & Shah, P. in *Proc. 3rd Machine Learning Healthcare Conf.* Vol. 85 161–226 (PMLR, 2018).
147. Padmanabhan, R., Meskin, N. & Haddad, W. M. Reinforcement learning-based control of drug dosing for cancer chemotherapy treatment. *Math. Biosci.* **293**, 11–20 (2017).
148. Wang, M., Scott, J. G. & Vladimirsky, A. Stochastic optimal control to guide adaptive cancer therapy. Preprint at *bioRxiv* https://doi.org/10.1101/2022.06.17.496649 (2022).
149. Tseng, H.-H. et al. Deep reinforcement learning for automated radiation adaptation in lung cancer. *Med. Phys.* **44**, 6690–6705 (2017).
150. Tardini, E. et al. Optimal treatment selection in sequential systemic and locoregional therapy of oropharyngeal squamous carcinomas: deep Q-learning with a patient–physician digital twin dyad. *J. Med. Internet Res.* **24**, e29455 (2022).
151. Niraula, D., Jamaluddin, J., Matuszak, M. M., Haken, R. K. T. & Naqa, I. E. Quantum deep reinforcement learning for clinical decision support in oncology: application to adaptive radiotherapy. *Sci. Rep.* **11**, 23545 (2021).
152. US National Library of Medicine. *ClinicalTrials.gov* https://clinicaltrials.gov/ct2/show/NCT02415621 (2023).
153. US National Library of Medicine. *ClinicalTrials.gov* https://clinicaltrials.gov/ct2/show/NCT03543969 (2023).
154. US National Library of Medicine. Pilot study of adaptive BRAF-MEK inhibitor therapy for advanced BRAF mutant melanoma. *ClinicalTrials.gov* https://clinicaltrials.gov/ct2/show/NCT03543969 (2018).
155. US National Library of Medicine. *ClinicalTrials.gov* https://clinicaltrials.gov/ct2/show/NCT03630120 (2021).
156. Smalley, I. et al. Leveraging transcriptional dynamics to improve BRAF inhibitor responses in melanoma. *EBioMedicine* **48**, 178–190 (2019).
157. Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C. & Faisal, A. A. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. *Nat. Med.* **24**, 1716–1720 (2018).
158. Dean, J. A. et al. Phase I study of a novel glioblastoma radiation therapy schedule exploiting cell-state plasticity. *Neuro Oncol.* **25**, 1100–1112 (2022).
159. Zhang, J. et al. A phase 1b adaptive androgen deprivation therapy trial in metastatic castration sensitive prostate cancer. *J. Clin. Oncol.* **40**, 5075 (2022).

Acknowledgements
The authors gratefully acknowledge support from the Dana-Farber Cancer Institute's Center for Cancer Evolution.

Author contributions
All authors contributed equally to the preparation of this manuscript.

Competing interests
F.M. is a co-founder of and has equity in Harbinger Health, has equity in Zephyr AI, serves as a consultant for Harbinger Health and Zephyr AI and is on the board of directors of Exscientia Plc. F.M. declares that none of these relationships are directly or indirectly related to the content of this manuscript. All other authors declare no competing interests.

Additional information
Peer review information *Nature Reviews Bioengineering* thanks Robert Noble, Blair Colyer and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.