Adaptive vs continuous cancer therapy: Exploiting space and trade-offs in drug scheduling

Jill A. Gallaher[1], Pedro M. Enriquez-Navas[2], Kimberly A. Luddy[2], Robert A. Gatenby[1], and Alexander R. A. Anderson[1]

[1]H. Lee Moffitt Cancer Center, Integrated Mathematical Oncology (Tampa, FL, United States)

[2]H. Lee Moffitt Cancer Center, Cancer Imaging and Metabolism (Tampa, FL, United States)

Abstract

The treatment of advanced cancers has greatly benefited from the introduction of new agents, such as targeted therapy and checkpoint inhibitors, to supplement or bypass conventional therapies. However, even the most effective therapy usually fails over time as cancer cells are able to deploy a wide range of molecular and microenvironmental resistance strategies. Here we propose that, while molecular dynamics largely govern response and resistance to therapy, evolutionary dynamics determine survival and proliferation of treatment-resistant cells. We hypothesize that understanding these evolutionary interactions may identify strategies to delay or prevent proliferation of the resistant population using conventional therapies thus prolonging time to recurrence.

Here we use an off-lattice, agent-based framework to model competition among sensitive and resistant populations during therapy in a spatially competitive resource-limited tumor microenvironment. Our model applies a classic evolutionary trade-off between fecundity (cellular proliferation) and survivorship (drug sensitivity). We simulate the application of an anti-proliferative drug on varying ratios of mixed sensitive and resistant cells using two general treatment strategies: a continuous schedule of maximum tolerated dose or an evolution-informed schedule that incorporates dose modulation and treatment vacations to sustain control of the tumor through competition between sensitive and resistant cell populations. We find tumors consisting only of sensitive cells can be cured with continuous treatment, but the presence of any significant population of resistant cells will lead to eventual recurrence. We identify two treatment strategies that control heterogeneous tumors: one emphasizes continuous dose modulation, and the other relies on treatment vacations. Both strategies control tumors over a wide range of resistant/sensitive population ratios but the average dose given is significantly lower with dose modulation while a more vacation-oriented schedule can control more aggressive tumors.

1. Introduction

Despite major advances in cancer therapies, most metastatic cancers remain fatal because tumor cells have a remarkable capacity to evolve drug resistance, both through genetic and non-genetic mechanisms (1). Most investigations of cancer treatment resistance have focused on identifying and targeting the molecular mechanisms that confer resistance. However, defeat of one resistance strategy often results in the deployment of another (2).

An alternative approach focuses on the population-level dynamics governed by Darwinian evolutionary principles that define the fitness of each cell within the local environmental context. For example, cancer cells often employ multidrug resistance pumps, in which the synthesis, maintenance, and operation require considerable investment of resources (up to 50% of the cell's total energy budget) (3). In the harsh tumor microenvironment this investment in survival will likely require diversion of resources that would ordinarily be devoted to invasion or proliferation. Thus, while tumor cells may possess the molecular mechanisms necessary for therapy resistance,

proliferation of resistant cells is governed by complex interactions that include the cost/benefit ratio of the resistance mechanism(s) and competition with other tumor subpopulations.

A common maxim in cancer treatment is to "hit hard and fast" through maximum dose dense strategies that administer the highest possible drug dose in the shortest possible time period. The maximum tolerated dose (MTD) principle has been the standard of care for cancer treatment for several decades and is the basis for clinical evaluation for most Phase I cancer drug trials. It has not, however, resulted in consistent cures in patients with most disseminated cancers (4). An evolutionary flaw in this strategy is the assumption that resistant populations are not present prior to therapy. It is now clear that cancer cells can be insensitive even to treatments that they have never "seen" before. Therefore, MTD therapy designed to kill as many cancer cells as possible, although intuitively appealing, may be evolutionarily unwise. This is because of a well-recognized Darwinian dynamic from ecology termed "competitive release". Observed, for example, when high doses of pesticide are applied for pest eradication (5), competitive release allows rapid emergence of resistant populations because of the combination of intense selection pressure and elimination of all potential competitors.

Despite the growing recognition of the importance of heterogeneity and evolution in driving treatment failure, explicit inclusion of Darwinian principles in clinical trial design is rare (6-9). However, both clinical and pre-clinical studies have shown promising results. Enriquez-Navas *et al.* used an evolution-guided treatment strategy to control breast cancer tumors in mice (10). They found that progression free survival can be prolonged when paclitaxel treatment schedules incorporate dose modulations and treatment holidays such that less drug is given to a responding tumor and more to a rebounding tumor. Two individual mice with controlled disease are shown in **Fig. 1A**. An ongoing clinical trial at Moffitt Cancer Center (NCT02415621) tests these evolutionary principles in patients with metastatic castration resistant prostate cancer to try to prevent the evolution of resistance to abiraterone therapy (11). In this trial, abiraterone is discontinued when the blood Prostate Specific Antigen (PSA) concentration falls below 50% of the initial value and does not resume until the PSA returns to the pre-treatment level. PSA changes during therapy are shown for 2 patients in **Fig. 1B**. It is important to note that: 1) each patient serves as their own control to calibrate the PSA as a relative value, and 2) the adaptive schedules effectively personalize the treatment to patient response so that while one patient has only 2 courses of treatment in a year, another gets 3 in 10 months. Yet both remain under control.

Figure 1. *Adaptive therapy in the laboratory (A) and the clinic (B). The original value of volume (A) or PSA (B) is shown across time as a red dashed line. The value at which treatment vacation occurs is shown as a shaded gray region. The dose over time is shown in purple. A) Two mice implanted with human breast cancer were treated with paclitaxel such that a 20% decrease in tumor volume results in a 50% dose decrease, where as a 20% increase in tumor volume results in a 50% dose increase and otherwise the dose remains the same. If the volume dips below 150 mm³, a treatment vacation occurs. These data were previously published in Enriquez-Navas et al (10). B) Two prostate cancer patients treated with abiraterone such that treatment is stopped if PSA falls below 50% of the original, and resumes when the PSA exceeds the original value.*

Two important questions emerge from these results: (i) For which cancers is continuous (MTD) treatment the best strategy and when is adaptive better? (ii) When should a treatment holiday driven adaptive therapy be given instead of a dose modulation one? While *selection* for resistance through application of continuous cytotoxic therapy seems inevitable, *proliferation* of those cells may be controlled using evolutionary principles. Importantly, multiple experimental models have shown that drug-resistant cancer cells proliferate slower than sensitive cells in the absence of drug (12-16). This is because resistant tumor cells, like most drug-resistant bacteria, incur a fitness cost due to the energy costs involved (17). Here, we computationally investigate these questions under the hypothesis that these costs can be exploited to delay or prevent proliferation of resistant cells in the tumor.

2. Methods

To investigate intratumoral evolutionary dynamics during treatment, we modify an off-lattice agent-based model built to investigate trait selection via spatial competition. We focus on the role of initial distributions of proliferation rates on both overall tumor growth and spatial distribution of subpopulations. We then examine outcomes from varying treatment strategies with an anti-proliferative drug focusing on progression of resistant populations and treatment failure.

Figure 2. *Model rules. A) Decision-making flowchart for each cell over each time step. Proliferation steps are shown in green, quiescence in yellow, and drug-induced death in red. The decisions that are multicolored reflect the phenotypic heterogeneity in cell cycle times and sensitivity to drug. Key shows the trade-off between cell cycle time and drug sensitivity. B) The off-lattice model assumes space limited proliferation, such that a dividing cell with insufficient space for non-overlapping offspring enters a quiescent state.*

The flowchart in **Figure 2** shows the decisions made for each cell at each time point. Green arrows and nodes represent proliferation decisions; yellow, quiescence; multicolored, decisions affected by heterogeneity; and red, exit of a cell due to drug-induced death. In this off-lattice model, we assume that cells take up space such that at carrying capacity, they enter quiescence due to contact inhibition. Space is deemed available if any contiguous set of integer angles are empty and sufficient to allow a cell to fit without overlap (described previously in (18)). When a cell is quiescent, we assume that the cell discontinues decision-making until space is available, i.e. it cannot be affected by drug and it cannot proliferate. We also assume that there is no cell death due to regular turnover and that resources are abundant when cells are not suppressed by lack of space.

Based on a number of experimental observations, we assume that the cost of resistance results in a decrease in the cell's proliferation rate (12-16). Thus, we assign the fastest proliferation rate to cells that are 100% sensitive, and in a linear fashion, completely resistant cells have the slowest proliferation rate. This provides a simple linear trade-off between a cell's fitness in the absence of drug and its fitness when the drug is present.

Regardless of the treatment strategy, a cell's response to drug exposure depends on its sensitivity and the dose applied. We define drug sensitivity, due to the imposed trade-off with the ability to reproduce, as a decreasing linear function of cell cycle time:

$$s(T) = \frac{T_{max} - T}{T_{max} - T_{min}} , \tag{1}$$

where T is the time it takes to go through a cell cycle, and T_{min} and T_{max} are the minimum and maximum allowed cell cycle times, which are 10 and 50 hours, respectively. We represent drug efficacy as a typical Hill function for simplicity and generality (19,20):

$$E(D) = \frac{D^n}{D^n + K_{0.5}^n} , \tag{2}$$

where $n=1.5$ and $K_{0.5}=0.125$. These values are not calibrated to any specific half maximal activity, but generically assume a response function that gives nearly 100% probability of death at the highest drug concentrations, dropping off slowly at mid ranges and then quicker as it gets to lower concentrations. The probability of death, P_{death}, combines the sensitivity and dose effects:

$$P_{death}(T, D) = s(T)E(D) = \frac{s(T)D^n}{D^n + K_{0.5}^n} . \tag{3}$$

We explore treatment strategies other than conventional MTD, which we term "adaptive therapy", wherein we adapt the next treatment based upon the tumors previous response. More specifically, we vary the dose D from its previous value D_0 based on the relationship of the current number of cells N to the previous number of cells N_0 according to the following equation:

$$D(\alpha, \beta) = \begin{cases} 0 & \text{if } N < 0.5N_0 \\ (1+\alpha)D_0 & \text{if } N > (1+\beta)N_0 \\ (1-\alpha)D_0 & \text{if } N \leq (1-\beta)N_0 \\ D_0 & \text{otherwise} \end{cases}, \qquad (4)$$

where α and β are fixed values that can affect responsiveness of the change in dose to the change in population size. In **Eq. 4**, if the number of cells is below half of the original, a treatment vacation is enforced. If the change in number of cells is within a specified window the dose remains the same as before. Otherwise, there will be a fractional increase in the dose by α if the fractional increase in population size is greater than β and a fractional decrease in dose by α if the fractional decrease in the population size is greater than β. The dose begins at the MTD, and never exceeds the MTD and each new dose is determined every 3 days in accordance with the experiments reported by Enriquez-Navas *et al* (10). If a cell is targeted by the drug, there is a 15-30h delay, randomly chosen, before it is removed from the domain, to account for the time it takes for apoptosis to occur (21,22).

3. Results

Fitness differences and space limitations affect competition and selection

The cost of resistance and the selection force imposed by space limitations can be demonstrated through a combination of experimental and mathematical models. The MCF7Dox cell line is highly resistant to many chemotherapy agents due to upregulation of the membrane efflux pump P-glycoprotein or Multiple Drug Resistance (MDR1) proteins. These cells were labeled with red fluorescence protein (RFP). The MCF7Dox cells are plated with the parental MCF7 cell line (labeled with green fluorescent protein, GFP), which does not express MDR1 and remains sensitive to chemotherapeutic agents. In these experiments, the cell lines were mixed in a 1:1 ratio to observe the relative dynamics of growth. The cells are grown for 3-4 days, harvested, counted, split and re-plated over 46 days. GFP and RFP were measured by flow cytometry at each interval (**Fig. 3A**). The resistant cells, because of the fitness cost of the MDR expression, are rapidly outcompeted by the parental MCF7 cells after only a few generations in the absence of chemotherapy (doxorubicin) in the media.

Figure 3. *Fitness differences and space limitation affect growth and selection. A) Sensitive (MCF7) and resistant (MCF7Dox) cells to doxorubicin are grown together in vitro and demonstrate the fitness cost of resistance. The plot shows over time the percent of drug-sensitive and drug-resistant cells. B) Model population dynamics showing fraction of sensitive phenotype (40h cell cycle) and resistant (60h cell cycle) cells at a ratio of 1:1 with 100 cells and grown up to*

20,000 cells. In one example, the cells were initially spread sparsely (dish), and in the other, cells are initially seeded in a tight cluster configuration. C) The spatial distribution for each example is shown at two time points, t_1=7 days and t_2=30 days, and for the cluster also at t_3=90 days. A zoomed-in view shows the distribution of proliferating (yellow) and quiescent (green if sensitive, red if resistant) cell types.

We then investigated these evolutionary dynamics through our computational model system (**Fig. 3B**) to examine competition for space with given fitness differences over time. We assume that the population mix consists of a faster growing sensitive phenotype and a slower growing resistant phenotype. We use the average growth rates found at the beginning of the experiment, which was a 40 hour cycle time for MCF7 and a 60 hour cycle time for MCF7Dox. We initially seeded 100 cells scattered throughout a circular domain with a 1.5 mm radius to grow distinct colonies, reflecting the distribution in the *in vitro* model. We grow the population to 20,000 cells and find that, over time, the sensitive cells quickly take over the space and prevent further proliferation of the resistant colonies. However, the *in vitro* environment with ample space and nutrients is distinct from the densely populated and resource depleted tumor environment. In solid tumors, the cells are less free to move and proliferate, which adds additional selection pressure. If we initialize the simulation with tumor cells in a more confined solid mass but with the same fractions of sensitive and resistant, we find that there is a significant change in the population dynamics; the limited space causes greater competition between the different populations, taking much longer to grow to the same population size. The spatial layout (**Fig. 3C**) shows that when the cells can grow as separate colonies, the "tumor edge" is large and convoluted. With the cells arranged in the dense cluster, there is greater selection at the boundary, and the sensitive cells quickly take over the expanding front trapping the sensitive cells within. Because proliferation is limited to the edge of the mass, the population grows much slower.

Phenotypic spread reduces coexistence during growth

Sensitivity to drugs is often viewed in terms of binary states, where a cell is simply sensitive or resistant and it would respond with either death or survival. However, in reality tumors have a more nuanced mix of phenotypes that differ in their sensitivity to a drug. Here we investigate how various mixtures of sensitive and resistant cells compete to form a solid tumor mass.

We begin with 100 cells, with phenotypes drawn from a normal distribution of drug sensitivities (each tumor has a unique mean cell cycle time τ and a standard deviation σ_τ). The simulation is stopped when the number of cells reaches 15,000, which produces a mass around 1.5 mm in diameter, representing a micrometastasis. The initial and final population distributions are shown in **Figure 4** for each case, with the percent makeup of the dominant phenotype labeled. We consider an array of initial distributions with increasing mean and standard deviation in cell cycle times (corresponding to decreasing sensitivities). The difference between the initial and final compositions can tell us something about the competition between the phenotypes. Toward the top of the array, where there is little heterogeneity, the compositions remain rather homogeneous, i.e. the sensitive tumors remain mostly sensitive and the resistant tumors retain resistance. However, with increasing initial heterogeneity, the more sensitive cells are seen to outcompete the more resistant cells. In many cases, the competition is so great that the final proportion of resistant cells is imperceptibly small. **Figure 4**A-D shows the final spatial configuration of some of the more extreme cases. They all show that the more resistant cells get trapped in the interior of the tumor by the more sensitive cells, which take over the invading edge. More coexistence between phenotypes is seen in B and D, but given enough time the more proliferative phenotypes at the invading edges will eventually take over. Overall, the trend shows

that for a heterogeneous distribution of phenotypes, the faster proliferators will take over the population, but the timing will depend on the relative fitness differences.

Figure 4. *Growing tumors from many different initial compositions of cell phenotypes. We display in pie charts the fraction of the population occupied by cells with different proliferation rates at the beginning (left) and end (right) of the simulation. In each, the initial distribution is defined by a mean cycle time τ and a standard deviation σ_τ. A pie piece with a dashed border and a number shows the dominant clone and the percent occupied at that time point. The color key gives the cycle time and sensitivity over the range of colors. Final spatial distributions for selected cases A-D are displayed below.*

Continuous treatments can cure some tumors; adaptive treatments can control most tumors

Treatment is applied to the previously grown tumors (number of original cells N_0=15,000) and stops when either the tumor is cured (i.e. the number of current cells N=0), recurs (i.e. $N=4/3N_0$=20,000 cells), or the tumor reaches an age of 2 years post-treatment. We assume tumor control if the final number of cells is below 10,000, which is half the value that determines recurrence and accounts for fluctuations during adaptive treatments. For simplicity, we do not simulate any short-term pharmacokinetics, so for the continuous therapy, we simply apply the maximum tolerated dose (MTD) the entire duration and each dose change in the adaptive strategy is instantaneously realized. Only when a cell is capable of dividing does it become sensitive to drug toxicity and die with the probability given in **Eq. 3**. Further, we define several metrics to describe adaptive treatments: the average dose, D_{ave}, gives the average dose over the total time period, and the equivalent dose, D_{eq}, gives the equivalent number of MTD doses that equal the sum of actual dose given over the total time period.

Figure 5. *CT and AT schedules applied to different compositions of tumors: initial distributions before growth with mean sensitivity s and standard deviation σ_s are A) s=100%, σ_s =5%; B) s=100%, σ_s =25%; C) s=60%, σ_s =5%; and D) s=60%, σ_s =25%. The top panels show the dose schedules (in % of MTD over time) for each treatment strategy. Below those, the population dynamics for each strategy is shown. The spatial configurations for the two strategies at several different time points are shown below the plots, starting from the same tumor configurations. A link to the movies is given in the Supplement.*

We start by comparing CT with an adaptive scheme that will change the dose by 25% (α =0.25) if the population size changes by 5% ($\beta=0.05$) and otherwise follows Eq. 4. In general, there is not one therapy strategy that works best for all tumors, but the response to the treatment depends on the tumor composition. In **Fig. 5** we show the range of response with both treatment strategies: continuous therapy (CT) and adaptive therapy (AT), using the tumors grown in **Fig. 4** (labeled A-D). We see that homogeneous tumors consisting of only sensitive cells (**Fig. 4A**) are curable with CT and can be controlled with AT. While cure by CT takes less time (114 days), AT gives a lower average dose (D_{ave}=34%), but a larger total dose (81 doses vs. 39 doses) over a longer time period (2 years). Tumor B consists primarily of sensitive cells but is initiated with a broader range of phenotypes, including some treatment-resistant phenotypes that have ended up in the tumor interior surrounded and suppressed by sensitive cells. When treated continuously, the tumor initially responds very well, but this is followed by recurrence of a completely resistant population after around 1 year. However, the same tumor under the adaptive strategy can be controlled with a lower average dose (D_{ave}=35%) and lower total dose (85 doses vs. 122 doses) over the full 2 years. Tumor C is mostly a homogeneous tumor with a mean survival rate of around 60% to the MTD, which should yield a net negative growth rate under a full drug dose. However, there is only a slight response to treatment at the maximum dose as the dominant sensitive phenotype is slowly eliminated and the moderately resistant cells survive under CT, which leads to

tumor progression. Using the AT strategy, the change in the population size during treatment is so slow that the dose never modulates. When the population halves ($N=0.5N_0$) at around 1 year, most of the sensitive cells, which were coexisting with the less sensitive cells, have already been eliminated. At this point, the treatment vacation begins, and because the tumor is diffuse, there is a large portion of non-quiescent cells, so we see a swift increase in growth as resistant cells (only slightly <50% sensitivity) take over. Overall, this causes a net loss of 1 month's time over the CT strategy. There is no real advantage of one treatment strategy over the other as both essentially give the same total dose (230 doses) over the full 2-year time period. In tumor D, when there is more phenotypic spread, the sensitive cells again outcompete the more resistant cells during the growth phase. We see here that both treatments lead to eventual recurrence, but AT keeps the total number of cells controlled longer than CT. Notably, AT gives a higher total dose (192 doses vs. 101 doses), but it recurs later (~2 years vs. 300 days) and at average of 79% of the MTD.

From the population dynamics of the whole array of tumor compositions, we can discern why different treatments fail or succeed. **Figure 6** shows how the sensitive and resistant populations change with CT and AT (**Fig. 6A**) for the same complete array of tumors in **Fig. 4**. The arrays show, for a single run, how the makeup of each tumor changes over the course of treatment. The asterisk marks, for CT, which tumors are cured and, for AT, which tumors are controlled. We see that for both cure and control the sensitive population must be maintained over the treatment period. In some cases, the resistant population starts to appear at the end, but the phenotypes representing most of the population are sensitive for the entire simulation.

Figure 6. *Changes of phenotypic compositions over time for each pre-growth normal distribution with mean cell cycle times τ from 10-30h (from left to right), and heterogeneity (standard deviation in cell cycle time σ_τ of 2-10h from top to bottom). Each square shows the distribution of phenotypes over time as depicted by the key shown below the figure. A) For CT, cure is labeled with an asterisk. For AT ($\alpha=0.25$, $\beta=0.05$), control is labeled with an asterisk. B) Best 2 out of 3 winning strategies for different tumor compositions. Each square is colored according to the best-case scenario based on 3 simulations according to the key. Cure by CT is favored first, control is favored next. If recurrence is most likely, the treatment with the most time gained before recurrence is favored, and the color indicates the average time gained by the winning strategy (time gained =|TTR_{CT}-TTR_{AT}|).*

The descending order of desired outcomes is to: cure, maintain, then gain the most time before recurrence. **Figure 6B** delineates these winning strategy choices for the best 2 out of 3 trials over the array. We find that the top left side of the plot gives cure by CT as the best and most likely outcome, where the initial distribution is mostly sensitive and homogeneous. On the lower left and some along the edge of the cure region, control by AT is most likely over cure and recurrence, which is a region where the initial distribution is still mostly sensitive but has some heterogeneity. In any other case we have recurrence, and the color bar on the bottom indicates

the time gained by the winning strategy. We see that for the top left region, treatments produce nearly the same time to recurrence. That is, the response with such low sensitivities is so poor that dose modulation is rare, such that the AT treatment is nearly the same as the CT treatment (constant MTD). The top right edge of the cure region favors CT by a few months, but the middle and lower right regions of the array favor the AT strategy, because enough sensitive cells will have trapped the resistant cells during growth and are maintained during treatment. The decoupled time to recurrence plots are shown in **Fig. S1A (CT and AT1)**.

A more vacation-oriented strategy can control the tumor at the expense of higher doses

The AT schedule used so far, which we now refer to as AT1 (α=0.25, β=0.05), was chosen somewhat arbitrarily. Changing α and β changes the amount the dose is modulated if the population changes by a threshold amount, respectively, and these values can have an impact on AT control. We test another schedule that boosts both values, requiring a larger change in population size (10%) to produce a larger change in the dose modulation (50%), which we will call AT2 (α=0.50, β=0.10).

Figure 7. Comparing CT, AT1 (α=0.25, β=0.05), and AT2 (α=0.50, β=0.10) treatments, where adaptive dose schedules are given by Eq. 4. A-C) Using a tumor from a pre-growth normal distribution of s=100 % and σ_s=25 %, we compare these treatment strategies. A) The dose schedules (in % of MTD) for each treatment. B) The population dynamics during each treatment. C) The spatial layout at various time points. D) Best 2 out of 3 winning strategies for different tumor compositions. Each square is colored according to the best-case scenario based on 3 simulations according to the key at the bottom. Cure by CT is favored first, and control is favored next with a preference toward a lower dose. If recurrence is most likely, the strategy with the most time gained before recurrence is favored, and the square is colored according to the average time gained by the winning strategy (time gained =$|TTR_{CT}$-$TTR_{AT}|$). E) The average dose D_{ave} over the treatment period for the array of initial distributions for CT, AT1, and AT2. A link to the movies is given in the Supplement.

We compare all treatment strategies (CT, AT1, and AT2) in Fig. 7 using a tumor with a pre-growth distribution that centers around a sensitive phenotype *(s=100%)* but with a large degree of heterogeneity (σ_s=25%). At the start of treatment the resistant cells are trapped in the interior by the sensitive cells. **Figure 7A** gives the dose schedules, **Fig. 7B** shows the population dynamics, and **Fig. 7C** reveals the spatial layouts at various time points. The tumor recurs with the CT strategy while both AT strategies control the tumor for the full 2 years. We can see in the AT1 dose schedule that the dose lowers at the start then incorporates both vacations and dose adjustments to control the tumor. In the AT2 schedule, however, we find that the population does not change sufficiently fast to invoke a dose change, so control is achieved solely by having treatment vacations. However, the AT1 schedule gives a much lower effective dose to the patient. Whilst the CT treatment recurs at 366 days after 122 doses of MTD, the adaptive therapies allow for tumor control over 2 years with an average of D_{ave}=35%(76%) of the MTD resulting in D_{eq}=85(185) doses for AT1(AT2). With this new AT2 schedule, we find similar winning outcomes when compared to CT as with the previous schedule, however, more importantly we see extended control for tumor compositions that are less sensitive and heterogeneous (**Fig. 7D**). The average doses for the whole array of tumor compositions are given for CT, AT1, and AT2 in **Fig. 7E**. We find that while the treatment vacations do reduce the dose, modulating the dose substantially lowers the average dose in regions that are controlled by adaptive schedules. When the proliferating cells exposed to the drug are very sensitive, cell kill can be achieved with a smaller dose. However, when less sensitive phenotypes are proliferating, the increased dose in AT2 grants better control.

More treatment vacations are needed to control heterogeneous invasive tumors

For the adaptive schedules to work, the sensitive cells must impede the proliferation of the resistant cells by competing for space and trapping the resistant cells inside the tumor. However, if the cells can move, the spatial structure that keeps the cells quiescent and hidden from the drug is disturbed. We now examine the effect of cell migration by allowing cells to move in persistent random walk. Each cell is given a persistence time (it will move for a length of time randomly chosen from a normal distribution with 80±40 minutes), after which it turns at a random angle and starts again. All cells move at a modest speed of 5 μm/h for as long as the cell is not in the quiescent state and doesn't contact another cell, in which case the cells simply turn by a random angle and start again with a new persistence time.

Figure 8. *Allowing migration (5 μm/h) reduces the efficacy of control by AT. A-C) An example of a tumor (with a pre-growth distribution of s=100 % and σ_s=25%) subject to strategies CT, AT1 (α=0.25, β=0.05), and AT2 (α=0.50, β=0.10). We show the dose schedules (A), the population over time (B), and the spatial layout (C) for each strategy. D) Best 2 out of 3 winning strategies for different tumor compositions. Each square is colored according to the best-case scenario based on 3 simulations according to the key at the bottom. Cure by CT is favored first, and control is favored next with a preference toward a lower dose. If recurrence is most likely, the strategy with the most time gained before recurrence is favored, and the square is colored according to the average time gained by the winning strategy (time gained =|TTR$_{CT}$-TTR$_{AT}$|). E) The average dose D_{ave} over the treatment period for the array of initial distributions for CT, AT1, and AT2. A link to the movies is given in the Supplement.*

Figure 8 shows an example with the same initial conditions as the previous example (s=100% and σ_s=25%), but with migrating cells. Again, the dose schedules are given (**Fig. 8A**), the population over time (**Fig. 8B**), and the spatial layout (**Fig. 8C**). We see that for the CT case, compared to the case without migration (**Fig. 7**), we have a much quicker decline in the population followed by a much faster relapse at 153 days. For this strategy, the patient gets 51 doses at MTD. Due to the spatial spreading from migration, very few cells are quiescent and the drug is able to affect more cells. This also means that resistant cells escape sooner leading to a faster return of growth. For AT1, the tumor population doesn't get small enough to trigger a vacation, but dose modulation occurs. The dose reduces then rises back up at the end when the resistant, unresponsive cells take over. This schedule ends up with a recurrent tumor at 258 days after D_{eq}=101 doses, delivered at D_{ave}=30% MTD. Using AT2, the tumor is controlled for the 2 year time period (not fully shown). The dose also modulates because the larger proliferating fraction (and therefore susceptible to the drug) causes larger fluctuations in the population size triggering a dose change. Importantly, by employing the quicker switching between fast growth during the

vacations and fast death with higher doses, the tumor is able to keep the sensitive population at the invasive edge. This schedule gives D_{eq}=163 doses at D_{ave}=67% MTD.

The full sweep of the array space is shown in **Fig. 8D** comparing CT vs. AT1 vs. A2. We find that while most of the original region cured by CT remains, allowing migration has destroyed the control by AT1 while the regions previously controlled with AT2 are somewhat preserved. The average doses for each strategy are shown in **Fig. 8E**. In the sensitive region, which was previously controlled by both adaptive strategies in tumors without migration, the doses are seen to be lower for both AT1 and AT2. However, control is only achieved with the AT2 strategy. As we increase the migration rate further to 10 μm/h, neither adaptive treatment is favored to control the tumor (**Fig. S1**).

Treatment vacations help delay recurrence of heterogeneous tumors with phenotypic drift

We have considered treatment responses to tumors that have pre-existing heterogeneity but assume that progeny directly inherit the same phenotypes as their initial parental lineage. However, this ignores the potential impact of subsequent mutations or epigenetic changes that may alter the cell phenotype over generations. To test how phenotypic drift affects the response to treatment strategy, we allow a cell the opportunity to change its proliferation rate slightly from its mother cell upon division. In particular, there is a 10% probability at each division that a cell will alter its cycle time by ±1 h or stay the same (ensuring that the IMT stays within the allowed range of 10-50h).

Figure 9. *Allowing phenotypic drift reduces the efficacy of AT. A-C) An example of a tumor (with a pre-growth*

In **Fig. 9**, we show an example with initial conditions of s=80% and σ_s=25%. The treatment schedules for CT, AT1, and AT2 are shown in **Fig. 9A** as a percent of the MTD. The population dynamics are shown in **Fig. 9B**. Finally, the spatial configurations for each treatment at various time points are displayed in **Fig. 9C**. In all cases, treatments eventually fail. CT, as in the previous case with migration, automatically selects for the resistant cells and delivers the MTD for 297 days, which corresponds to 99 doses of MTD. AT1 lasts nearly the same amount of time as the CT strategy, but by using a smaller dose (D_{ave}=66% over 285 days with D_{eq}=63 doses). Because of the drifting phenotype, the sensitive cells can now become more resistant over generations. We can see that for AT1 there are individual hotspots on the tumor's edge where selection for more resistant phenotypes emerges. This results in separation of the mass into several individual tumor clumps centered around these regions. AT2 is able to control the tumor a little longer (387 days) by having treatment vacations and no dose modulation (D_{ave}=85% and D_{eq}=109 doses).

The array of different tumor compositions is evaluated to compare CT vs. AT1 vs. AT2 (**Fig. 9D**). The region originally cured by CT is significantly reduced, and in both cases the ability to control the tumor through adaptive means is completely eliminated. CT actually does a better job with just a few months gain over AT1 in most regions. But again, the AT2 schedule delays recurrence for longer over the CT schedule for heterogeneous tumor compositions. Regardless, neither treatment sustains the control for long, as the net gain with AT2 is at the most 6 months. The average doses for CT, AT1, and AT2 are shown in **Fig. 9E**. In the sensitive region, which was controlled when there was no drift, the doses are seen to be higher for both AT1 and AT2. Increasing the drift rate further to occur at every division destroys any control whatsoever using CT, AT1, or AT2 (**Fig. S2**).

Why do treatment vacations maintain tumor sensitivity for longer?

When migration and phenotypic drift were not allowed, dose modulation granted a way to control a tumor while lowering the dose rate. However, an emphasis on treatment vacations can partly preserve this control. To understand why, it is helpful to introduce the effective net proliferation rate: $\rho_{eff}=\rho(1-P_{death}/0.5)$, where P_{death} is the probability of death, and ρ is the proliferation rate unperturbed by treatment. This value gives a better understanding of how dose and sensitivity combine to give a net proliferation rate across the range of phenotypes. It essentially modulates the proliferation rate by the probability of death due to a certain dose (% of MTD), such that if P_{death}=0, $\rho_{eff}=\rho$, if P_{death}=0.5, ρ_{eff}=0 and at P_{death}=1.0 the cells die off as fast as they divide. The effective net proliferation is shown in the background heat map of **Fig. 10**. The highest sensitivities have both the highest growth rates (without drug) and the highest death rates (at MTD), which are the darkest colors at the top row of the heat map. Remaining in the red zone means maintaining net cell kill over proliferation.

Figure 10. *Effect of modulating dose on the growth rates of different phenotypes (sensitivities) for different treatment strategies: CT, AT1 (α=0.25, β=0.05), and AT2 (α=0.50, β=0.10). We use the trajectories for the invasive (left) and drifting (right) phenotypes from the examples in Fig. 8A-C and Fig. 9A-C, respectively. The background heat map displays the effective growth rate of a phenotype, which modulates the given growth rate by the probability of death due to a specific dose. The gray dashed line indicates where net growth is zero. Each colored line shows how the sensitivity of the tumor population changes with the given dose for each strategy. A point is placed for each month in time at the average sensitivity and the average dose of each treatment over that month. Points get larger in size at each represented time point.*

If we plot the previous examples as trajectories though the space of average sensitivity over average dose, we can see why failures occur for the invasive (left) and evolving (right) phenotypes (same examples as from **Fig. 8A-C** and **Fig. 9A-C**). Using the CT strategy means a constant dose. This is shown with the black line that does not vary in the x-direction. If all phenotypes were to lie above the dashed line (net zero proliferation), cure is possible. However, because there are a variety of existing phenotypes, the mean sensitivity decreases over time. As the most sensitive cells are eliminated and space becomes available the line is crossed where the more resistant cells can have a net positive growth and the tumor recurs. With the invasive phenotype AT1 kills off the most sensitive clones sequentially as the dose is modulated to keep the tumor at a constant size. The dose modulation strategy will shift to doses lower than the MTD to find the lowest dose that still kills the most sensitive phenotypes. Eventually, the most sensitive cells will be killed, but because the dose is low the next most proliferative cells will have a net positive proliferation rate and outgrow. The increased growth then causes a dose increase to keep proliferation in check. This essentially "picks off" sequentially sensitive phenotypes with ever increasing doses until the more resistant cells take a hold of the invasive edge. At that point, the resistant cells have a small (albeit positive) growth rate and there is limited chance of control. With the AT2 strategy that relies more on treatment vacations, there is a sharp transition between the quick drug-free expansions and MTD-induced contractions. In this case, with the most sensitive cells at the proliferating edge, the net decline during treatment is approximately equal to the net increase when the dose is zero. As long as mostly sensitive cells continue to lie on the outer rim, the tumor can be controlled.

With a drifting phenotype, cell sensitivity can become either more sensitive or more resistant. This allows movement along the y-axis. For all treatment strategies, the drug kills off the more sensitive cells, so the phenotypes eventually end up further into the resistant region where

the net proliferation rate is always positive. While this does not change the outcome much for CT, the outgrowth of just a few resistant phenotypes causes a quicker dose increase for AT1, which causes a quicker decline of the most sensitive cells. This results in an eventual recurrence. By using the AT2 schedule that shifts between extreme doses, the fastest proliferators are maintained during vacations that are also very sensitive during drug application in order to suppress other cells. Unfortunately, the effect is not sustained for very long.

4. Discussion

Our simulations demonstrate the necessity of matching any cancer treatment regimen to the corresponding intratumoral evolutionary dynamics. The term "precision medicine" is often applied to strategies that match tumor treatment to specific predictive biomarkers. While this approach increases the probability of success it neglects the reality that virtually all tumor responses are followed by evolution of resistance and, thus, treatment failure. Our results indicate that precision medicine also needs to encompass the complex evolutionary dynamics that govern emergence and proliferation of resistant phenotypes. Here we present the diverse Darwinian interactions that can lead to resistance but also corresponding treatment strategies that can exploit these dynamics to delay the time to progression.

It is clear from our analysis that there is no one-size-fits-all solution evolutionary strategy for all tumors. Using a modified off-lattice agent based model, we find conventional MTD application of cancer therapy *can* cure homogeneous tumors that consist entirely or almost entirely of cells that are sensitive to the applied treatment. While this is improbable at advanced stages, it is observed clinically as some tumors that histologically appear homogeneous, such as testicular cancer and some lymphomas, are frequently cured by conventional chemotherapy. However, decades of clinical experience, consistent with our model predictions, have found cure is not typically achievable in most common cancers that are highly heterogeneous and contain cells that are therapy resistant due to their inherent phenotypic properties or environmentally-mediated mechanisms (23).

In tumors that are not curable by available therapy, our models find that evolution-based strategies that exploit the fitness cost of resistance can delay treatment failure and tumor progression. Recent experimental studies have shown adaptive therapy can enforce prolonged tumor control (10), however, none of the prior theoretical models have included spatial dynamics and thus assume that the different tumor subpopulations are well mixed (14,15). This is in contrast to radiological and pathological images of tumor that show marked spatial heterogeneity (24-27). Spatial structure can affect the emergence of resistant phenotypes due to space limitation (18,28) and limited drug perfusion (29,30). We present some insight into how spatial context might influence disease control, which reflects similar ideas presented in a recent study with growing bacterial colonies of *E. Coli* (28)*.* In this study, Fusco *et al.* showed how less fit resistant clones can be trapped due to spatial constraints, but high intensity drug exposure leads to the competitive release of dormant mutants. Our model shows the same nature of competition, and we are able to test multiple treatment strategies on different compositions of phenotypes.

We consider phenotypic heterogeneity over a wide spectrum of tumor compositions. A larger variation in cell phenotypes beyond just sensitive and resistant (with corresponding fitness differences) can drive changes in the dominating population at different doses. By systematically testing an array of different initial phenotypic distributions we can delineate different regions where some schedules work better than others. We also find that some AT strategies are better

than others for different situations. For example, with cell migration and phenotypic evolution, we can apply a strategy with less dose modulation and more emphasis on treatment vacations to keep the population sensitive. The extreme dose changes keep the tumor responding quickly during both the growth phase and the treatment phase to either maintain the spatial structure in the case of migration or to prevent the evolution of less sensitive cells in the case of phenotypic drift.

Figure 11. *The model checks that tumors with dissimilar compositions can be controlled by adaptive scheduling (left). We tested the response for CT, AT1 (α=0.25, β=0.05), and AT2 (α=0.50, β=0.10) using the sum of all tumor cells to determine the dose for the AT treatments. The total tumor burden dynamics from all metastatic sites are shown at the top, while the individual spatial compositions at the end of the simulation are given below for each treatment strategy. Clinically, this model can be used to couple the total systemic tumor burden to the individual spatial distributions found from imaging (right).*

Comparing widely disseminated cancer, where the distribution of tumor cell subpopulations will likely vary among the metastatic sites, to our model, which only considers the response of single tumors may seem disconnected. It is thus important to ask how much variation in clinical response to CT and AT should be expected. Testing the treatments on a set of tumor metastases with mixed compositions shows that control can be achieved (**Fig. 11**). By using the total numbers of cells across all tumors to represent a systemic measure of burden, we find that with the CT strategy the more sensitive tumors respond with complete eradication, while the less sensitive tumors eventually recur. In contrast, both AT schedules can still control the disease. While the more sensitive tumors shrink, the less sensitive ones grow in size keeping the total population relatively constant. The full dose schedules are shown in **Fig. S3** along with a case where the metastases are heterogeneous but have all the same composition. In the latter situation, we get very similar results to what would happen with single independent tumors. Modeling the different spatial distributions of multiple tumors can also be used in conjunction with the periodic systemic measure of tumor burden to be tested, calibrated, and validated using blood biomarkers and imaging data. Clearly, there are timescale issues when using data from in vitro models (days), in vivo models (months), and clinical trials (years) (Fig. 1), but a more rigorous approach to fit data could be done to extract specific growth and response rates for conversion across scales.

The methods for actually measuring tumor burden non-invasively and frequent enough to direct treatment decisions are not well developed for all cancers. In the ongoing adaptive therapy trial on prostate cancer (11), PSA is used as a marker for disease burden. PSA may not be a perfect indicator, but it is readily available, standardized, and utilized (31). Other biomarkers are used for surveillance of various cancers with degrees of specificity and sensitivity (32). As technologies for serum biomarkers, circulating tumor cells, and cell-free DNA continue to be developed (33), we hope this challenge can be addressed to better manage advanced and disseminated cancers through periodic blood draws.

The model presented here is greatly simplified and abstracted to study the impact of spatial intratumoral heterogeneity on controlling tumor progression using adaptive therapy. We have ignored pharmacokinetics (34), spatially weighted dose dependence (35,36), microenvironmental influence (37), therapy-induced acquired drug resistance (37), and of course the 3rd dimension (38). All of these are possible extensions and modulations of this work, but we have chosen a simple starting point to understand how a spectrum of cell phenotypes compete for space under different drug strategies. Our work illustrates clearly the importance of using treatment response as a key driver of treatment decisions, rather than fixed strategies that ignore it. We strongly believe that the future of precision medicine shouldn't be in the development of new drugs but rather in the smarter evolutionary enlightened application of preexisting ones.

Acknowledgements

The authors gratefully acknowledge funding from both the Cancer Systems Biology Consortium (CSBC) and the Physical Sciences Oncology Network (PSON) at the National Cancer Institute, through grants U01CA151924 and U54CA193489.

Bibliography

1.   Lackner MR, Wilson TR, Settleman J. Mechanisms of acquired resistance to targeted cancer therapies. Future Oncology. 2012 Aug;8(8):999–1014.

2.   Saunders NA, Simpson F, Thompson EW, Hill MM, Endo-Munoz L, Leggatt G, et al. Role of intratumoural heterogeneity in cancer drug resistance: molecular and clinical perspectives. EMBO Mol Med. EMBO Press; 2012 Jun 25;4(8):675–84.

3.   Kam Y, Das T, Tian H, Foroutan P, Ruiz E, Martinez G, et al. Sweat but no gain: Inhibiting proliferation of multidrug resistant cancer cells with "ersatzdroges." Int J Cancer. 2014 Sep 2;136(4):E188–96.

4.   Crawford S. Is it time for a new paradigm for systemic cancer treatment? Lessons from a century of cancer chemotherapy. 2013 Jun 24;:1–18.

5.   Zeilinger AR, Olson DM. Competition between stink bug and heliothine caterpillar pests on cotton at within-plant spatial scales. … Experimentalis et Applicata. 2011.

6.   Hiley CT, Swanton C. Pruning Cancer's Evolutionary Tree with Lesion-Directed Therapy. Cancer Discovery. American Association for Cancer Research; 2016

Feb 4;6(2):122–4.

7.   Gillies RJ, Verduzco D, Gatenby RA. Evolutionary dynamics of carcinogenesis and why targeted therapy does not work. Nat Rev Cancer. 2012.

8.   Korolev KS, Xavier JB, Gore J. Turning ecology and evolution against cancer. Nat Rev Cancer [Internet]. Nature Publishing Group; 2014 Apr 17;14(5):371–80. Available from: http://dx.doi.org/10.1038/nrc3712

9.   Aktipis CA, Boddy AM, Gatenby RA, Brown JS, Maley CC. Life history trade-offs in cancer evolution. Nat Rev Cancer. Nature Publishing Group; 2013 Nov 11;13(12):883–92.

10.   Enriquez-Navas PM, Kam Y, Das T, Hassan S, Silva AS, Foroutan P, et al. Exploiting evolutionary principles to prolong tumor control in preclinical models of breast cancer. Science Translational Medicine. 2016 Feb 24;8(327):327ra24–4.

11.   Moffitt Cancer Center. Adaptive Abiraterone Therapy for Metastatic Castration Resistant Prostate Cancer

[Internet]. 2015. pp. 1–4. Available from: https://clinicaltrials.gov/ct2/show/NCT02415621

12.   Kreso A, O'Brien CA, van Galen P, Gan OI, Notta F, Brown AMK, et al. Variable Clonal Repopulation Dynamics Influence Chemotherapy Response in Colorectal Cancer. Science. 2013 Jan 31;339(6119):543–8.

13.   Chmielecki J, Foo J, Oxnard GR, Hutchinson K, Ohashi K, Somwar R, et al. Optimization of Dosing for EGFR-Mutant Non-Small Cell Lung Cancer with Evolutionary Cancer Modeling. Science Translational Medicine. American Association for the Advancement of Science; 2011 Jul 6;3(90):90ra59–9.

14.   Silva AS, Kam Y, Khin ZP, Minton SE, Gillies RJ, Gatenby RA. Evolutionary Approaches to Prolong Progression-Free Survival in Breast Cancer. Cancer Research. 2012 Dec 16;72(24):6362–70.

15.   Silva AS, Durand A, Ribeiro MC, Alsina M, Shain K, Baz R. An Evolutionary Approach for Personalized Therapy in Multiple Myeloma. AM. 2016;07(02):159–69.

16.   Moore N, Houghton JM, Lyle S. Slow-cycling therapy-resistant cancer cells. Stem cells and development. 2011.

17.   Melnyk AH, Wong A, Kassen R. The fitness costs of antibiotic resistance mutations. Evol Appl. 2015.

18.   Gallaher JA, Anderson ARA. Evolution of intratumoral phenotypic heterogeneity: the role of trait inheritance. Interface Focus. 2013;3(4):20130016.

19.   Gesztelyi R, Zsuga J, Kemeny-Beke A, Varga B. The Hill equation and the origin of quantitative pharmacology. Archive for history of …. 2012.

20.   Gardner SN. A Mechanistic, Predictive Model of Dose-Response Curves for Cell Cycle Phase-specific and -nonspecific Drugs. Cancer Research. American Association for Cancer Research; 2000 Mar 1;60(5):1417–25.

21.   Gelles JD, Chipuk JE. Robust high-throughput kinetic analysis of apoptosis with real-time high-content live-cell imaging. Cell Death & Disease. 2016.

22.   van Nieuwenhuijze AEM, Van Lopik T, Smeenk RJT, Aarden LA. Time between onset of apoptosis and release of nucleosomes from apoptotic cells: putative implications for systemic lupus erythematosus. Annals of the Rheumatic Diseases. BMJ Publishing Group Ltd; 2003 Jan 1;62(1):10–4.

23.   Morris L, Riaz N, Desrichard A, Şenbabaoğlu Y. Pan-cancer analysis of intratumor heterogeneity as a prognostic determinant of survival. Oncotarget. 2016.

24.   Gerlinger M, Rowan A, Horswell S. Intratumor Heterogeneity and Branched Evolution Revealed by Multiregion Sequencing. New England Journal of Medicine. 2012.

25.   Gillies RJ, Anderson ARA, Gatenby RA, Morse DL. The biology underlying molecular imaging in oncology: from genome to anatome and back again. Clinical Radiology. Elsevier; 2010 Jul 1;65(7):517–21.

26.   Gatenby RA, Grove O, Gillies RJ. Quantitative Imaging in Cancer Evolution and Ecology. Radiology. 2013 Oct;269(1):8–14.

27.   O'Connor JPB, Rose CJ, Waterton JC. Imaging intratumor heterogeneity: role in therapy response, resistance, and clinical outcome. Clinical Cancer …. 2015.

28.   Fusco D, Gralka M, Kayser J, Anderson A. Excess of mutational jackpot events in expanding populations revealed by spatial Luria-Delbruck experiments. Nature. 2016.

29.   Fu F, Nowak MA, Bonhoeffer S. Spatial Heterogeneity in Drug Concentrations Can Facilitate the Emergence of Resistance to Cancer Therapy. Antia R, editor. PLoS Computational Biology. 2015 Mar 19;11(3):e1004142.

30.   Shah AB, Rejniak KA, Gevertz JL. Limiting the development of anti-cancer drug resistance in a spatial model of micrometastases. Mathematical biosciences and …. 2016.

31.   Wallace TJ, Torre T, Grob M, Yu J, Avital I, Brücher B, et al. Current approaches, challenges and future directions for monitoring treatment response in prostate cancer. Journal of Cancer. 2014.

32.   Duffy MJ. Tumor markers in clinical practice: a review focusing on common solid cancers. Medical Principles and Practice. 2012.

33.   Lowes L, Bratman SV, Dittamore R, Done S. Circulating tumor cells (CTC) and cell-free DNA (cfDNA) workshop 2016: scientific opportunities and logistics for cancer clinical trial incorporation. International Journal of …. 2016.

34. Groh CM, Hubbard ME, Jones PF. Mathematical and computational models of drug transport in tumours. Journal of The …. 2014.

35. Rejniak KA, Estrella V, Chen T, Cohen AS. The role of tumor tissue architecture in treatment penetration and efficacy: an integrative study. Frontiers in …. 2013.

36. Kim MJ, Gillies RJ, Rejniak KA. Current advances in mathematical modeling of anti-cancer drug penetration into tumor tissues. Frontiers in oncology. 2013.

37. Sun X, Bao J, Shao Y. Mathematical Modeling of Therapy-induced Cancer Drug Resistance: Connecting Cancer Mechanisms to Population Survival Rates. Sci Rep [Internet]. 2016 Mar 1;6:22498–13. Available from: https://www.researchgate.net/publication/296627868_Mathematical_Modeling_of_Therapy-induced_Cancer_Drug_Resistance_Connecting_Cancer_Mechanisms_to_Population_Survival_Rates

38. Tang L, van de Ven AL, Guo D, Andasari V, Cristini V, Li KC, et al. Computational Modeling of 3D Tumor Growth and Angiogenesis for Chemotherapy Evaluation. Pappalardo F, editor. PLoS ONE. Public Library of Science; 2014 Jan 3;9(1):e83962–12.