Cornell University

Modeling Subclonal Evolution in Cancerous Tumors as a Strategic Game

Submitted May 2020, in partial fulfillment of
the conditions for the award of the degree **M.Eng Computer Science.**

Shubhom Bhattacharya; Sameer Lal
sb2287; sjl328

Supervised by Eva Tardos

Department of Computer Science
Cornell University

I hereby declare that this dissertation is all my own work, except as indicated in the text:

Signature _____

Date _____ / _____ / _____

I hereby declare that I have all necessary rights and consents to publicly distribute this dissertation via the Cornell University.

Abstract

A hallmark of the cancer tumor ecosystem is the rapid evolution of malignant cells which become resistant to treatment. Current treatment methods apply intense therapy to target the largest decrease in subclone populations. However, these treatment methods fail to account for the evolutionary dynamics of subclones, and therefore potentially expose patients to harmful side effects. In this project, we propose a model to account for the strategic interaction of the physician and the evolving subclones based on their resistance to therapy. Using this model, we describe treatment schedules that take advantage of inter-subclone competition for tumor ecosystem resources as well as prevent the most resistant subclones from dominating the tumor.

Acknowledgements

Thank you to the Cornell M.Eng degree program and especially to our advisor Eva Tardos for her guidance in this project.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Cancer is an incredibly complex and diverse group of disease. The American Cancer Society estimates that 1.8 million new cases of cancer will be diagnosed in the United States alone, and over 600,000 cancer patients will die in 2020[9]. The progression of cancer varies depending on type, treatment, and patient. However, all are characterized by an uncontrolled replication of cells [4]. One of the most dangerous pathological characteristics of a cancer is the mechanism of how malignant tumor cells invade other tissue or organs, a process known as *metastasis*[1]. The primary challenge of the physician is to limit cancer growth and protect healthy tissue as much as possible mitigating risk of metastasis.

1.1 Motivation

The primary motivation of this project is to provide a framework that extends an existing mathematical model by accounting for the selective force induced on subclones by physicians' actions. In this sense, such a framework can generalize to The following broad questions provided guidance to our research:

- Therapy is expensive. Can better "strategized" chemo reduce the cost of chemotherapy?

- Chemotherapy may cause other health problems and will almost certainly lower

[1] A non-comprehensive glossary of related biological terms is compiled for the readers' convenience in Appendix A

patients' standard of life. Can strategized chemo make therapy less strenuous for patients?

- Tumors of the same cancer may manifest various subclonal phenotypes in different patients. Can we demonstrate certain therapeutic strategies to be
- Does game theoretic modeling support current assumptions about terminal vs. acute cancer diagnoses?
- Can we use existing game theory framework, such as Price-of-Anarchy and Optimality, to estimate how much worse treatments are relative to each other for a particular patient?
- The best case for the physician is that all cancer cells are destroyed. Is there a sequence of actions that makes progress towards this outcome?

1.2 Aims and Objectives

In this project we attempt to answer these motivating questions from the previous section. The first objective is to realize a reasonable scope for modeling interactions at the cell, tissue, organ, and body level all carry unique assumptions. Next, we aim to define a framework that models subclonal growth and fitness through the course of physician treatment. Lastly, we hope to enable flexibility to adding multiple subclone populations, multiple treatments, and time series analysis. We hope to propose a model to reason about treatment in a more coherent way. An added benefit of this generalization is that other sequential, strategic problems can be modeled in a similar manner.

Chapter 2

Background and Related Work

2.1 Tumor Heterogeneity and Growth

2.1.1 Clonal Evolution

The groundwork for studying the somatic evolution of cancer was laid out by Peter Nowell [7]. One result in his study is that tumors initially consisted of a small number of progenitor cells that had been damaged by a carcinogen. The study further outlined that while the same cancer disease may share similarities between patients, selective forces can create quite different tumor environments. Furthermore, cancer tumors show successive adaption at a faster rate than healthy cells, demonstrating that the tumor ecosystem is rapidly changing. Davis et. al. compared several different patterns of subclonal evolution: linear, branching, neutral, or punctuated as depicted in the **Figure 2.1** , [1].

2.1.2 Overview of Therapy

There exist many treatment strategies for cancerous tumors. Surgical removal of tumor tissue can directly remove cancerous cells, but remnants can remain and metastasize. Furthermore, surgery is difficult in several sensitive anatomical regions. Radiotherapy can localize treatment by damaging genetic material in the effected region. Like surgery, radiotherapy can be harmful in certain regions of the body or for certain groups of people such as children. Furthermore, radiotherapy can harm healthy body cells at a high rate,

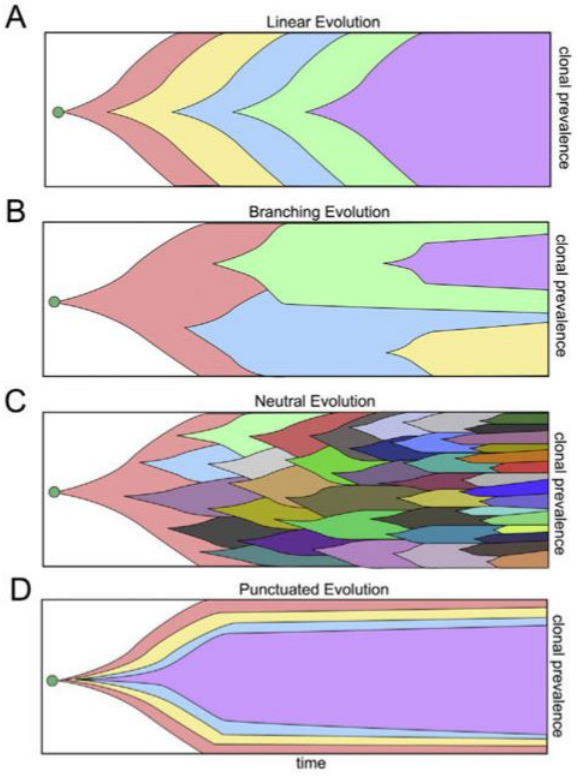

Figure 2.1: The patterns of evolution observed in cancer populations

causing potentially chronic aftereffects for patients. Chemotherapies deliver chemicals to rapidly growing cells in the body without discrimination (i.e. healthy cells such as hair cells are affected). Chemotherapy can be toxic in high doses to the body and also generally decreases the patient's standard of life. Other notable treatments include immunotherapy, which strengthens the body's immune response against cancerous cells, and the use of angiogenesis inhibitors, which prevent blood vessel formation needed for tumors to obtain essential oxygen and nutrients.

In this study, we focus on a current area of research involving the subclass of therapies called *targeted therapies*. Targeted therapies are designed to harm cancer cells based on their unique phenotype characteristics that healthy body cells do not share.

2.1.3 Subclonal Detection

Through recent developments in biopsy analysis, researchers can infer tumor population composition, especially through massively-parallel sequencing methods [10, 8] and ma-

chine learning-based approaches [2]. Massively-parallel sequencing methods show promise in detecting subclonal mutations without knowing the mutation *a priori* with a subclonal proportion as low as 1:5000 [8], i.e. even rare subclones can be identified using these methods.

2.2 Strategic Approaches

2.2.1 Game Theory

Game theoretic approaches have been used in the past to study competition between cancerous cells and healthy somatic cells. For example, [6] uses evolutionary game theory to model a cooperation game between osteoclasts, osteoblasts, and cancerous cells in multiple myeloma, a blood cancer. Similarly, [5] models an "evolutionary double bind" between two drug-resistant subclones and one susceptible subclone based on fitness functions that are parametrized by drug susceptibility and resistance. The aforementioned approaches suffer from several common deficiencies that we attempt to resolve. First, they assume a constant number of subclonal phenotypes (1 and 3, respectively), which is unrealistic. Next, while [5] studies the effect of physician therapy, it determines a therapeutic regime *a priori* rather than allowing physicians to adjust treatments and therefore fails to consider the physician's decision as a strategic interaction. Characterizing a treatment regime strategically allows definition of better treatment that *scales to as many subclones as are present or may become present* rather than relying on knowledge prior to the start of treatment.

2.2.2 Hamiltonian medicine

Another proposed approach borrowing from inclusive fitness theory is Hamiltonian medicine [3]. In addition to considering the traditional fitness effects of clonal reproduction, Hamiltonian medicine considers the beneficial collaborative reward obtained from cellular behaviors that in turn promote a specific set of genes. From this perspective, the growth of a subclone population could, in part, be linked to its propensity for "altruistic" behavior.

Though we do not account for these interactions in our proposed model, they present an added dimension of strategic interaction between subclones that merit further research.

Chapter 3

Framework

In this section, we first propose a general framework to model doctor-subclone evolution. Then, we provide a class of doctor strategies for each graph topology. Finally, we implement various greedy doctor heuristics and alternative treatment strategies and compare their progress towards minimizing average fitness of subclone populations.

3.1 Fitness Functions

In targeted therapy, doctors use drugs to target cancer cells without disrupting normal cells. We assume *well-mixing* in an oncological sense; i.e., a given cell in the tumor has an equal probability for interacting with each other cell. This allows us to ignore physical promixity on the tumor in our analysis and further assume that all populated subclones can interact with nonzero probability. As a result, we assume the administered medicine interacts with both the intended subclone colony and colonies with similar genetic makeup.

We extend the fitness equation described by Bastanta et. al. [5] to any k therapies. Formally, let the doctor's choice of treatment at time t be donated $d_t \in \mathbb{R}^k$ and $d_t^{(j)} \in [0, 1] \forall j$. The j-th index of d_t therefore corresponds to the dosage of treatment j at time

t. Subclone i is characterized by two time-invariant parameters:

$$c_i \in \mathbb{R} \in [0,1] \qquad \alpha_{\mathrm{i}} \in \mathbb{R}^{\mathrm{d}} \in [0,1]$$

c_i is a "cost of resistance" that represents phenotypic fitness as it relates to non-treatment related aspects of the tumor microenvironment. For example, mutations that provide resistance to some drugs for i may render it less able to perform important cellular functions, independent of whether a treatment is being administered. The other parameter, $\alpha \in [0,1]$, is a vector representing i's susceptibility to each drug, where $\alpha_i^j = 0$ signifies that subclone i is completely resistant to treatment j and $\alpha_i^j = 1$ signifies complete susceptibility. We define the *phenotypic vector* of a subclone i to contain these phenotypic parameters:

$$v_i = [c_i, \alpha_i^{(0)}, \alpha_i^{(1)}, ...\alpha_i^{(d)}]$$

Then, we define the fitness of a subclone i to be a scalar:

$$W_t(i) = \max(0, 1 - c_i - \alpha_i \cdot d_t) \in \mathbb{R}$$

Lastly, in the following discussion, we call the set of subclones N.

3.2 Proportion Update

At each step of the evolution, we update the proportion of each subclone $i \in N$ using the equation

$$p_{i,t+1} = \min(1, p_{i,t} \cdot \frac{W_i(t)}{\bar{W}})$$

maintaining the invariant that the sum of the proportions is 1. Here, $p_{i,t}$ denotes the proportion at time step t, $W_i(t)$ is the fitness of colony i at time t, and $\bar{W}$, defined to be

$$\bar{W} = \sum_i p_i W_i$$

is the average fitness of the tumor.

3.3 Doctor-Subclone Interaction

To extend previous work, we allow doctors to respond to changes in the relative populations and fitnesses of subclones. The evolution of the system is described as follows, where T_d is the set of times where the doctor is allowed to administer treatment.

for $t \in \{1, 2, \ldots\}$ **do**
  if $t \in T_d$:
  then
    `doctor applies treatment` $\mathtt{d}_t$
  end if
  `Recalculate subclone fitness`
  `Recalculate subclone proportions`
end for

At the end of the simulation, we calculate the average fitness and if it is below a certain threshold, we can conclude that we were able to treat the cancer. Biologically, this corresponds to the elimination of most cancer cells and/or a small likelihood of metastasis. If we are not able to achieve a low average fitness across many iterations of treatment, we conclude that we cannot tame the cancer and the doctor has lost.

3.4 Graphical Extension

We can further model the set of subclones as a weighted graph G(V;E) with vertex set V and edge set E. V corresponds to the subclone populations, and each weighted edge $e_{ij} = \cos(v_i, v_j)$ between nodes i and j represents the phenotypic similarity between the two subclones. The weight of the edge is determined through cosine similarity of each subclone's phenotypic vector, and thus constrained between 0 and 1. Intuitively, the graph models the genetic similarity of subclone populations in the tumor. If E is thresholded by $\gamma \in [0, 1]$, a number of connected components are formed. The thresholding operation

for each edge $e \in E$ between nodes i and j is:

$$e' = \mathbb{1}(e > \gamma) \cdot e$$

A graph representation allows physician strategies to be dependent on structural properties- for example, to promote diversity, a physician may want to maximize the number of components.

Chapter 4

Evaluation

4.1 Three Node System

First, we attempt to replicate the system used by Bastanta et. al. [5]. In this system, there are two medicines A and B. Subclone colony R_A is resistant to A and subclone colony R_B is resistant to R_B. Colony S is susceptible to both medicine A and medicine B. In our framework, this structure is represented as three nodes, R_A, R_B and S.
Using the same input parameters as described in [5] we were able to replicate the results achieved in the paper as seen in **Figure 5.1**.
The doctor's strategy is to attack R_A during the first half of treatment and then attack R_B during the second half of the treatment. Over time, the p_S diminishes (as it is susceptible to both treatments), the p_{R_A} population first decreases as p_{R_B} increases, and then p_{R_A} increases as p_{R_B} decreases.

4.2 Preliminary Strategical Reasoning

To compare the effect of treatment, the first realization is that without any treatment, this model favors subclone $i = \arg\min_{i \in N} c_i$ in the long term [1], as drug-resistance traits

[1] we use the phrase "long term" to denote behavior at times after the effects of initial population and fitness are eclipsed by the dynamics of therapy

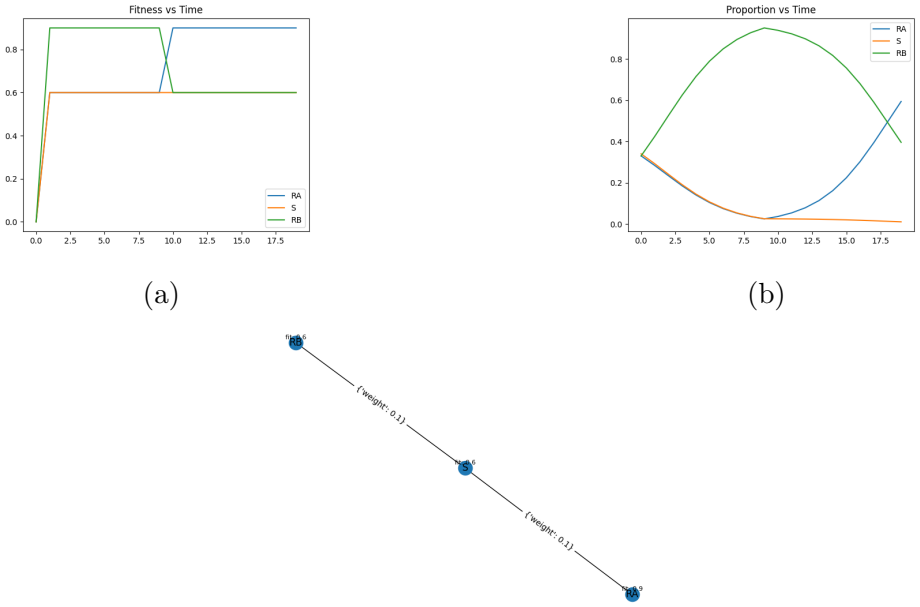

(a) (b)

(c)

Figure 4.1: Implementing topology by Bastanta et al.

render no fitness advantage. Furthermore, if i and j are subclones such that $\alpha_i^{(k)} \leq \alpha_j^{(k)}\ \forall k$ treatments, then no sets of doctor strategy can favor j over i. If $c_i \leq c_j$ then i dominates j and in the long term we can guarantee $p_i \geq p_j$, as $w_i \geq w_j$ for all time t. Lastly, average fitness with any treatment will be less than or equal to average fitness with no treatment at all, as fitness functions are monotonically decreasing with treatment dosage.

4.3 Greedy Doctor Strategies

A natural first set of doctor strategies to analyze are those that are greedy with respect to some metric of the tumor micro-environment. Intuitive metrics include subclone population, subclone fitness, or the product of the two. Recall that is N the set of subclones present, and that the dosage of treatment j at time t $d_t^(j)$ is constrained to [0,1].

A proportion-based greedy strategy will target the subclone with the highest tumor proportion. If $i = \arg\max_{i \in N} p_i$, then the drug that i is most susceptible to is $x = \arg\max \alpha_i$, and the doctor applies the drug at time t as $d_t[x] = 1$. Fitness follows as i is chosen by $i = \arg\max_{i \in N} W_i$ and the proportion-weighted fitness extends this to $i = \arg\max_{i \in N} p_i W_i$.

The simulation yielding Figure 5.2 and Figure 5.3 saw a greedy strategy with respect to

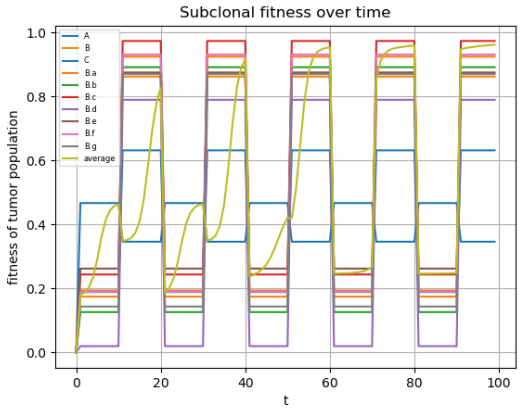

(a) Fitness

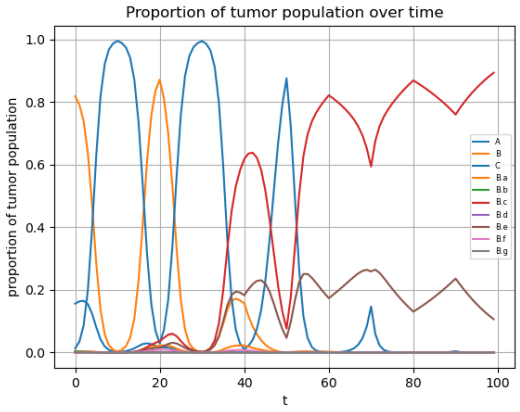

(b) Proportions

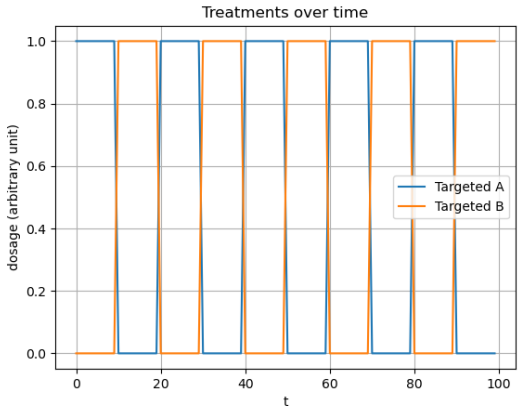

(c) Treatments

Figure 4.2: Greedy with respect to fitness metric

the proportion only and proportion-fitness product yield the same output. They differ only after $t = 60$ strategically, but this difference has an important effect on average fitness and the proportions observed. The first observation is that in both cases, after $t = 60$, the subclonal fitness has a periodic behavior, as a result of the changing treatments at each timestep. In the greedy fitness, the population of subclone B.c (a descendent of progenitor subclone B) is increasing on average, on average, while B.e is decreasing on average. In the other figure, there is a more overlapping behavior between the populations of subclones B.c and C.

We found that proportion-weighted fitness is, perhaps unsurprisingly, a better indicator of the danger a subclone can cause. Especially when resorting to a greedy strategy for the duration of treatment, a fitness-only strategy will often harm one subclone population while letting another dangerous subclone population grow. Meanwhile, a proportion-only

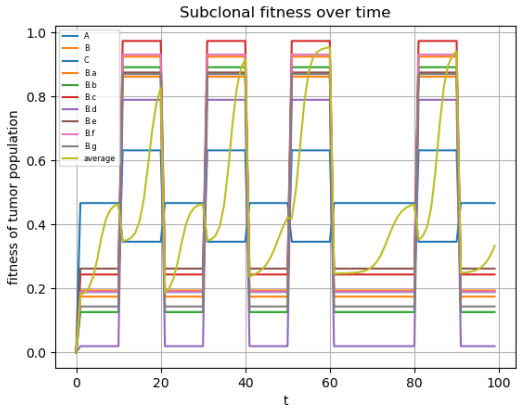

(a) Fitness

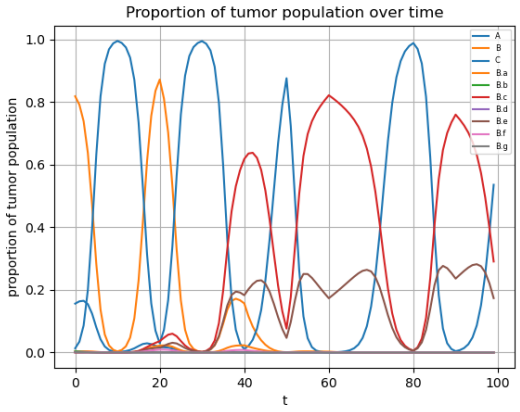

(b) Proportions

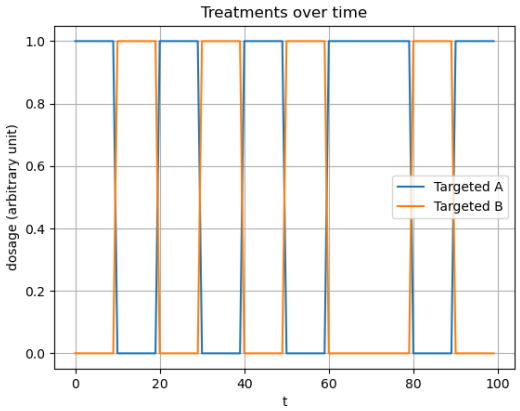

(c) Treatments

Figure 4.3: Greedy with respect to proportion and proportion-fitness metric (identical)

strategy can be useful as far as proportion being a signal of fitness in the long-term; indeed, because of this relationship between fitness and population, it is often the case that the targeted subclone is the fittest and most populous. In some cases, the choice of greedy strategy did not significantly alter subclonal growth and was thus inconsequential. This makes sense in situations where a given subclonal population is more fit than the others regardless of therapy applied, so that all greedy-strategies target it eventually.

4.4 Medicinal Takeaways

A qualitative result that follows biological modeling is that after several treatments, in most simulated cases, only 1 or 2 subclones remained, even with many subclones present to begin with. In a long-term perspective, this resembles linear and branching evolu-

tionary patterns, and is consistent with the fact that most mutations are not, in fact, advantageous.

We further attempted to simulate the dynamics of chemotherapy as opposed to targeted therapies. In these cases, greedy strategies were not employed; instead, the doctor acted according to human input responding to randomized initial conditions and graph structures. Most subclones are susceptible to chemotherapies, but as mentioned before, these treatments are harsh on the patients' healthy cells. A promising result we found is that a chemotherapy treatment regime punctuated by target therapies can maintain interclonal competition.

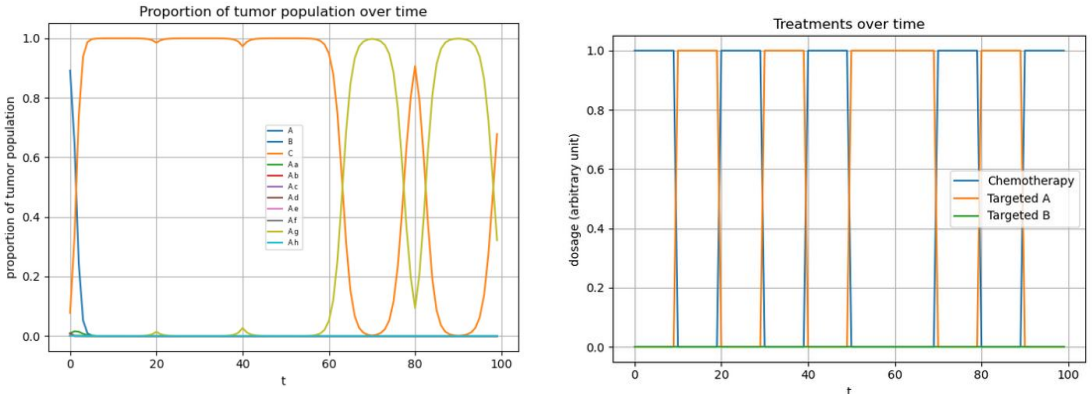

(a) After t=60, two subclones reach a stasis (b) Punctuated chemotherapy regime

Figure 4.4: Alternating therapies can maintain subclonal competition

We further attempted to assess the impact of using *dominated therapies*. In a strategic sense, such a therapy j is dominated by another therapy k if for all subclones, $\alpha_i^k > \alpha_i^j$; i.e. no subclone is more susceptible to j than k. In our example, we chose to represent this by a weak immunotherapy regime, which can vary significantly in effect based on the patient.

Fitness for all subclones increases at t=40 when the immunotherapy regimes as all subclones are more fit, but a targeted treatment at t=80 quickly reverses one subclone from dominating the tumor.

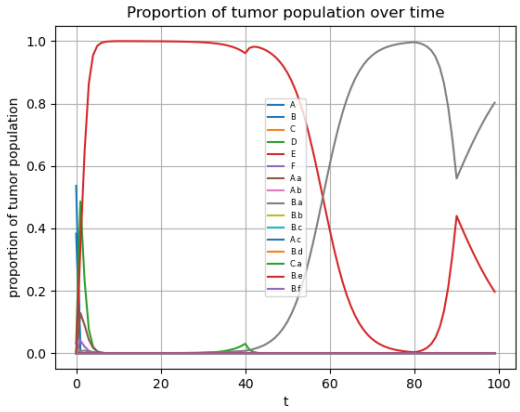

(a) Proportion effects of immunotherapy

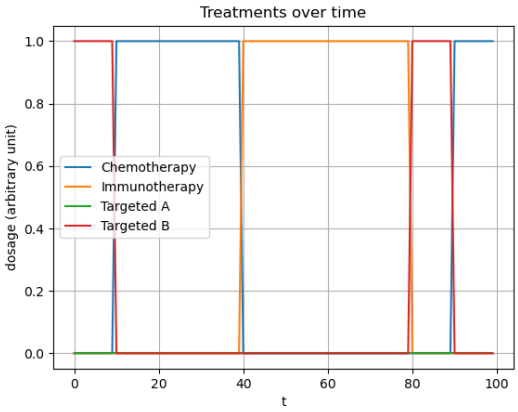

(b) Immunotherapy regime t=40 to t=80

Chapter 5

Conclusions

We have proposed a framework to describe the strategic interaction of a physician with subclones in the cancer tumor microenvironment. This framework extends an existing model of tumor growth by defining a strategic therapeutic regime dependent on subclonal phenotypic characteristics. We demonstrate the viability of this framework using some intuitive greedy heuristics as well as some treatments based on targeted therapy. While these simulations are not conclusive, they provide a basis for building patient-centered strategic therapies.

Chapter 6

Future Work

This framework provides many useful avenues for future research. First, generalizing to many subclones allows us to predict subclonal evolution using statistical or biological methods. Thus, this framework can continue to evolve in time as new nodes are added and possible long-term effects of therapy can be incorporated. Furthermore, the process by which subclonal populations evolve may be a stochastic process that can be incorporated into the birth and death of novel subclonal populations. For example, if descendant k varies from parent i enough that $|\alpha_k - \alpha_j| > \epsilon$, for some fixed ϵ, k can form a new component in the graph. An added advantage of this graph-based formulation is that strategic interaction based on graph properties, such as degree and clustering, can be represented, for instance, to suppress or encourage tumor diversity. Therefore, metrics that modify or replace the average-fitness used in our simulations may better characterize desirable properties such as diversity and growth rate.

Secondly, while the parameters of the subclone gene vector are defined to be time-invariant, future analysis or simulation can vary these by time to model added dependencies, for example, to model the previously mentioned altruistic behaviors that can build resistance with added population.

Lastly, we mentioned several strategies doctors can take, but our in phrasing the strategic problem of cancer treatment, we were unable to fully characterize a formal game. There is an inherent tradeoff in attempting to characterize the complexities of cancer and making reasonable assumptions in a computational approach, which we attempted via rigorous

background research and considering alternative modeling methods. Still, further work phrasing solution concepts in game theory (or, perhaps, new ones) can refine this work. For example, the long-term population dynamics of subclones could be understood in terms of evolutionarily stability concepts, and doctors' best responses could be understood in terms of other objectives. We could also analyze each strategy's effect on the potential function for this game, and look to decrease it at each treatment or come within an approximation of an optimal therapy. Designing this complex mechanism is difficult, but we believe this work is a reasonable foray into modeling interclonal competition in cancer therapies.

Chapter 7

Implementation

The simulations discussed were implemented in Python using standard libraries (numpy, pandas, and matplotlib). The graph analysis and structuring was performed using the Networkx package. Code can be accessed at this Github repository

Bibliography

[1] Alexander Davisa, R. G., and Navin, N. Tumor evolution: Linear, branching, neutral or punctuated? *Biochim Biophys Acta 369*, 1867 (2017), 151–161.

[2] Artem Solokov, Evan O.Paull, J. M. S. One-class detection of cell states in tumor subtypes. *Pac Symp Biocomput. 194* (2016), 405–416.

[3] Bernard Crespi, K. F., and Úbeda, F. First principles of hamiltonian medicine. *Philosophical Transactions of the Royal Society B 369*, 1642 (2014).

[4] Cooper, G. *The Cell: A Molecular Approach. 2nd edition.* Sinauer Associates, 2000.

[5] David Basanta, R. A. G., and Anderson, A. R. A. Exploiting evolution to treat drug resistance: Combination therapy and the double bind. *Molecular Pharmaceutics 9*, 4 (2012), 914–921.

[6] Jorge M. Pacheco, F. C. S., and Dingli, D. The ecology of cancer from an evolutionary game theory perspective. *Interface Focus* (2014).

[7] Nowell, P. The clonal evolution of tumor cell populations. *Science 194*, 4260 (1976), 23–28.

[8] Peter J. Campbell, Erin D. Pleasance, P. J. S. E. D. R. R. I. G. G. A. F. A. R. G. P. A. F., and Stratton, M. R. Subclonal phylogenetic structures in cancer revealed by ultra-deep sequencing. *Proc Natl Acad Sci U S A. 105*, 35 (2008), 13081—13086.

[9] Rebecca L. Siegel MPH, Kimberly D. Miller MPH, A. J. D. P. Cancer statistics, 2020. *ACS Journals* (2020).

[10] Roman K. Thomas, Elizbeth Nickerson, M. M. Sensitive mutation detection in heterogeneous cancer specimens by massively parallel picoliter reactor sequencing. *Nature Medicine 12* (2006), 852–855.

Appendix A

Glossary

The definitions in this glossary are based off of online websites (namely, Wikipedia and dictionary.com) and modified for context in this paper.

- Adaptive Therapy: different paradigm of treatment that depends on spatial and temporal evolution of cancer cells
- Carcinogen: A substance or agent causing cancer.
- Clone: a set of cells that all descend from a common ancestor cell. A clone is usually distinguished through inheritance of a distinctive genetic lesion (mutation) that occurred in the ancestor cell.
- Copy number variation: different DNAs can have different numbers of the same sequence expressed. The CNV characterizes these differences.
- *De novo*: from the beginning, anew
- Driver mutation: a mutation that gives a selective advantage to a clone in its microenvironment, through either increasing its survival or reproduction. Driver mutations tend to cause clonal expansions.
- Genotype: the genetic makeup of an organism, or, in this paper, a subclone
- *In silico*: performed on a computer i.e. computer simulation/experimentation

- Maximum Tolerated Dose (MTD): Treatment method in which patients are given the maximum dose they can handle without unintentional side effects
- Neoplastic progression: the somatic evolutionary process by which normal tissue changes into malignant (cancerous) tissue.
- Passenger mutation: a mutation that has no effect on the fitness of a clone but may be associated with a clonal expansion because it occurs in the same genome with a driver mutation. This is known as a hitchhiker in evolutionary biology.
- Phenotype: genotype (G) + environment (E) + genotype/ environment interactions (GE) phenotype (P)
- Progenitor: an ancestor in the direct line, used in this paper to refer to those cells whose interactions with a carcinogen generated the initial tumor
- Somatic evolution: The accumulation of mutations and epimutations in somatic cells (the cells of a body, as opposed to germ plasm and stem cells) during a lifetime, and the effects of those mutations and epimutations on the fitness of those cells.
- Subclone: A set of cells that shares the same phenotype and is distinguished from other cells by a different phenotype