NOVEL APPROACHES FOR OPTIMAL THERAPY DESIGN IN DRUG-RESISTANT POPULATIONS

by

DAVIS T WEAVER

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy

Dissertation Advisor: Dr. Jacob Scott

Department of Nutrition
Systems Biology and Bioinformatics

CASE WESTERN RESERVE UNIVERSITY

May 2023

CASE WESTERN RESERVE UNIVERSITY

SCHOOL OF GRADUATE STUDIES

We hereby approve the thesis of

Davis T Weaver

Candidate for the **Doctor of Philosophy degree***.

Committee Chair

Mehmet Koyutürk, PhD

Committee Member

Michael Hinczewski, PhD

Committee Member

George Dubyak, PhD

Committee Member

Jacob Scott, MD, DPhil

Date: April 6, 2023

*We also certify that written approval has been obtained for any proprietary material contained therein.

Table of Contents

List of Tables . . . iv
List of Figures . . . v
Acknowledgements . . . xii
Dedication . . . xiii
Abstract . . . xi

I Bioinformatic Methods for Drug Target Identification and Optimal Therapy Design in cancer. 1

1 Graph methods for drug-target identification. 2
1.1 Introduction . . . 3
1.2 A primer on Graph Theory . . . 4
1.3 Protein-Protein interactions networks . . . 6
1.4 Drug Target Identification . . . 11
1.5 Discussion . . . 16

2 Network potential identifies therapeutic *miRNA* cocktails in Ewing sarcoma 18
2.1 Introduction . . . 19
2.2 Results . . . 21
2.3 Discussion . . . 36
2.4 Methods . . . 39

3 An open-source R package to facilitate drug target identification. 49
3.1 Introduction . . . 50
3.2 Design and Data Sources . . . 52
3.3 Example Uses . . . 55
3.4 Discussion . . . 62
3.5 Algorithm Definitions . . . 63

II AI-enabled treatment of evolving cell populations 67

4 Reinforcement Learning informs optimal treatment strategies to limit antibiotic resistance 68
4.1 Introduction 69
4.2 Methods 71
4.3 Results 75
4.4 Discussion 96

III Discussion and Conclusion 99

5 Other contributions to clinical optimization 100
5.1 Introduction 101
5.2 Modeling outcomes for patients with co-morbid conditions 102
5.3 Modeling Treatment Outcomes in Ovarian Cancer 103
5.4 Computing the rate of progression for pre-malignant pancreatic lesions . . . 105
5.5 Forecasting the trajectory of the opioid pandemic 106

6 Discussion 108
6.1 Summary of Contributions 110
6.2 Next Steps 111
6.3 Conclusions 112

IV Bibliography 114
Bibliography 115

List of Tables

2.1 **Top protein targets for each cell line.** We ranked potential targets by predicted change in network potential when each protein was modeled as repressed. 34

2.2 **Top cancer-associated protein targets for each cell line.** We ranked potential targets by predicted change in network potential when each protein was modeled as repressed, limited to proteins causally associated in cancer according to the Cosmic database. Proteins that appear in the same position for ≥ 3 cell lines are **bolded**. 34

2.3 **Genes ranked by network potential are enriched for several biological pathways related to cancer as well as the miRNA bio-genesis pathway** Pathways with an adjusted p-value < 0.05 are shown above. "ES" refers to enrichment score and "NES" refers to the normalized enrichment score. "nMoreExtreme" refers to the number of random gene sets (out of 500) that were more enriched than the test set. Size refers to the number of genes in the pathway that were also present in our mRNA expression dataset. 35

3.1 **Top 10 proteins ranked by affinity to Ewing Sarcoma-related seeds.** . . 60

4.1 Reference codes for drugs under study 71

4.2 **Example drug sequences.** Here, we show the first 10 selected drugs for representative episodes of the three top-performing replicates. 77

4.3 Key Hyperparameters for reinforcement learner 86

List of Figures

1.1 **Examples of directed and undirected graphs** 5

1.2 **Examples of simple and multigraphs.** 5

2.1 **Cartoon describing rationale for focusing on miRNA combination therapy.** With single-agent therapy, both target mRNA and non-target mRNA are inhibited an equal amount, potentially resulting in toxicity due to off-target effects. With miRNA combination therapy, the common target mRNA is inhibited to a greater degree than any individual non-target miRNA. . . . 20

2.2 **Network potential demonstrates a different distribution compared to mRNA expression.** Main panel: scatterplot comparing mRNA expression and network potential for all genes in our 18 Ewing Sarcoma cell lines. For each gene, we averaged across all samples for both mRNA expression and network potential. Unlike network potential (top axis histogram), mRNA expression (right axis histogram) has a bimodal distribution. 22

2.3 **TRIM25, APP, ELAVL1, AND RNF4, and XPO1 are the top protein targets ranked by predicted disruption following *in silico* repression. Panel A:** Box and whisker plot describing the change in network potential following *in silico* repression for each of the top 50 proteins. 99.99% confidence interval from the permutation test are displayed alongside the box and whisker plots. It is notable that EWSR1, the kinase associated with Ewing sarcoma development, is considered highly influential in the cell signaling network by this method, even in comparison to the computed null distribution. Genes that have previously been causally implicated in cancer according to the Cosmic database are highlighted in red [1]. Essential housekeeping genes (excluding those that are causally implicated in cancer) are highlighted in blue. The heat-map on the x-axis corresponds to the protein-mRNA correlation of each gene in the Cancer Cell Line Encyclopedia [2]. **Panel B:** Histogram depicting the distribution of Pearson correlation between mRNA expression and protein expression from the Cancer Cell Line Encyclopedia for all nodes included in our final Ewing sarcoma cell signaling networks. Proteins that were ranked particularly highly in panel A were labeled in panel B. **Pancel C:** Bar chart describing the most frequently observed genes among the top projected targets for the 15 patient tumor samples we analyzed. There was very little overlap in top projected targets between the cell line and patient data, reflecting the transcriptional heterogeneity present in Ewing Sarcoma. 25

2.4 **Many of the most promising miRNA candidates repress large numbers of essential housekeeping genes.** We identified the top miRNA for treatment of Ewing sarcoma, ranked by their predicted disruption of the Ewing sarcoma cell signaling network. **A:** Boxplot showing the projected disruption in network potential for the top miRNA candidates (averaged across all samples). The heatmap on the x-axis describes the number of essential housekeeping genes that each miRNA is predicted to target. **B:** Scatterplot showing the relationship between projected network disruption and the number of putative mRNA targets for a given miRNA. **C:** Heatmap showing z-score normalized miRNA expression for 622 of the evaluated miRNA for the 6 cell lines under study. The Y axis is clustered by the projected ΔG associated with a given miRNA. There doesn't seem to be a clear pattern of miRNA expression based on projected ΔG. 26

2.5 **We identified miR-483-3p, miR-5695, and miR-4514s as the optimal 3-miRNA cocktail for Ewing Sarcoma therapy.** We identified cocktails that are predicted to maximally downregulate target genes (red shading on the figure), while avoiding downregulation of essential housekeeping genes to limit toxicity (blue shading on the figure). **Panel A.** shows the targeting heatmap for the best predicted cocktail for cell line A673. The miRNA that make up the cocktail are presented on the y-axis. Putative gene targets are highlighted on the x-axis. Lines that span multiple miRNAs occur when a gene is downregulated by 2 or more miRNAs in the cocktail. **Panel B.** shows a histogram of the number of microRNA that target a given housekeeping gene in the best cocktail. **Panel C.** displays the targeting heatmap for the worst-performing cocktail for cell line A673 among those tested for reference. **Panel D.** shows a histogram of the number of microRNA that target a given housekeeping gene in the worst predicted cocktail. **Panel E** shows a bar graph showing the miRNA that most frequently appear in either the bottom or top 10 predicted cocktails (averaged across cell lines) for Ewing Sarcoma therapy. 28

2.6 **Network potential describes different features of a cell signaling network compared to mRNA expression alone. Panel A:** Histogram of mRNA expression for each gene (averaged across all samples). **Panel B:** Histogram of the network potential for each gene (averaged across all samples) mRNA transcripts with an expression level of zero were excluded from both histograms to better visualize the distribution of genes that are expressed. **Panel C:** Box plot showing the total mRNA expression for each cell line and patient sample (patient samples begin with SJEWS). **Panel D:** Box plot showing the total network potential for each cell line and patient sample. 31

2.7 **Protein targets ranked by contribution to network stability.** When averaging across cell lines, XPO1, LMNA, EWSR1, HSP90AA1, and CUL3 were identified as the most important proteins in the Ewing sarcoma cell signaling network (when limiting our analysis to proteins causally implicated in cancer [1]). When each protein was simulated as completely repressed *in silico*, network potential was increased by 654, 456, 429, 425, and 399, respectively. The heatmap at the bottom of the plot describes the protein-mRNA correlation for each gene in the cancer cell line atlas. Grey indicates no data was available. It is reassuring that EWSR1, the kinase associated with Ewing sarcoma development, is identified as highly influential in the cell signaling network by this method. 32

2.8 **Overview of our findings, using the stringdb protein-protein interaction network. Panel A:** Scatterplot with marginal histograms comparing mRNA expression to network potential. **Panel B:** Box and whisker plot showing the change in network potential for the top 50 genes, as well as 99.99% confidence intervals from the permutation test. We also show a histogram comparing the top 50 genes identified by our pipeline using stringdb compared to biogrid as the protein-protein interaction network. . . 33
2.9 **Simplified schematic of our computational pipeline.** We defined a measure of tumor state, which we term network potential (Equation 2.1), based on both mRNA gene expression and the underlying protein-protein interaction (PPI) network. Next, we ranked mRNA targets based on their contribution to network potential of each cell line, aiming to approximate the relative importance of each mRNA to network stability. After identifying these mRNA targets, we then identified miRNA and miRNA cocktails that preferentially acted to repress the most influential of the ranked mRNA targets, with the aim of defining synthetic miRNA-based therapy for downregulation of these targets. 39

3.1 **Example workflow for crosstalkr software package.** crosstalkr supports all 3 of the most common steps in an interactomic pipeline; interfacing with online PPI repositories, pruning the resulting network, and ranking nodes to produce scored gene sets. 52
3.2 **Protein-protein interaction subnetwork for GAPDH, HSP90AA1, EEF1A1,HNRNPC, and TPT1. HSP90AA1 was identified as a critical hub in the computed subnetwork.** 59

4.1 **Schematic of artificial intelligence system for controlling evolving cell populations. A:** *E. coli* population evolving on fitness landscapes under the strong selection, weak mutation evolutionary regime. At each time step, a reward signal r and a measure of system state s are sent to the replay memory structure. **B:** Replay memory array stores (s, a, r, s') tuples where s' is state s+1. These are then used to batch train the neural network. **C:** Deep Neural network estimates the value of each action given information about the environment's state. The action with the largest estimated value is then applied to the evolving cell population. 72

4.2 **Performance of RL agents** in a simulated *E. coli* system. **A:** Line plot showing the effectiveness of the average learned policy as training time increases on the x-axis for RL agents trained using fitness (red) or genotype (blue). **B:** Boxplot showing the effectiveness of 10 fully trained RL-fit replicates as a function of noise. Each data point corresponds to a single episode. The width of the distribution provides information about the episode by episode variability in RL-fit performance. **C:** Density plot summarizing the performance of the two experimental conditions (measured by average population fitness) relative to the three control conditions. **D:** Signal to noise ratio associated with different noise parameters. Increasing noise parameter decreases the fidelity of the signal that reaches the reinforcement learner. 76

4.3 **Drug cycling policies learned by RL-genotype and RL-fit. A:** Heatmap depicting the learned policy for 100 replicates (on the x-axis) of the RL-genotype and 100 replicates of RL-fit. Far left column (enlarged) corresponds to the optimal policy derived from the MDP condition. The Y-axis describes the β-lactam antibiotics each RL agent could choose from while the color corresponds to the probability that the learned policy selected a given antibiotic. Bottom heatmap shows the median fitness benefit observed under the policy learned by a given replicate. **B:** Heatmap showing the average learned policy for RL-fit and RL-genotype. RL-genotype learns a more consistent mapping of state to action compared to RL-fit. 79

4.4 **Comparison of evolutionary trajectories seen under different regimes A-C:** Selected Pairwise comparisons of state transition frequency under different experimental conditions. State transition frequency is nearly identical for the RL-genotype and MDP conditions (R=0.97). In contrast, state-transition-frequency for the RL-fit and MDP conditions are related but less strongly correlated (R=0.75). As expected, state transition frequency is not very similar between the RL-fit and random conditions (R=0.62). **D:** Bar chart comparing the frequency that states are observed under different experimental conditions. The value of each state (to the learner) is highlighted for each state by the bottom heatmap. High value states are observed more frequently in RL-fit, RL-genotype, and MDP conditions compared to the random condition. 80

4.5 **Movement of simulated *E. Coli* population through the genomic landscape. Top row:** Heatmap depicting the joint probability distribution for each state transition under the different experimental conditions. The second two show the difference in state transition probability compared to the MDP condition. **Bottom row:** Graph depicting the fitness landscape, beginning with the wild type (bottom) all the way to the quadruple mutant (top). Size of arrow depicts the frequency with which a state transition was observed under the labeled experimental condition. The color of each node corresponds with the expected value (to the learner) of being in that state. As above, the second two arrows correspond to the observed difference between RL-Fit or RL-genotype and the MDP condition. . . . 82

4.6 **MDP value function closely matches opportunity landscape for drugs commonly used under MDP policy.** Panels A and B show the 16-genotype fitness landscape under study, starting with the wild type at the top, progressing throw the single mutants, double mutants, triple mutants, and finally the quadruple mutant at the bottom. **A:** Opportunity landscape for the 5 drugs most commonly used under the MDP policy (CTX, CPR, AMP, SAM, and TZP). **B:** Observed state transitions under the MDP policy. The node color corresponds to the value function estimated by solving the MDP. Lower values correspond to states the MDP policy attempts to avoid while higher values correspond to states the MDP policy attempts to steer the population. **C:** Scatter Plot showing the distribution of fitness with respect to genotype for the 15 β -lactam antibiotics under study. The drug selected by RL-genotype in a given genotype is highlighted in light blue. In cases where the MDP selected a different drug than RL-genotype, that drug is highlighted in orange. **D:** Number of genotypes with fitness above or below 1 for each drug under study. Drugs that are used by both the MDP and RL-genotype are highlighted in orange. Drugs that are used by only the MDP are highlighted in green. Drugs that are used by only RL-genotype are highlighted in blue. 85
4.7 **Performance of RL-fit and RL-genotype for each replicate. A:** Fitness observed under RL-fit policy compared to random drug cycling condition. **B:** Fitness observed under RL-fit policy compared to fitness observed under optimal policy. **C:** Fitness observed under RL-genotype policy compared to random drug cycling condition. **D:** Fitness observed under RL-genotype policy compared to fitness observed under optimal policy. 87
4.8 **Identification of policy groups using PCA and kmeans clustering.** RL-fit replicates were separated into 2 distinct replicates. RL-genotype replicates were separated into 3 distinct groups but it is unclear if these groups are meaningfully different. 88
4.9 **policy heatmaps for groups identified using PCA and kmeans clustering.** Groups 1,2, and 4 correspond to RL-genotype policies. Groups 3 and 5 correspond to RL-fit policies. color gradient represents probability that a given drug (x-axis) will be selected when population is in a given state (y-axis). 88
4.10 **Opportunity Landscape for MDP-derived policy.** Opportunity landscape is an optimistic combination of 5 empirically measured drug landscapes. Just 1/16 genotypes is near a fitness peak on the opportunity landscape, helping to explain the extremely low fitness observed in the simulated E.Coli population when the MDP-derived policy is applied. . . 90
4.11 **Opportunity landscape for most common policy identified in the RL-genotype condition.** As in the MDP-derived policy, just 1/16 genotypes is near a fitness peak in the opportunity landscape. 91
4.12 **Opportunity landscape for the most common policy identified in the RL-fit condition.** The most common RL-fit policy relies on AMC and CTX to control the E. Coli population. Assuming the most optimistic combination of these two drug landscapes, 4/16 genotypes are near a fitness peak. 92
4.13 **MDP-derived optimal policy.** 93
4.14 **Effect of variation in gamma on optimal policy performance and composition** 94

4.15 Two-drug cycling policies. 95

5.1 Projected benefit in life years as a function of patient comorbidity and age. 103

5.2 Life expectancy (LE) differences (in months) between two strategies (“test” – “standard-of-care”) for treatment selection in patients with stage IIIC ovarian cancer, across varied performance assumptions for a hypothetical test. The black line represents test characteristics where the life expectancy gain = 0. High test specificity (correct triage of patients with resectable disease to primary cytoreductive surgery) is more important than high test sensitivity (correct triage of patients with unresectable disease to neoadjuvant chemotherapy plus interval cytoreductive surgery) for achieving life expectancy gains. Even with a perfect test (top right), the life expectancy gain is modest (1.2 months) 104

Acknowledgements

I would like to first thank my PI and mentor, Dr. Jacob Scott. His joy for science is infectious. As far as I can tell, the only thing he cares more about than discovering new things about the universe is the well-being of his trainees. Thank you to Dr. Andrew Dhawan and Dr. Jeff Maltas for being friends and mentors over the last 4 years. I would also like to thank my first scientific mentors, Dr. Pari Pandharipande and Dr. Amy Knudsen, who taught me (among many other things) how to write science.

I would also like to thank all the members of Theory Division (Jake's lab). In particular, thank you to my office mates - Eshan King, Dr. Jess Scarborough and Dr. Rowan Barker-Clarke for making work fun to go to.

I would like to acknowledge the NIH (T32GM007250) and the Case Comprehensive Cancer Center (Cancer Center Trainee Award for Cancer Research) for providing financial support.

Finally, I could never have written this thesis without the love and support of my wife, Emily Weaver.

Dedication

This work is dedicated to memories of Lynn Grogan, Michael Tucker, Barbara Tucker, and Bob Weaver. Their battles with cancer inspired me to spend my career studying it.

Novel Approaches for Optimal Therapy Design in Drug-Resistant Populations

Abstract

by

DAVIS T WEAVER

The current maximum tolerated dose treatment paradigm for cancer and bacteria fails to account for the capacity of these disease agents to evolve. When treatment fails to achieve rapid extinction, drug resistant clones rapidly proliferate into an uncontrollable tumor. To make significant progress for cancer patients, we need to better understand the evolutionary processes that drive cancer, and design treatments that explicitly account for them. In this dissertation, I will describe 3 projects that support the design of evolutionary therapies that explicitly account for the capacity of cancer (and bacteria) to evolve resistance in response to drug therapy. In **Chapter 2**, we developed two novel methods to support precision targeting of tumors; a novel leave-one-out style method for node ranking and a novel algorithm for ranking miRNA combinations that maximizes tumor disruption while minimizing toxicity. In **Chapter 3**, we described crosstalkr, an open-source software package to facilitate interactomic analyses. In **Chapter 4**, we described a novel approach for designing evolutionary therapies that leverages reinforcement learning to learn drug cycling policies given only limited information about an evolving system.

Part I

Bioinformatic Methods for Drug Target Identification and Optimal Therapy Design in cancer.

Chapter 1

Graph methods for drug-target identification.

1.1 Introduction

In the last 25 years, graph or network-based methods have transformed biological science. Interactomics, or the use of large-scale biological network data, has enabled more effective drug target identification, disease gene prioritization, drug repurposing, and allowed researchers to more deeply understand the signaling context of specific proteins and phenotypes [3]. The rapid expansion of tools and applications that leverage interactomic resources has led some to declare that we are living in the biological network era [4]. Interactomics includes the analysis of varied biological networks, including coexpression networks [5–8], metabolic networks [9, 10], drug-target networks [11], and protein-protein interaction networks (PPIs) [12]. Some resources combine different evidence channels to generate integrated "Functional Association Networks" [13, 14].

We will focus our discussion on genome scale interaction networks, particularly those provided by centralized databases. These interaction networks attempt to represent known interactions among proteins, genes, drugs, and more to facilitate integrative analyses that leverage the molecular context of a given protein or pathway [15]. Genome-scale functional interaction networks have become the default interactomic data resource for the vast majority of new methods for drug target identification [13]. Pipelines for drug target identification can be summarized as three-step processes. First, an interaction database is queried and structured as a graph. Second, the genome-wide graph is pruned to a more tractable subgraph. Third, a node scoring or ranking algorithm is applied to generate potential targets for drug therapy. Additional phenotypic data (i.e. RNA sequencing) can be incorporated during steps 2 or 3.

In this chapter, we will introduce graph-based data structures, notation, and algorithms. Next, we will discuss the landscape of genome-scale interaction databases. We will then describe graph-based methods that contextualize genome-scale interaction databases and identify phenotype-specific subnetworks. Finally, we will highlight major advances in drug development that have been made possible by these methods.

1.2 A primer on Graph Theory

A graph (usually called a ‘network’ in biology) is a relational data structure defined by pairwise interactions between objects. Graphs are composed of a set of nodes or vertices (V). Nodes are connected by a group of edges (sometimes called interactions, links, or arcs) E. A graph can therefore be described as $G = (V, E)$. Imagine a simple signaling network comprised of 5 proteins, EGFR, EGF, KRAS, PI3K, and PIP3. We can represent this process as a graph by describing a set of nodes $G(V)$ and edges $G(E)$:

$$\begin{aligned} G(V) &= \{EGFR, EGF, KRAS, PI3K, PIP3\} \\ G(E) &= \{(EGFR, EGF), (EGFR, KRAS), (EGFR, PI3K), \\ &\qquad (KRAS, PI3K), (PI3K, PIP3)\} \end{aligned}$$

.

As shown in **Fig 1.1**, an edge $e_{i->j}$ represents either a one-way (in the case of a directed graph) or a two-way binary interaction (in the case of an undirected graph) between two nodes v_i and v_j. Interactions can be either weighted or unweighted. Any node v_j that is connected by an edge to a node v_i is considered to be a neighbor of v_i. For example, $N(EGFR) = \{PI3K, KRAS, EGF\}$ while $N(PIP3) = \{PI3K\}$. Degree is a related concept, and describes the number of edges incident to a given node. In simple graphs

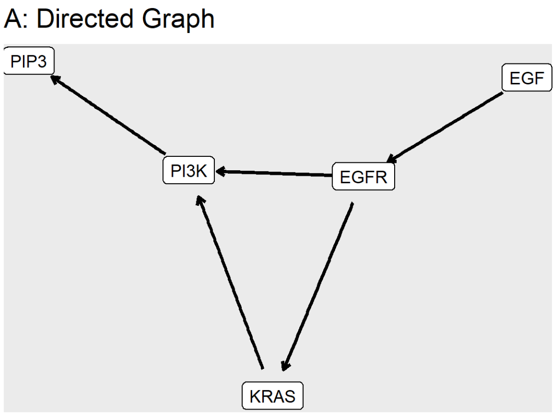

Figure 1.1: Examples of directed and undirected graphs

(i.e. no loops or multi-edges), $deg(v)$ is equal to the number of neighbors of node v. For example, $deg(EGFR) = 3$. Another key concept is the idea of a simple graph. A simple graph is constrained to only allow a single edge between a given pair of nodes. A multi-graph, by comparison, allows multiple edges between a given pair of nodes as well as "loops" (an edge between a node and itself). A comparison between a simple graph and a multi-graph is provided in **Fig 1.2**.

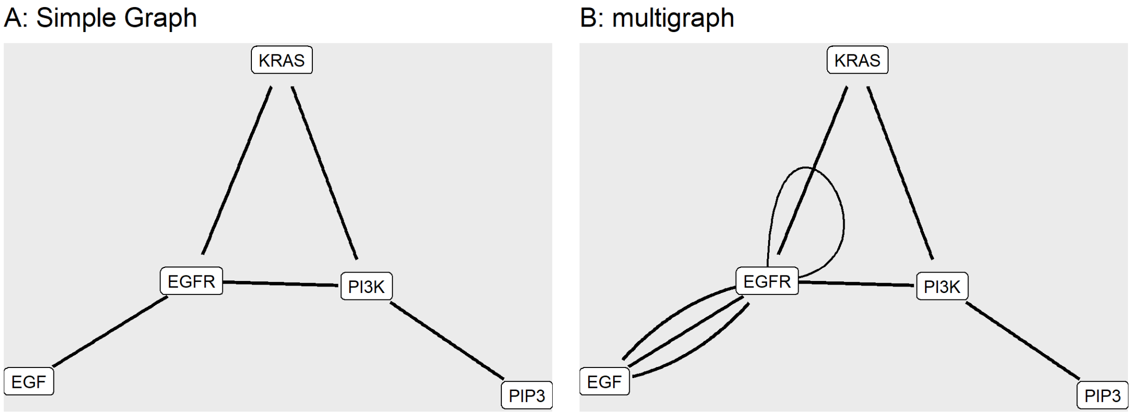

Figure 1.2: Examples of simple and multigraphs.

While there are certainly examples of loops and multi-edges in biology, most analyses of biological networks treat them as simple graphs. Graph methods make use of the relational structure of graphs. One key concept that is used frequently is a walk. A walk is a sequence of nodes defined by a starting node v_i and an ending node v_j. In the directed version of our graph G (shown in **Fig 1.1A**), the longest walk traverses all nodes and is of length 4 ($EGF- > EGFR- > KRAS- > PI3K- > PIP3$). A random walk is then a walk where the edge (if there are multiple options) that is traversed is chosen based on a given probability distribution (typically uniform). A biased random walk is a walk in which the edge probabilities are assigned based on some non-uniform probability distribution. For a more thorough introduction to the mathematical field of graph theory, see "Introduction to Graph Theory" [16].

Graphs are an excellent proxy for biological systems due to the widespread presence of interacting components. Proteins, metabolites, genes, and RNA complexes can all be represented as nodes. Biological processes like RNA interference, protein binding, metabolic cascades, and many other signaling processes can then be faithfully represented using edges. Once biological networks have been represented using graphs, additional experimental attributes such as protein or gene expression can be integrated.

1.3 Protein-Protein interactions networks

Protein-protein interaction networks (PPIs) are some of the most widely applied biological networks. PPIs represent biological signaling networks as graphs where the proteins are nodes and known interactions between proteins are edges [17]. Beginning in the early 2000s, many research groups around the world developed ambitious PPI databases that were meant to capture all known interactions between all known proteins [17–19]. These invaluable resources enable users to analyze proteins or pathways of interest in their local context. The graph-based structure also facilitates incorporation into neural networks [20].

In this section, we will discuss network inference, the process by which modern interaction databases are constructed from the primary literature. Next, we will summarize and evaluate freely available interaction databases. Finally, we will touch on the most pressing limitations that users must account for in their analyses.

1.3.1 Network Inference

Expert manual curation of the literature remains the gold standard for generating and updating interaction databases [21]. In this process, new edges are added to the interaction database when proteins are linked by a gold standard functional association assay in a peer-reviewed research paper [12,22]. Gold standard functional assays are based on the "guilt by association principle", and are designed to record an undirected interaction whenever two proteins come in close contact. Typically, researchers will tag a set of "bait proteins" that fluoresce when they are within a certain distance of a prey protein. One of the most common methods relies on Affinity Purification and Mass Spectrometry (AP-MS) to identify the proteins involved in a given binary interaction [23].

Expert manual curation is limited by the curators ability to find relevant papers and the supply of relevant papers in the literature. Text mining applications, including natural language processing, have helped reduce the first bottleneck by scoring and recommending relevant publications to teams of expert curators [24–26]. Some database projects have taken to using NLP to directly predict functional associations from the literature without relying on manual human checks [14, 15, 27].

In addition to direct evidence of interactions that are curated from the literature, there are a number of indirect evidence streams that are used to infer interactions. First, RNA sequencing experiments and high throughput mass spectrometry are used to infer protein interactions based on coexpression [15]. In addition, relative genomic position can be used to infer protein interactions, particularly in bacterial species where genes are known to be clustered based on shared functions [13]. Gene co-evolution can also be used to infer protein interaction based on the hypothesis that alleles that evolve together across many lineages are likely to be involved in the same function [13].

Finally, researchers have developed Bayesian and machine learning methods to predict protein interactions that are not adequately covered by other evidence streams. In one study, researchers incorporated protein structural information to predict more than 300,000 human protein protein interactions as accurately as comparable high-throughput screens [28]. Others have incorporated recent advances in machine learning like XGBoost to improve protein protein interaction prediction [29–31]. In addition, dramatic improvements in the prediction of protein structure [32] have been incorporated into structure-based prediction of protein interactions [33] to yield even more accurate predictions of protein complexes. As computational efficiency continues to increase, machine learning predictions of protein structure should enable dramatic improvements in genome-wide prediction of protein-protein interactions.

1.3.2 Available Resources

The 2000s and 2010s saw a dramatic expansion in publicly available tools and databases for interactomic analyses. Dozens of groups began ambitious curation efforts with the goal of converting thousands of academic papers into high-quality, machine-readable protein-protein interaction databases. According to pathguide.org, more than 700 functional interaction resources have been deployed [34]. Despite the large number of potential resources, the vast majority of bioinformaticians rely on a few primary interaction resources and a few meta-databases.

Primary Databases

Primary databases house empirical protein-protein interactions and are typically built through expert curation of the biomedical literature. The most used resources are Biogrid [12], the Database of Interacting Proteins [17], the Human Protein Reference Database [35], IntAct [22], and Mint [36]. As noted above, many many more of these resources have been created over the years, but they tend to be extremely difficult to find and use due to lack of ongoing support and development. IRefIndex helps to alleviate some of these concerns by providing a centralized warehouse for 10 of the largest available empirical PPI repositories [19]. The latest release of Biogrid contains 842,363 edges in the human binary interactome (after removing loops and multi-edges), more than twice as many as were available in 2019 [37].

Meta-databases

Primary databases remain limited to a subset of the true interactome. Empirical methods to identify PPIs have known biases and error rates that limit the ability of expert curation to capture all PPIs [29]. Meta-databases attempt to work around this shortcoming by integrating additional sources of evidence beyond expert-curated literature. These sources are described in section 1.3.1, and include RNA or protein co-expression, model predicted interactions, and gene co-evolution [13]. The most popular resources for genome-wide functional association data are String [13], HumanNet [14], and FunCoup [27]. These resources incorporate the same primary data, apply similar transformations, and even provide somewhat similar tools for the analysis of their data [13, 14, 27]. The vast majority of interactomic methods rely on these critical resources [13,14,27]. Stringdb stores 5,934,147 binary functional associations for homo sapiens, 5.2 million of which are estimated using methods other than expert curation of the primary literature.

1.3.3 Limitations

Despite new and cheaper assays for characterizing protein interactions, the literature clearly does not cover all possible functional interactions. As such, databases that rely solely on expert manual curation will suffer from selection bias. "Themed Curation Drives" - disease-specific curation efforts only exacerbate this problem [12]. Further, the high throughput assays that measure the vast majority of interactions that populate publicly available databases are not perfectly accurate, leading to possible error propagation throughout the literature [27]

The most comprehensive PPIs are available as undirected graphs. Directed graphs, which encode both unidirectional and bidirectional interactions, would be much more faithful representations of the underlying biology. However, they are far less comprehensive than their undirected counterparts due to limitations in PPI detection assays, and are likely to suffer from extreme selection bias. The vast majority of methods to measure PPIs empirically rely on bio-luminescence which is triggered when two proteins come in very close contact [23]. Direction of interaction cannot be inferred from these data.

Further, PPIs are typically measured in a generic context, or outside the context of the cell entirely. For example, Affinity Purification followed by Mass Spectrometry (AP-MS) allows researchers to measure thousands of interactions simultaneously. Unfortunately, these interactions must me measured in a cellular lysate, an unnatural extracellular "soup" of proteins [23]. Other popular methods are limited to certain cellular compartments, leading to enrichment of measured interactions in these areas and a comparable blind spot in other parts of the cell. For example, the yeast-two-hybrid method (Y2H) relies on the expression of a transcription factor that promotes the production of luciferase. The

Y2H method is therefore limited to PPIs that occur in the nuclear compartment [23]. Finally, the vast majority of methods for measuring PPIs are much more accurate when measuring durable PPIs compared to transient ones. Overall, PPI databases are likely to be biased towards strong interactions that occur in the nuclear or submembrane regions in the cell [23].

PPIs are also not typically measured in the context of a given disease state. This is particularly problematic when studying cancer due to the known re-wiring of the healthy cell signaling network that occurs during oncogenesis [38]. Tissue-specific and disease-specific PPIs can be generated by combining the full PPI network with gene expression or protein expression data. Edges between nodes that are not jointly expressed in the target phenotype are pruned, resulting in a more specific network. Unfortunately, this process cannot account for the generation of new interactions in the disease state, a particularly damaging limitation in the context of cancer. In one study, researchers measured the EGFR PPI network for two cancer cell lines, one with low KRAS expression and one with high KRAS expression [38]. Of the roughly 4,420 interactions they discovered in total, less than 20% were available in genome-scale PPI repositories. Further, more than 1,000 of the discovered interactions were edges in just one of the conditions, highlighting the context-dependent nature of cell signaling networks [38].

1.4 Drug Target Identification

Researchers have leveraged functional association databases to develop novel methods for drug target identification and disease gene prediction. While drug target identification may be used to describe the process of predicting an interaction between a specific compound and a physical target, we will focus here on the process of identifying critical proteins that are likely to be useful targets for drugs in a given disease. This process is sometimes called "disease-gene prediction" or "disease-gene association". Many of these methods make use of multiple phenotypic measurements, using an interaction database as a scaffold [39–41]. In **Table 1.4**, we highlight some of the graph-theoretic approaches that have been developed for this purpose.

Category	Methods	citations
Centrality Measures	betweenness, degree centrality, closeness, local average connectivity, DiffSLC	[42–52]
Path-based	Random-walk with Restarts, network flow, shortest path	[53–59]
Thermodynamic	Gibb's Free Energy, Entropy, Complexity	[39, 40, 60–63]
Leave one out	in-silico knockdown, network rewiring	[39, 40, 45, 63, 64]
Machine Learning	GNN, random forest	[41, 65–67]

1.4.1 Graph Reduction

As mentioned above, a key limitation of publicly available PPI databases is a lack of context-specificity. As a result, bioinformaticians have developed a number of methods to identify phenotype-specific subnetworks using a publicly available genome-wide PPI as a starting point. Smaller, phenotype-specific subnetworks also tend to be more easily interpreted and computationally tractable.

Graph reduction techniques can be stratified into many categories, and have been summarized and benchmarked extensively elsewhere [68]. For the purposes of this chapter, we will briefly summarize path-based approaches for graph reduction. It should be noted that unsupervised module identification can be accomplished with a number of another methods, including kernel-based methods, machine learning methods, modularity optimization, and more [65, 68].

Path-based approaches take advantage of the local structure of the protein-protein interaction network to score proteins based on affinity to one another. Initial attempts to detect phenotype-specific subnetworks focused on simple nearest neighbor and shortest-path based approaches [54]. Using known disease-gene interactions as seeds, these approaches categorized direct neighbors or near neighbors as candidate genes for a novel disease-specific subnetworks. Modern approaches instead rely on graph search algorithms such as random walk with restarts to traverse the interactome [54, 55, 68]. Random Walk with restarts is an extremely well studied algorithm that has been used to prioritize nodes in wide-ranging graph applications [69]. Many formal descriptions of random walk with restarts are available in the literature [70]. Briefly, random walk with restarts performs many random walks on a given graph G. A random walk is a stochastic process that begins at a given node V_i and traverses the edges E of a graph according to some edge probability such that the walk visits nodes V_j. Each step, there is a probability of restarting the process at the original node V_i. Nodes V_j are then scored based on the frequency with which they are visited on the walk that began at V_i. These scores can be used as a measure of node V_j's affinity for node V_i. V_i can be either a single node or a set of several nodes, in which case the random walk begins at a randomly selected node from among the node set V_i.

RWR can be solved extremely quickly using matrix multiplication [70], making it an attractive method for subnetwork identification on extremely large networks. In one study, researchers used a random-walk based approach to generate a colorectal-cancer specific subnetwork based on previously defined colorectal cancer-specific seed proteins

[55]. Network flow is a related approach that models information flow through a graph from sources to sinks [57]. Network flow has been used extensively to combine multiple biological networks into integrated analyses.

Many bioinformaticians use prior data to inform the subnetwork detection process. The most commonly integrate form of phenotypic data are previously identified disease genes. Previously identified disease genes can be used as seeds to bias path-based methods like random walk with restarts [55]. Selection of seeds can be informed through experimental evidence or publicly available repositories of known disease-gene interactions such as OMIM [71]. Phenotype-specific RNA-sequencing data is also commonly used to inform the subnetwork detection process [40,41,65]. For example, scientists have filtered nodes in the full PPI based on an RNA expression cutoff. Researchers have also limited the PPI to only genes that are significantly differentially expressed in relevant RNA sequencing data. Finally, researchers have computed a number of derivative metrics from RNA expression that have then been used to reduce the size of the PPI [39, 40]. Protein expression data would be even more informative for this context, but is rarely measured genome-wide due to technical constraints [2].

1.4.2 Node Ranking

In the last 20 years, researchers have developed or repurposed a number of different approaches to rank nodes for the purpose of drug target identification or disease-gene prioritization. Many researchers have applied or developed measures of centrality to use for node ranking [42–52]. Centrality measures are hypothesized to effectively rank proteins by importance due to the scale-free nature of PPI networks [48]. Scale-free networks tend to have hub and spoke architectures. Such networks are highly susceptible to targeted disruption of the hub nodes [48, 72].

For example, researchers developed the radiation sensitivity index (RSI) by ranking differentially expressed genes by degree to build a gene signature out of critical hubs in the cell signaling network [52]. RSI has since been validated as a predicted signature of radiation response across cancer subtypes, and has the potential to transform radiation dosing [73]. In another study, researchers developed a network-science pipeline to identify key proteins associated in resistance to first-line therapy to glioma, the most common brain cancer. They first identified proteins that were associated with the glioma phenotype in the literature. Next, they used these proteins to construct a glioma-specific subnetwork using shortest-path methodologies to filter the genome-wide interaction database. Finally, they used degree and betweenness to rank nodes to identify proteins critical to treatment-resistant glioma and validated these targets *in silico* [44].

The goal of improving disease-gene prioritization has led some researchers to develop novel measures of node centrality [49–51]. For example, researchers in 2014 defined hierarchical closeness, a structural centrality measure that outperformed other centrality measures at predicting known disease genes [49]. In another study, researchers defined local average connectivity to fulfill a similar funtion [51].

Thermodynamic measures have also been used to prioritize nodes in biological networks [39, 40, 61–63, 74, 75]. These measures attempt to represent fundamental properties of a network like energy or complexity using formulae imported from chemistry or physics. Finally, these ranking functions can be used in combination with "leave one out" style network disruption analyses to attempt to quantify the relative importance of each node in a given network. We will discuss leave one out and thermodynamic methods in greater detail in **Chapter 2**.

1.5 Discussion

Network biology is transforming biological research and biotechnology. Significant investment in PPI repositories has led to the wide-spread availability of genome-scale interaction networks with millions of putative edges [14, 15]. Researchers and industry scientists are using these resources to add signaling context to a wide variety of analyses. One of the most common use cases for network biology is disease-gene prioritization, a critical step in modern drug development pipelines. In this review, we described the three main components of a bioinformatic pipeline for drug target identification or disease-gene prioritization; online PPI repositories, subnetwork identification, and node ranking. We first defined core concepts in graph theory. Next, we surveyed available online PPI repositories, including the process of network inference and its limitations. Finally, we described methods for phenotype-specific subnetwork identification and node ranking.

While significant progress in network biology has been made in the last 20 years, limitations remain. Even with nearly a million curated interactions available in the most recent iteration of Biogrid, new PPI studies routinely measure 80% novel interactions, suggesting relatively poor coverage of the true interactome. New research suggests the relatively low coverage is due to the context-specific nature of protein protein interactions [38]. Measuring most of the true interactome will require a large number of PPI detection studies in every cellular context imaginable. Further, most genome-wide association databases do not provide information about the context in which interactions were measured, forcing researchers to rely on subnetwork identification methods to contextualize the interaction network and reduce noise. One notable exception is the Integrated Interactions Database (IID), which provides functional annotations along with binary interactions [76].

Despite the cambrian explosion of methods and online tools to facilitate interactomic analysis, comprehensive benchmarking for different node ranking methods is not available in the literature. In contrast, researchers have attempted to systematically benchmark algorithms for subnetwork identification. In one study, researchers set up a Kaggle-style data science competition to evaluate dozens of subnetwork identification methods using GWAS-detected gene sets as a ground truth. They then reported which methods and classes of methods were most effective at identifying these network modules across a diverse set of complex diseases. They found that Kernel-based approaches, random-walk based methods, and modularity optimization methods to be the most popular and effective [68]. Similar benchmarking studies to assess the many novel node ranking methods are a major unmet need in network biology. Despite these limitations, network biology continues to improve the drug development process and help scientists bring advances to patients.

Chapter 2

Network potential identifies therapeutic *miRNA* cocktails in Ewing sarcoma

Published as: **Weaver, D. T.**, Pishas, K. I., Williamson D., Scarborough J. A., Lessnick S. L., Dhawan A., &Scott J. G. (2021, October). Network potential idetifies therapeutic *miRNA* cocktails in Ewing sarcoma. *In Plos Computational Biology* DOI: `https://doi.org/10.1371/journal.pcbi.1008755`

2.1 Introduction

Ewing sarcoma is a rare malignancy arising from a gene fusion secondary to rearrangements involving the EWS gene [77]. There are 200-300 reported cases each year in the United States, disproportionately affecting children [78]. High levels of inter-tumor heterogeneity are observed among Ewing sarcoma patients despite a shared EWS gene fusion initiating event [79]. Ewing sarcoma is also extremely prone to developing resistance to available chemotherapeutics [80]. These features make it an ideal system to develop personalized therapies for resistant tumors or to avoid the development of resistance altogether.

MicroRNA (miRNA)-based therapeutics, including anti-sense oligonucleotides, are an emerging class of cancer therapy [81]. Recent work has highlighted the critical importance of miRNAs in the development and maintenance of the cancer phenotype [80–82]. MiRNA dysregulation has been implicated in the development of each of the hallmark features of cancer [83], and restoration of expression of some of these critical downregulated miRNAs has been studied as a potential treatment for several different cancers [82, 84, 85]. In particular, in the past decade, anti-sense oligonucleotide inhibitors of the STAT3 transcription factor have shown promise in the settings of lymphoma [86, 87] and neuroblastoma [88]. MiR-34 has shown to be effective in pre-clinical studies for treatment of both lung cancer [89–91] and prostate cancer [92]. Finally, miR-34 and let-7 combination therapy has been shown to be effective in pre-clinical studies of lung cancer [91].

MiRNAs have been recognized as potential high-value therapeutics in part due to their ability to cause widespread changes in a cell-signaling network [81]. A single miRNA molecule can bind to and repress multiple mRNA transcripts [82, 93–95], a property that can be exploited when designing therapy to maximally disrupt a cancer cell signaling network. This promiscuity of miRNA binding may also increase the risk of off-target effects

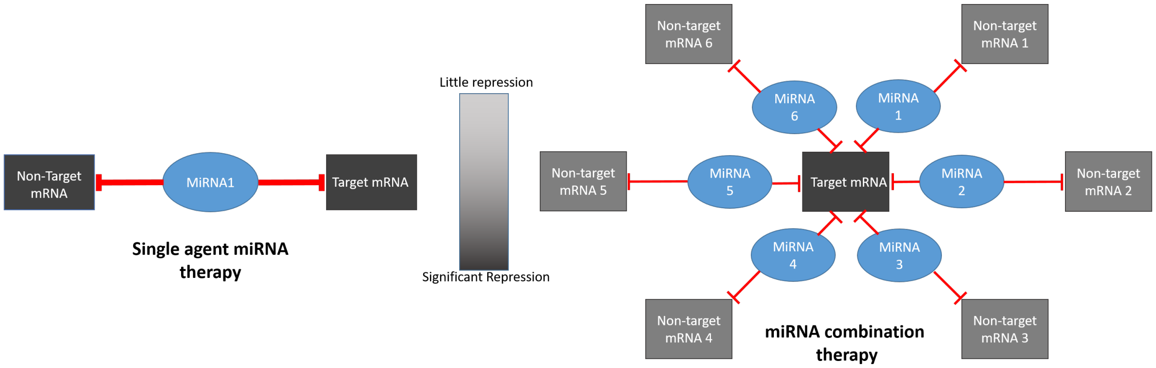

Figure 2.1: Cartoon describing rationale for focusing on miRNA combination therapy. With single-agent therapy, both target mRNA and non-target mRNA are inhibited an equal amount, potentially resulting in toxicity due to off-target effects. With miRNA combination therapy, the common target mRNA is inhibited to a greater degree than any individual non-target miRNA.

and toxicity (Figure 2.1). For example, miR-34 was effective in pre-clinical studies for the treatment of a variety of solid tumors [89–92], only to fail in a phase I clinical trial due to “immune-related serious adverse events” [96]. To capitalize on the promise of miRNA-based cancer therapy while limiting potential toxicity, we developed a systematic, network-based approach to evaluate miRNA cocktails. We focused on miRNA cocktails rather than single miRNA therapeutics due to the potential for miRNA cocktails to minimize toxicity compared to single miRNA regimens [97] (**Figure 2.1**).

In this work, we build on previous studies applying thermodynamic measures to cell signaling networks in the field of cancer biology [74, 98, 99], as well as works that describe a method to use gene homology to map miRNAs to the mRNA transcripts they likely repress [82, 93, 94]. Reitman et al. previously described a metric of cell state analogous to Gibbs free energy that can be calculated using the protein-protein interaction network of human cells and corresponding transcriptomic data [98]. Gibbs free energy has been correlated with a number of cancer-specific outcomes, including cancer grade and patient survival [99]. Additionally, Reitman et al. leveraged Gibbs and other network measures to identify personalized protein targets for therapy in a dataset of low-grade glioma patients from The Cancer Genome Atlas (TCGA) [98]. Previous work has also highlighted the critical importance of miRNAs to maintenance and development of the oncogenic phenotype, and demonstrated the utility of applying miRNA-mRNA mappings. [82] In this work, we developed and applied a computational pipeline that leverages these network principles to identify miRNA cocktails for the treatment of Ewing sarcoma.

2.2 Results

2.2.1 Network Overview

We calculated the network potential, a unitless measure of cell state, for each protein in the cell signaling networks for each of the six Ewing sarcoma cell lines in our experiment. An overview comparing network potential to normalized mRNA expression can be found in **Figure 2.2**. An additional overview of the total network potential for each cell line and biological replicate compared to total mRNA expression is presented in **Figure 2.6**. The histograms of network potential and mRNA expression demonstrate markedly different distributions (**Figure 2.2**), indicating that network potential describes different features of a cell signaling network compared to mRNA expression alone. Notably, network potential and mRNA expression for these cell lines are stable across different biological replicates, as demonstrated by the low interquartile range (**Figure 2.6C** and **2.6D**). There were larger differences in both mean expression and network potential across cell lines (**Figure 2.6C** and **2.6D)** when compared to between-replicate differences. The global average network potential across all samples was -3.4×10^5 with a standard deviation of 1605.

The included patient samples (beginning with SJEWS on **Figure 2.6**) demonstrate substantially more variation in both mRNA expression and network potential. Much of

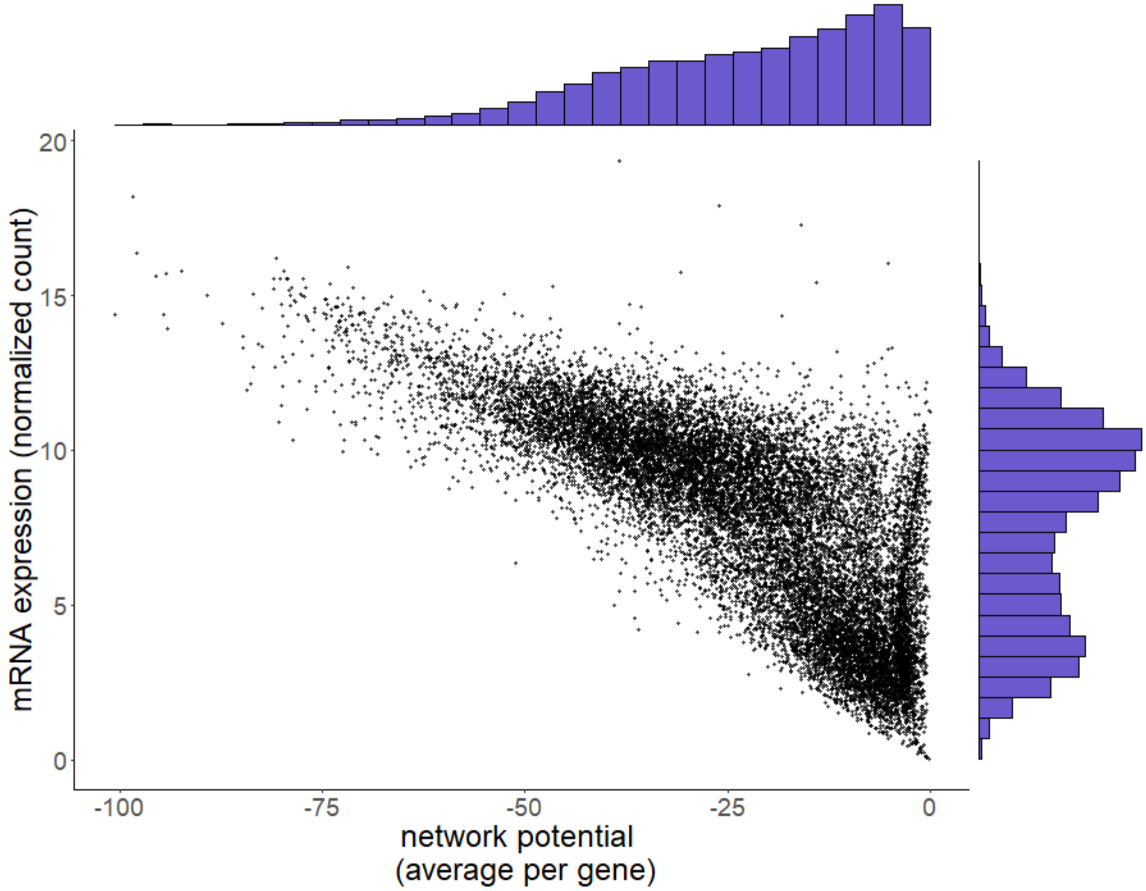

Figure 2.2: Network potential demonstrates a different distribution compared to mRNA expression. Main panel: scatterplot comparing mRNA expression and network potential for all genes in our 18 Ewing Sarcoma cell lines. For each gene, we averaged across all samples for both mRNA expression and network potential. Unlike network potential (top axis histogram), mRNA expression (right axis histogram) has a bimodal distribution.

what we are capturing here can likely be attributed to batch effects, as these patient samples may have been sequenced years apart on machines with dramatically different capabilities.

2.2.2 Identification of Protein Targets

We identified TRIM25, APP, ELAV1, RNF4, and XPO1 as top 5 targets for therapy for each of the 6 cell lines based on the degree of network disruption induced following *in silico* repression of each gene. Of these 5 genes, only XP01 has been previously implicated in oncogenesis [1], while only ELAVL1 is a known essential housekeeping gene. There was a high degree of concordance between cell lines among the top predicted targets (**Table 2.1**). Of the top ten predicted targets, all 10 targets are conserved for all 6 cell lines. The top 50 protein targets are presented in **Figure 2.3**. Many of the top identified genes fell within the 99.99% confidence interval of the computed null distribution (Figure 2.3), suggesting that these genes are highly connected hub genes that are likely to score high in ΔG regardless of the tumor-specific RNA-sequencing information provided. As a sensitivity analysis, we repeated our protein identification pipeline using stringdb as the PPI network provider rather than biogrid. Using stringdb, we identified broadly similar targets, with the majority of identified genes being either essential housekeeping genes or genes causally implicated in cancer. 13 out of the 50 identified targets were shared between stringdb and biogrid. In contrast to our main results using biogrid, the vast majority of top targets identified with stringdb fell within the 99.99% confidence interval of the computed null distribution (**Figure 2.8**) In addition, the distribution of network potential across all genes looks extremely similar regardless of which PPI was used (**Figures 2.2** and **2.8**).

Surprisingly, there was very little overlap between the top predicted targets between the cell line data and the patient samples (**Figure 2.3**). The genes that appeared most frequently among the top projected targets for the 15 patient tumor samples were SLC24A1, ARTN, DHRSX, TEX261, and FRMD8. Compared to the cell line samples, fewer of the most frequent projected targets were cancer-associated or defined housekeeping genes (**Figure 2.3**)

Some of these identified genes in the cell line data are likely essential housekeeping genes highly expressed in all or most cells in the body, making them inappropriate drug targets (**Figure 2.3**). TRIM25, and ELAV1, for example, are involved in protein modification and RNA binding, respectively [100]. We therefore repeated this analysis, limiting our search to gene targets that have been causally implicated in cancer [1]. With this limitation in place, we identified XPO1, LMNA, EWSR1, HSP90AA1, and CUL3 as the top 5 targets for therapy when $\mathit{\Delta G}$ was averaged for all cell lines. The top 10 cancer-related targets for each cell line can be found in (**Table 2.2**). The top 50 protein targets (limited to those causally implicated in cancer) can be found in **Figure 2.7**.

We also conducted gene set enrichment analysis for the all the genes represented in our cell signaling network (averaged across all samples). We ranked genes by network potential (averaged across all samples) and compared our gene set to the "hallmarks" pathways set, downloaded from the Molecular Signatures Database (MSigDB) [101, 102]. This analysis was conducted using the fGSEA package in R, which uses the Benjamini - Hochberg procedure to correct the false discovery rate [103,104]. Our gene set was enriched (adjusted p-value < 0.05) in 24 of the 50 pathways included in the hallmarks set; including apoptosis, DNA repair, mTOR signaling, MYC signaling, and WNT β-catenin signaling. Our gene set was also highly enriched (normalized enrichment score $= 1.73$) in the miRNA bio-genesis pathway. The full results are presented in **Table 2.3**.

2.2.3 Identification of miRNA Cocktails

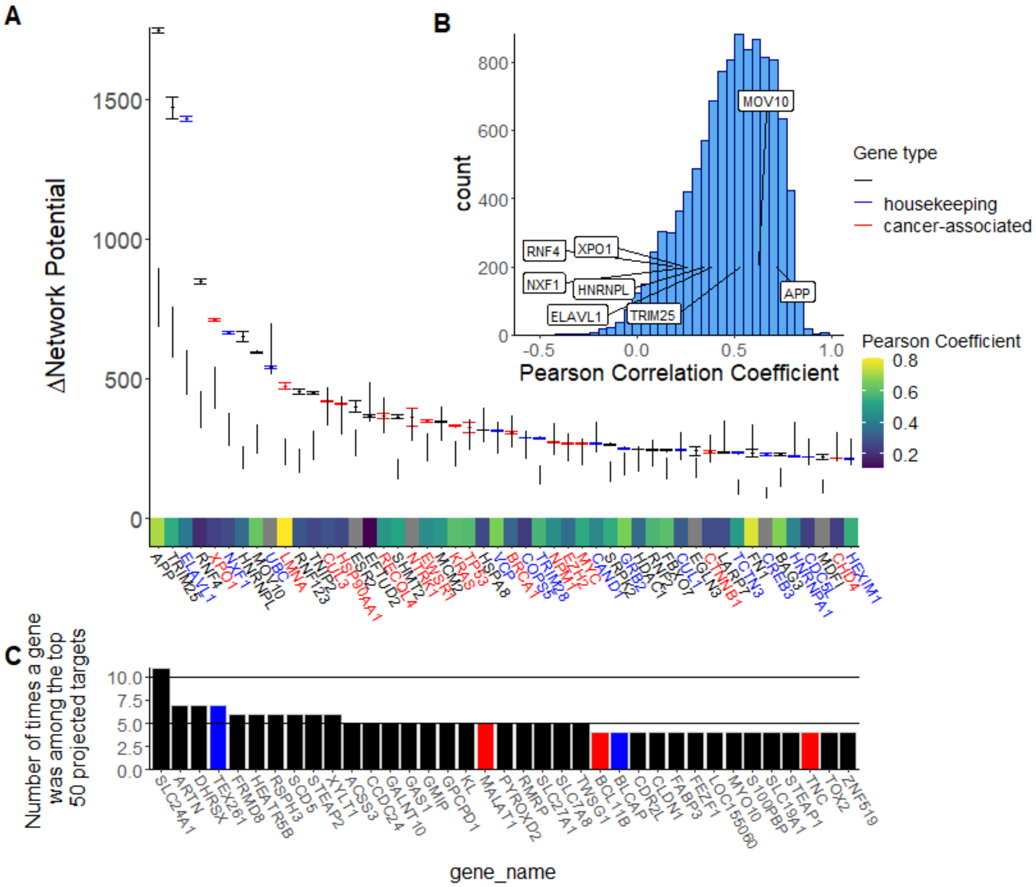

Figure 2.3: TRIM25, APP, ELAVL1, AND RNF4, and XPO1 are the top protein targets ranked by predicted disruption following *in silico* repression. Panel A: Box and whisker plot describing the change in network potential following *in silico* repression for each of the top 50 proteins. 99.99% confidence interval from the permutation test are displayed alongside the box and whisker plots. It is notable that EWSR1, the kinase associated with Ewing sarcoma development, is considered highly influential in the cell signaling network by this method, even in comparison to the computed null distribution. Genes that have previously been causally implicated in cancer according to the Cosmic database are highlighted in red [1]. Essential housekeeping genes (excluding those that are causally implicated in cancer) are highlighted in blue. The heat-map on the x-axis corresponds to the protein-mRNA correlation of each gene in the Cancer Cell Line Encyclopedia [2]. **Panel B:** Histogram depicting the distribution of Pearson correlation between mRNA expression and protein expression from the Cancer Cell Line Encyclopedia for all nodes included in our final Ewing sarcoma cell signaling networks. Proteins that were ranked particularly highly in panel A were labeled in panel B. **Pancel C:** Bar chart describing the most frequently observed genes among the top projected targets for the 15 patient tumor samples we analyzed. There was very little overlap in top projected targets between the cell line and patient data, reflecting the transcriptional heterogeneity present in Ewing Sarcoma.

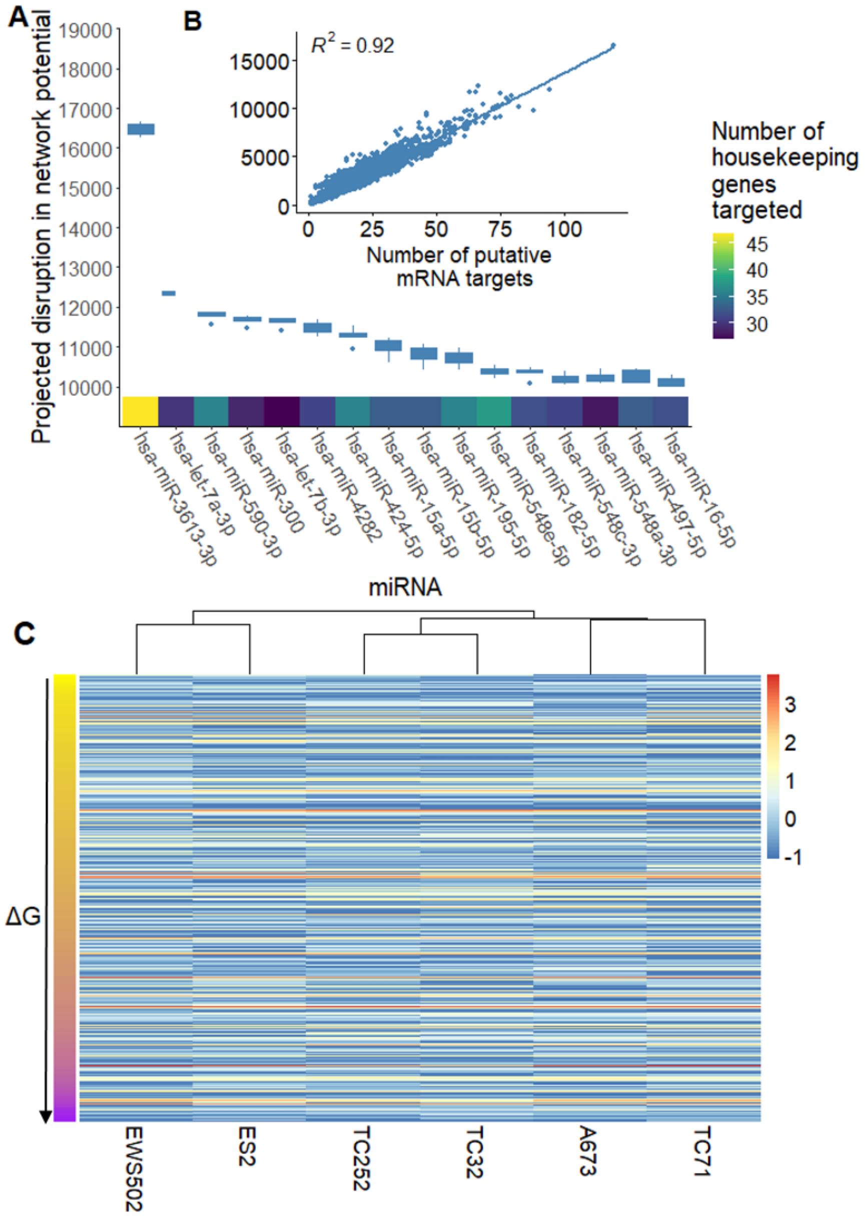

Figure 2.4: Many of the most promising miRNA candidates repress large numbers of essential housekeeping genes. We identified the top miRNA for treatment of Ewing sarcoma, ranked by their predicted disruption of the Ewing sarcoma cell signaling network. **A:** Boxplot showing the projected disruption in network potential for the top miRNA candidates (averaged across all samples). The heatmap on the x-axis describes the number of essential housekeeping genes that each miRNA is predicted to target. **B:** Scatterplot showing the relationship between projected network disruption and the number of putative mRNA targets for a given miRNA. **C:** Heatmap showing z-score normalized miRNA expression for 622 of the evaluated miRNA for the 6 cell lines under study. The Y axis is clustered by the projected ΔG associated with a given miRNA. There doesn't seem to be a clear pattern of miRNA expression based on projected ΔG.

We identified several miRNAs that were predicted to dramatically disrupt the Ewing sarcoma cell signaling network (**Figure 2.4**). When averaging all cell lines, we identified miR-3613-3p, let-7a-3p, miR-300, miR-424-5p, and let-7b-3p as the ideal miRs for preferential repression of proteins predicted to be important for Ewing sarcoma signaling network stability. miR-3613-3p, let-7a-3p, miR-300, miR-424-5p, and let-7b-3p were predicted to cause an average network network potential increase (driving the system less negative) of 17382, 13034, 12746, 12364 and 12280, respectively (see **Figure 2.4**). It should also be noted that we were able to identify a substantial number of miRNAs with potential activity against the Ewing sarcoma cell signaling network. We identified 27 miRNAs with an average predicted network potential disruption of greater than $10,000$. For comparison, the largest network change in network potential that could be achieved with a single gene repression across all cell lines was just 2064 (TRIM25). miRNA sequencing of the 6 cell lines under study did not reveal any clear pattern of miRNA expression based on predicted network potential disruption.

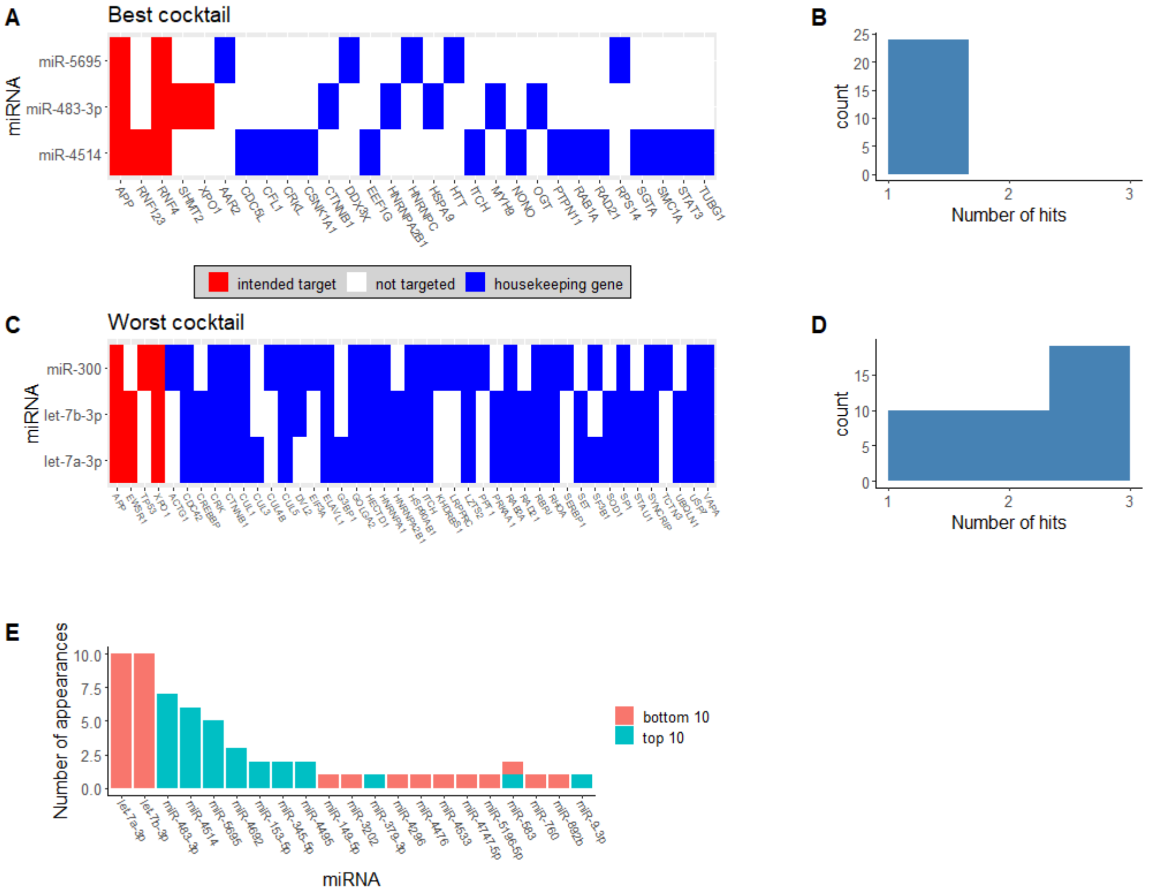

Figure 2.5: We identified miR-483-3p, miR-5695, and miR-4514s as the optimal 3-miRNA cocktail for Ewing Sarcoma therapy. We identified cocktails that are predicted to maximally downregulate target genes (red shading on the figure), while avoiding downregulation of essential housekeeping genes to limit toxicity (blue shading on the figure). **Panel A.** shows the targeting heatmap for the best predicted cocktail for cell line A673. The miRNA that make up the cocktail are presented on the y-axis. Putative gene targets are highlighted on the x-axis. Lines that span multiple miRNAs occur when a gene is downregulated by 2 or more miRNAs in the cocktail. **Panel B.** shows a histogram of the number of microRNA that target a given housekeeping gene in the best cocktail. **Panel C.** displays the targeting heatmap for the worst-performing cocktail for cell line A673 among those tested for reference. **Panel D.** shows a histogram of the number of microRNA that target a given housekeeping gene in the worst predicted cocktail. **Panel E** shows a bar graph showing the miRNA that most frequently appear in either the bottom or top 10 predicted cocktails (averaged across cell lines) for Ewing Sarcoma therapy.

These individual miRNAs target large numbers of transcripts in the cell and therefore may be difficult to administer as single-agents due to extreme toxicity. For example, the top miR candidate, miR-3613-3p, was predicted to repress 144 distinct mRNA transcripts in the full target set. We therefore sought to identify cocktails of miRNA that could cooperatively down-regulate key non-housekeeping genes while avoiding cooperative down-regulation of housekeeping genes that may be associated with toxicity. When targeting the top 10 predicted proteins from our *in silico* repression experiments, a 3 miRNA cocktail of miR-483-3p, miR-379-3p, and miR-345-5p was predicted to be the most optimal across all cell lines (**Figure 2.5A** and **Figure 2.5B**). Under the same conditions, a 3-miR cocktail of miR-300, let-7b-3p, and let-7a-3p was predicted to be the least optimal among 16,215 tested combinations (**Figure 2.5C** and **Figure 2.5D**). Notably, the most and least optimal miRNA combinations had similar activity against the 10 targets (**Figure 2.5A** and **Figure 2.5C**). The worst cocktail was defined by high levels of cooperative downregulation of housekeeping genes rather than lack of efficacy against putative targets (**Figure 2.5C** and **Figure 2.5D**). Let-7b-3p and let-7a-3p were heavily represented in the least optimal cocktails tested, appearing in 10 of the 10 worst 3 miRNA cocktails (**Figure 2.5E**). These highly promiscuous miRNA target large numbers of housekeeping genes, limiting their therapeutic utility alone or in combination **(Figure 2.4B**).

Notably, many of the most promising miRNA when considering only their total predicted network disruption tend to appear in the least optimal cocktails (**Figure 2.4**). This likely occurs because these miRNA tend to target large numbers of housekeeping genes and large numbers of genes overall. In contrast, the best miRNA cocktails tend to be composed of miRNA that target relatively few genes overall but exhibit some degree of target specificity. Put another way, they target the desired target genes while repressing relatively few essential housekeeping genes. An extreme example of this is the case of miR-483-3p. MiR-483-3p is in the bottom 50% of all miRNA when ranked by predicted network disruption, and is only predicted to repress 10 different transcripts. However, because it selectively targets several of our targets of interest, this relative small total projected network disruption is actually an attractive feature that makes it easy to build effective cocktails that include miR-483-3p. As a result, miR-483-3p appears in 7 of the top 10 predicted 3 miRNA cocktails. To assess the stability of our results, we repeated this analysis, focusing on the top 5 or top 15 predicted protein targets. We also repeated this analysis, assuming 10 % and 50% repression per miRNA that target a given mRNA. The top and bottom predicted cocktails were similar across these conditions and across all six cell lines. We have included the full ranked list of all miRNA cocktails tested across all conditions on Github

2.2.4 Additional Analyses

First, we provide some additional summary plots describing mRNA expression and network potential in our samples.

As described in the main text, we ranked proteins according to their contribution to network stability by calculating the change in network potential following complete *in silico* repression of each protein. In the main text, we limited our analysis to proteins that had been causally implicated in cancer according to the cosmic database [1]. Here, we present the top 50 proteins (when network potential for all 6 cell lines was averaged) ranked by contribution to network stability, not limited to proteins that were causally implicated in cancer (Figure 2.3).

We also analyzed each cell line individually to identify the top protein targets for each cell line. In the main text, we limited this analysis to proteins that had been causally implicated in cancer [1]. Here, we present the top protein targets for each cell line, not

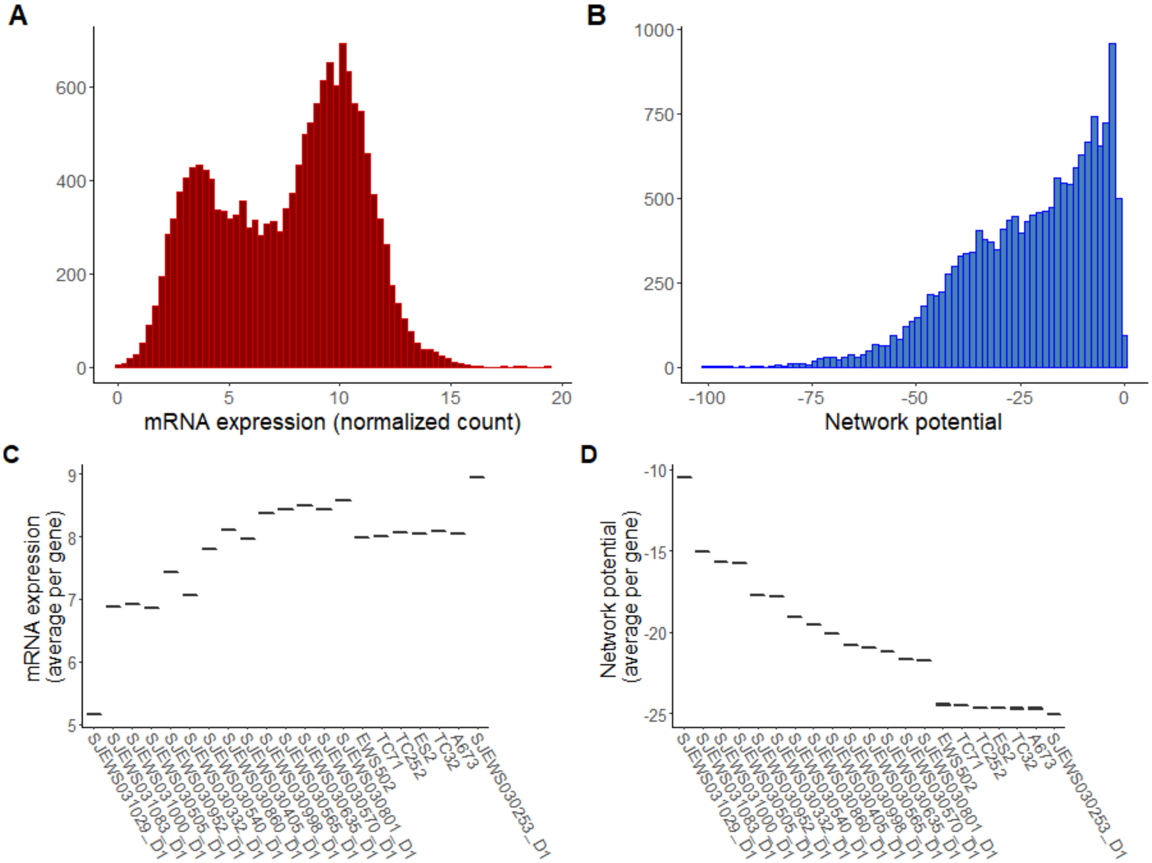

Figure 2.6: Network potential describes different features of a cell signaling network compared to mRNA expression alone. Panel A: Histogram of mRNA expression for each gene (averaged across all samples). **Panel B:** Histogram of the network potential for each gene (averaged across all samples) mRNA transcripts with an expression level of zero were excluded from both histograms to better visualize the distribution of genes that are expressed. **Panel C:** Box plot showing the total mRNA expression for each cell line and patient sample (patient samples begin with SJEWS). **Panel D:** Box plot showing the total network potential for each cell line and patient sample.

limited to those proteins that had previously been causally implicated in cancer (Table 2.1).

To evaluate the impact of PPI choice on our findings, we repeated the protein target selection portion of our analysis using an entirely different PPI, stringdb (**Figure 2.8**). To ensure a similar number of total edges to the biogrid database, we modulated the provided "interaction score" until the resulting stringdb derived network contained about 400000 edges. We ultimately used an interaction score cutoff of 700 (indicating a high confidence

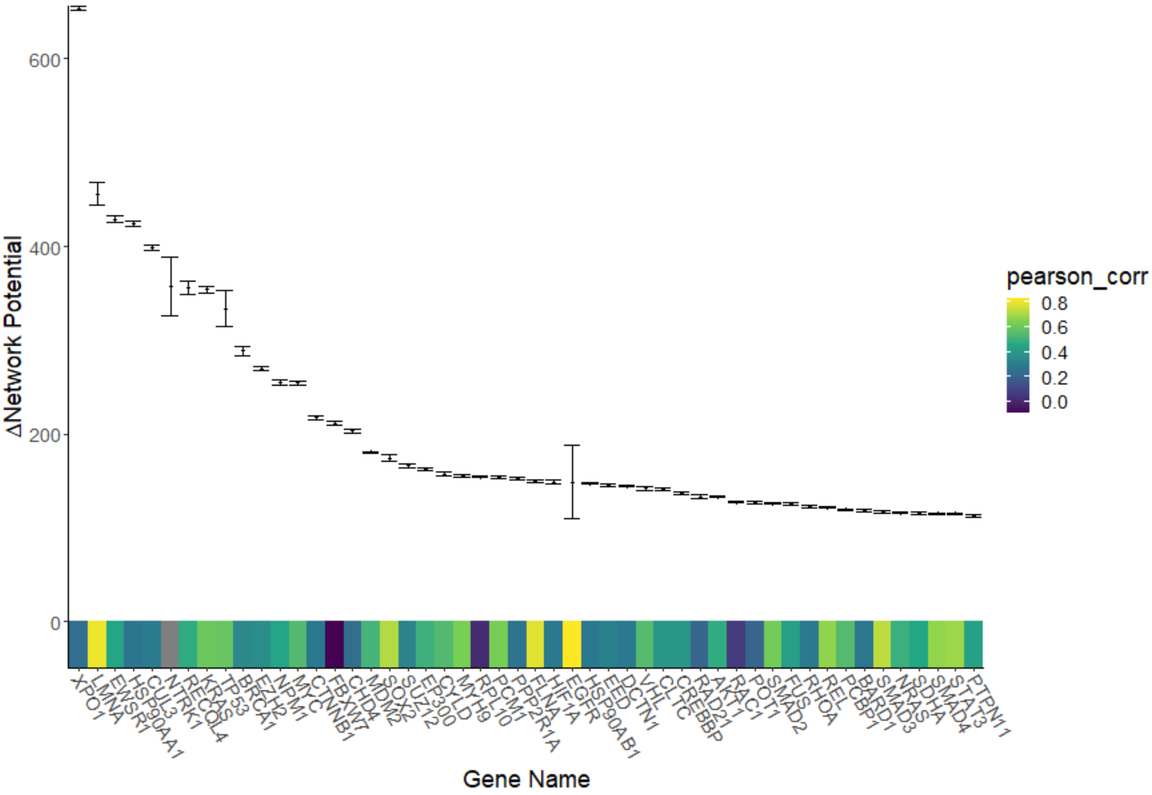

Figure 2.7: Protein targets ranked by contribution to network stability. When averaging across cell lines, XPO1, LMNA, EWSR1, HSP90AA1, and CUL3 were identified as the most important proteins in the Ewing sarcoma cell signaling network (when limiting our analysis to proteins causally implicated in cancer [1]). When each protein was simulated as completely repressed *in silico*, network potential was increased by 654, 456, 429, 425, and 399, respectively. The heatmap at the bottom of the plot describes the protein-mRNA correlation for each gene in the cancer cell line atlas. Grey indicates no data was available. It is reassuring that EWSR1, the kinase associated with Ewing sarcoma development, is identified as highly influential in the cell signaling network by this method.

in the observed interaction).

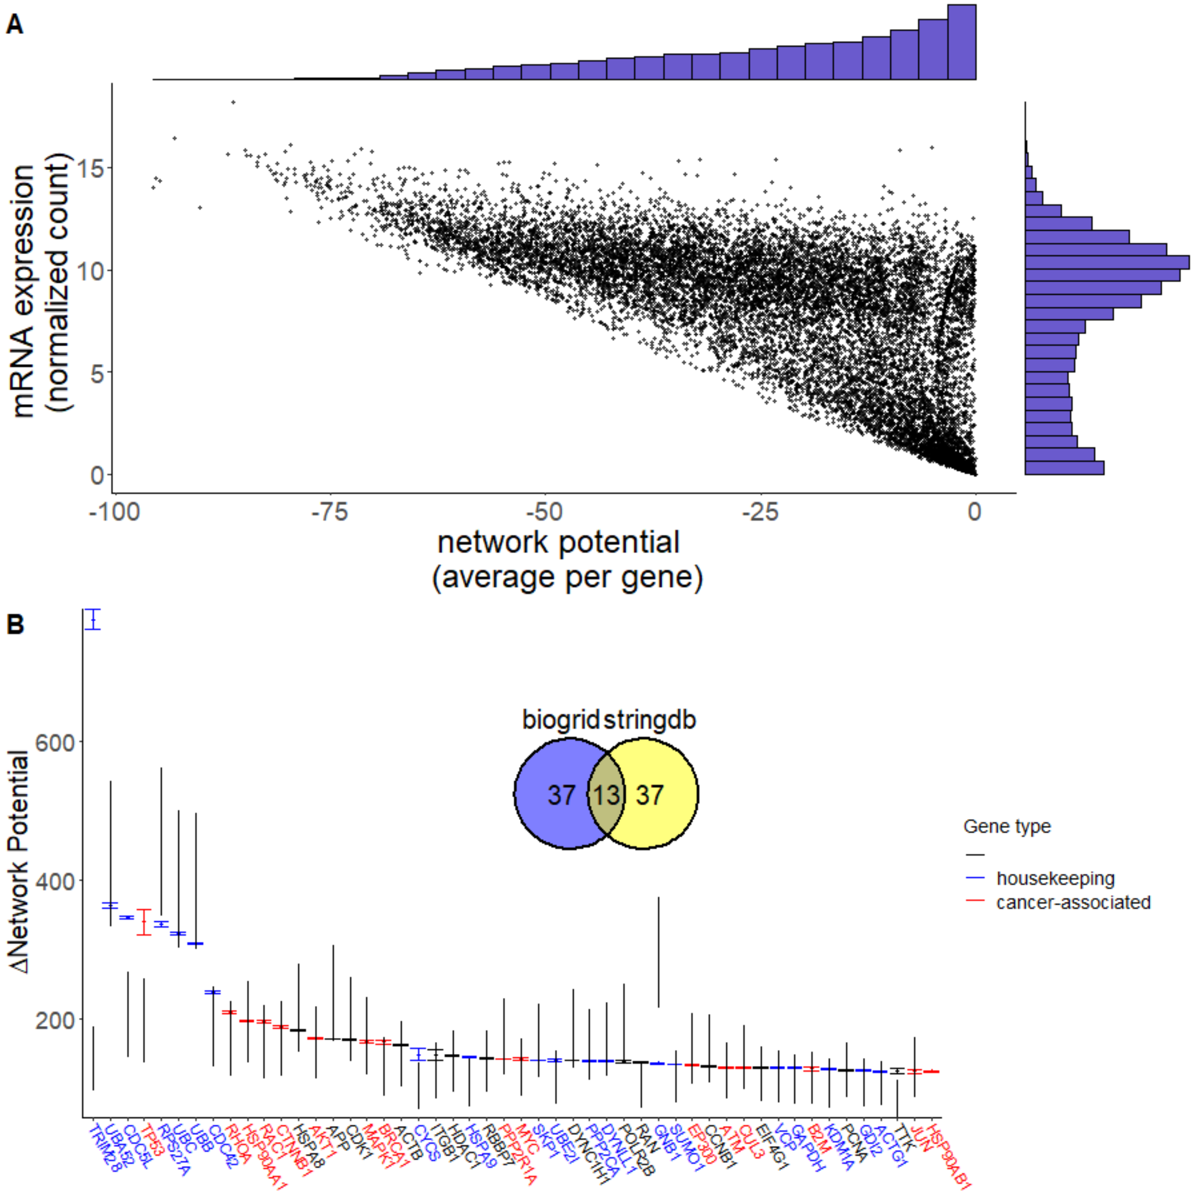

Figure 2.8: Overview of our findings, using the stringdb protein-protein interaction network. Panel A: Scatterplot with marginal histograms comparing mRNA expression to network potential. **Panel B:** Box and whisker plot showing the change in network potential for the top 50 genes, as well as 99.99% confidence intervals from the permutation test. We also show a histogram comparing the top 50 genes identified by our pipeline using stringdb compared to biogrid as the protein-protein interaction network.

	TC252	ES2	A673	TC32	EWS502	TC71
1	TRIM25	TRIM25	TRIM25	TRIM25	TRIM25	TRIM25
2	APP	APP	APP	APP	APP	APP
3	ELAVL1	ELAVL1	ELAVL1	ELAVL1	ELAVL1	ELAVL1
4	RNF4	RNF4	RNF4	RNF4	RNF4	RNF4
5	HNRNPL	HNRNPL	HNRNPL	HNRNPL	HNRNPL	HNRNPL
6	XPO1	XPO1	XPO1	XPO1	XPO1	XPO1
7	NXF1	NXF1	NXF1	NXF1	NXF1	NXF1
8	UBC	TNIP2	UBC	UBC	UBC	UBC
9	TNIP2	UBC	TNIP2	TNIP2	TNIP2	TNIP2
10	MOV10	MOV10	MOV10	MOV10	MOV10	MOV10

Table 2.1: Top protein targets for each cell line. We ranked potential targets by predicted change in network potential when each protein was modeled as repressed.

	TC252	ES2	A673	TC32	EWS502	TC71
1	**XPO1**	**XPO1**	**XPO1**	**XPO1**	**XPO1**	**XPO1**
2	**LMNA**	**LMNA**	**LMNA**	**LMNA**	NTRK1	EWSR1
3	EWSR1	**HSP90AA1**	**HSP90AA1**	EWSR1	**HSP90AA1**	LMNA
4	**HSP90AA1**	**EWSR1**	**EWSR1**	**HSP90AA1**	**EWSR1**	**HSP90AA1**
5	**CUL3**	**CUL3**	**CUL3**	**CUL3**	LMNA	**CUL3**
6	**NTRK1**	KRAS	**NTRK1**	**NTRK1**	CUL3	RECQL4
7	**TP53**	**TP53**	RECQL4	**TP53**	RECQL4	KRAS
8	RECQL4	RECQL4	KRAS	KRAS	TP53	NTRK1
9	KRAS	EGFR	TP53	RECQL4	KRAS	BRCA1
10	**BRCA1**	**BRCA1**	**BRCA1**	**BRCA1**	**BRCA1**	TP53

Table 2.2: Top cancer-associated protein targets for each cell line. We ranked potential targets by predicted change in network potential when each protein was modeled as repressed, limited to proteins causally associated in cancer according to the Cosmic database. Proteins that appear in the same position for ≥ 3 cell lines are **bolded**.

Gene set enrichment analysis We conducted gene set enrichment analysis, using all genes in our Ewing sarcoma cell signaling network as the gene set. We ranked this set of genes by change in network potential and used the "hallmarks" pathways set from the Molecular Signatures Database as the genomic background [101, 102]. We also included a gene set corresponding to the miRNA biogenesis pathway [84]. We used the "fgsea" R package version 1.8.0 to conduct the analysis using the following settings: nperm $= 500$, minSize $= 1$, maxSize $= \infty$, nproc $= 0$, gseaParam $= 1$, BPPARAM =NULL [103]. We found our gene set to be significantly enriched in several pathways related to oncogenesis, including DNA repair, apoptosis, and MTOR signaling (Table 2.3). These results indicate that network potential can identify a cancer-specific signal from mRNA expression data.

	pathway	pval	padj	ES	NES	nMoreExtreme	size
1	MITOTIC_SPINDLE	0.00	0.01	0.62	1.47	0.00	197
2	DNA_REPAIR	0.00	0.01	0.59	1.37	0.00	146
3	G2M_CHECKPOINT	0.00	0.01	0.72	1.68	0.00	187
4	APOPTOSIS	0.00	0.01	0.62	1.44	0.00	158
5	PROTEIN_SECRETION	0.00	0.01	0.63	1.44	0.00	94
6	APICAL_SURFACE	0.00	0.01	0.73	1.57	0.00	42
7	UNFOLDED_PROTEIN_RESPONSE	0.00	0.01	0.62	1.43	0.00	106
8	PI3K_AKT_MTOR_SIGNALING	0.00	0.01	0.69	1.58	0.00	104
9	MTORC1_SIGNALING	0.00	0.01	0.61	1.43	0.00	193
10	E2F_TARGETS	0.00	0.01	0.72	1.69	0.00	195
11	MYC_TARGETS_V1	0.00	0.01	0.80	1.89	0.00	193
12	OXIDATIVE_PHOSPHORYLATION	0.00	0.01	0.61	1.42	0.00	184
13	ALLOGRAFT_REJECTION	0.00	0.01	0.54	1.27	0.00	191
14	MIRNA_BIOGENESIS	0.00	0.01	0.80	1.73	0.00	40
15	WNT_BETA_CATENIN_SIGNALING	0.01	0.02	0.68	1.47	2.00	42
16	ANGIOGENESIS	0.01	0.02	0.72	1.53	2.00	34
17	TGF_BETA_SIGNALING	0.01	0.03	0.65	1.43	4.00	53
18	MYC_TARGETS_V2	0.01	0.03	0.64	1.41	4.00	58
19	P53_PATHWAY	0.01	0.03	0.51	1.19	5.00	194
20	UV_RESPONSE_UP	0.01	0.04	0.53	1.23	6.00	153
21	SPERMATOGENESIS	0.02	0.04	0.54	1.24	7.00	126
22	ADIPOGENESIS	0.02	0.04	0.50	1.18	8.00	191
23	INTERFERON_GAMMA_RESPONSE	0.02	0.04	0.50	1.18	8.00	193
24	INTERFERON_ALPHA_RESPONSE	0.02	0.05	0.56	1.28	10.00	91

Table 2.3: Genes ranked by network potential are enriched for several biological pathways related to cancer as well as the miRNA bio-genesis pathway Pathways with an adjusted p-value < 0.05 are shown above. “ES” refers to enrichment score and “NES” refers to the normalized enrichment score. “nMoreExtreme” refers to the number of random gene sets (out of 500) that were more enriched than the test set. Size refers to the number of genes in the pathway that were also present in our mRNA expression dataset.

2.3 Discussion

In this work, we described a novel methodology for the identification of potential miRNA cocktails for Ewing sarcoma therapy. First, we performed paired miRNA and mRNA sequencing on six Ewing sarcoma cell lines (GEO accession GSE98787). We then defined a metric of cell state, network potential, based on mRNA expression and signaling network topology. Using *in silico* repression and change in network potential, we identified the most important proteins in the cell signaling network for each of the 6 cell lines. We then repeated this process for 15 patient tumor samples derived from the St. Jude Cloud [105]. Notably, this set of proteins was enriched in 24 of the 50 pathways included in the "halmarks" gene set [101, 102]. The ranked protein set was also enriched for genes involved in the canonical miRNA biogenesis pathway [82]. We then evaluated more than 16000 3-miRNA cocktails (per cell line) based on predicted ability to disrupt key proteins in the Ewing Sarcoma cell signaling network while avoiding cooperative down-regulation of essential housekeeping genes. We ranked these 3-miRNA cocktails to identify promising miRNA combinations for therapy of Ewing Sarcoma.

The protein targets and miRNA candidates we identified in our dataset are consistent with the literature on Ewing sarcoma and cancer cell signaling, suggesting biological plausibility of our methodology. Of the top 50 protein targets that we identified, 15 were previously causally implicated in cancer [1], including EWSR1, the proposed driver of Ewing sarcoma development. In addition, our network-based approach suggests that known oncogenic hub genes such as KRAS and MYC are prime targets for disruption in cancer cells. We also identified a number of plausible targets that were not previously implicated in cancer, such as MOV10. MOV10 is an RNA helicase involved in the RNA-induced silencing complex (RISC), a key complex involved in epigenetic signaling by miRNA [106]. As mentioned previously, our findings suggest that the miRNA biogenesis pathway is enriched in the setting of Ewing Sarcoma. The central role of MOV10 in the EWS cell signaling network provides further evidence for the importance of miRNA signaling in EWS oncogenesis.

Many of the miRNA we identified as potential therapeutic candidates have been previously studied due to their association with cancer outcomes, including members of the let-7 family, miR-300, miR-424-5p, miR-4282, miR-15a-5p, and miR-590-3p. Loss of expression of the let-7 family of miRNA has been widely implicated in cancer development [107–110]. In Ewing sarcoma specifically, low levels of let-7 family miRNA have been correlated with disease progression or recurrence [107]. The let-7 family of miRNA have also been studied as treatment for non-small cell lung cancer in the preclinical setting [91]. Loss of miR-300 has been previously correlated with development and aggressiveness of hepatocellular carcinoma [111] as well as in oncogenesis of pituitary tumors [112]. Reduced expression of miR-424-5p and miR-4282 have each been implicated in the development of basal-like breast cancer [113, 114]. MiR-15a-5p has been shown to have anti-melanoma activity [115]. In addition, miR-590-3p has been show to suppress proliferation of both breast cancer [116], and hepatocellular carcinoma [117]. The broad literature linking many of our proposed miRNA candidates for Ewing sarcoma treatment to the development and maintenance of cancer highlights the ability of our computational pipeline to identify potentially promising therapeutic candidates in this setting. Prior to application of these findings for treatment of Ewing sarcoma or any other disease, specific *in vitro* and *in vivo* validation is needed.

The process by which putative miRNA targets were selected was based on sequence homology rather than direct experimental validation. As a result, it is possible that we included false positive miRNA targets in our analysis. For this study we relied on a protein-

protein interaction network presumably curated from analyzing normal human cells. It is possible that the derangements observed in cancer cells could change the underlying interaction network of a tumor cell. In the future, it may be possible to utilize protein-protein interaction networks specific to cancer or even specific to the cancer type under study. In addition, we did not consider specific binding sites that these miRNA may use to repress target mRNA. Certain miRNA may share binding sites on their target mRNA (i.e. the let-7 family of miRNA), which may make our assumption of linear additive miRNA effects invalid. We also used mRNA concentration as a surrogate for protein concentration in designing our cell signaling network. While this is not true in all cases, it is likely a reasonable approximation under steady state conditions [118–121] (see Section 2.4.2 for more details). In addition, protein-mRNA correlations in the cancer cell line atlas for the top proteins identified by our pipeline were fairly good, ranging from 0.07 to 0.8 for the top 50 identified protein targets. [2] (**Fig 2.3**).

Despite these limitations, our findings may facilitate the development of novel therapies for patients suffering from Ewing Sarcoma. To this point, severe toxicity has limited the translation of miRNA-based cancer therapies to the clinical setting. Our pipeline may enable the development of better miRNA therapies that clear this hurdle and open up this promising avenue of therapy for patients suffering from cancer. In addition, this novel method can facilitate the rapid identification of key proteins in any cancer cell signaling network for which mRNA sequencing data is available. This may facilitate more rapid drug discovery and assist in the discovery of proteins and miRNA that play a significant role in the cancer disease process.

2.4 Methods

2.4.1 Overview

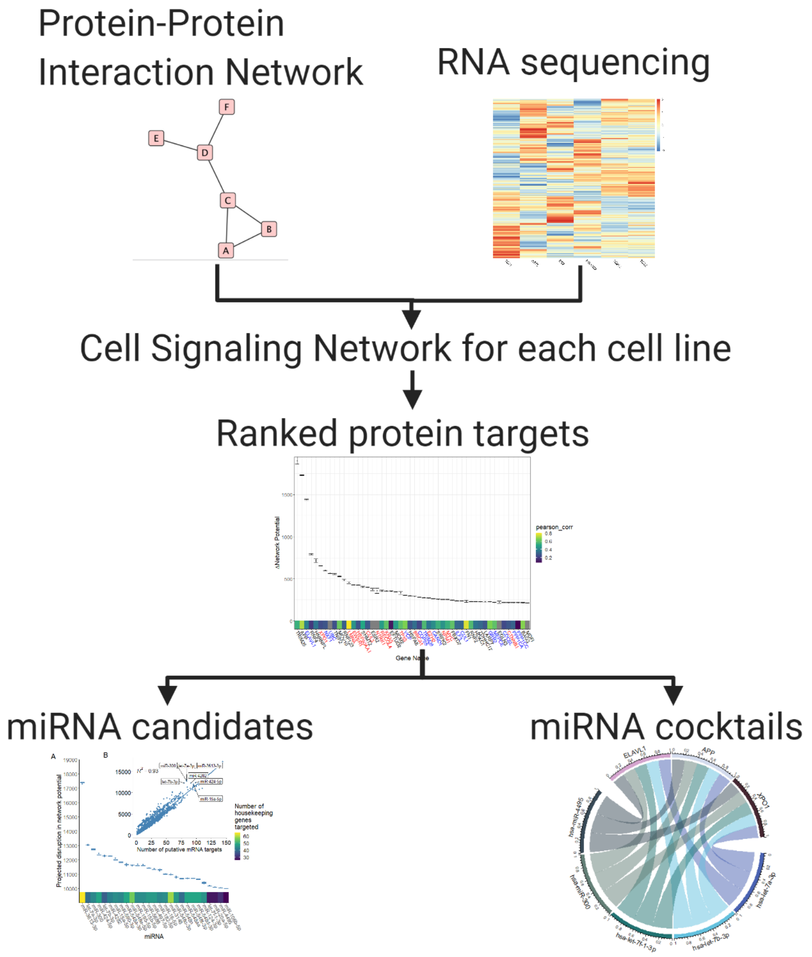

Figure 2.9: Simplified schematic of our computational pipeline. We defined a measure of tumor state, which we term network potential (Equation 2.1), based on both mRNA gene expression and the underlying protein-protein interaction (PPI) network. Next, we ranked mRNA targets based on their contribution to network potential of each cell line, aiming to approximate the relative importance of each mRNA to network stability. After identifying these mRNA targets, we then identified miRNA and miRNA cocktails that preferentially acted to repress the most influential of the ranked mRNA targets, with the aim of defining synthetic miRNA-based therapy for down-regulation of these targets.

We characterized six previously described Ewing sarcoma cell lines in triplicate [122] – A673, ES2, EWS502, TC252, TC32, and TC71 – using paired miRNA and mRNA sequencing. By evaluating 6 distinct cell lines, we aimed to assess the heterogeneity inherent to Ewing sarcoma *in-vitro*. We also utilized mRNA sequencing data for 15 ewing sarcoma patient tumor samples made available on the St. Jude Cloud [105]. We then defined a measure of tumor state, which we term network potential (**Equation 2.1**), based on both mRNA gene expression and the underlying protein-protein interaction (PPI) network. Next, we ranked mRNA targets based on their contribution to network potential of each cell line, aiming to approximate the relative importance of each mRNA to network stability. Relative importance of each mRNA to network stability was determined by calculating the change in network potential of each network before and after *in silico* repression of each mRNA (ΔG, described in **Section 2.4.5**). After identifying these mRNA targets, we then identified miRNA and miRNA cocktails that preferentially acted to repress the most influential of the ranked mRNA targets, with the aim of defining synthetic miRNA-based therapy for down-regulation of these targets. Our computational pipeline is schematized in **Figure 2.9**.

2.4.2 Data sources

We utilized three data sources to develop our Ewing sarcoma cell signaling networks: the BioGRID protein-protein interaction database [37], mRNA expression data from 6 Ewing sarcoma cell lines, which are available on GEO (accession GSE98787), and mRNA expression data from 15 Ewing sarcoma patient samples, which are available on the St. Jude Cloud [105].

Protein Protein Interaction Databases The BioGRID interaction database contains curated data detailing known interactions between proteins for a variety of different species, including Homo sapiens. The data were generated by manual curation of the biomedical literature to identify documented interactions between proteins [37]. To assist in manual curation, the BioGRID project uses a natural language processing algorithm that analyzes the scientific literature to identify manuscripts likely to contain information about novel PPIs. The dataset is therefore limited to protein interactions that are reliably reported in the scientific literature. As new research accumulates, substantial changes to the PPI network may occur. For example, between 2016 and 2018, the number of documented PPIs in *Homo sapiens* grew from 365,547 to 449,842. The 449,842 documented interactions in 2018 were identified through curation of 27,631 publications [37]. Importantly, the PPI network is designed to represent normal human tissue. To assess the importance of the specific PPI used to our results, we repeated much of our analysis using stringdb, another publicly available PPI with millions of documented interactions. To maintain some consistency with biogrid, we modulated the provided "interaction score" until the resulting network had around the same number of edges as the biogrid network. Interaction score is a number ranging from 0 to 1000 that describes the likelihood that a given protein protein interaction is biologically relevant. For our analysis, we included only interactions with an interaction score ≥ 700.

Ewing sarcoma transcriptomics Second, we utilized mRNA expression data from *in vitro* experiments conducted on six Ewing sarcoma cell lines (3 biological replicates per cell line). RNA/miRNA extraction was performed with a Qiagen kit with on-column DNase digestion. These mRNA and miRNA expression data were then normalized to account for between sample differences in data processing and further adjusted using a regularized log (Rlog) transformation [123, 124]. In order to extend our study to patient samples, we repeated our analysis on 15 patient tumors from the St. Jude Cloud for which RNA sequencing data was available. The St. Jude Cloud is a comprehensive, cloud-based data-sharing ecosystem that provides genomic data on thousands of samples from patients with pediatric cancer [105].

Notably, methods for calculating network potential from this type of data require protein concentrations rather than mRNA transcript concentrations. For the purposes of this analysis, we assumed that concentration of protein in an Ewing sarcoma tumor was equivalent to the concentration of the relevant mRNA transcript. A large body of work suggests that mRNA levels are the primary driver of protein levels in a cell under steady state conditions (i.e. not undergoing proliferation, response to stress, differentiation etc) [118–121]. However, recent work in a 375 cancer cell lines has shown that mRNA expression may not be predictive of protein expression in the setting of malignancy [2]. For this reason, we included the protein-mRNA correlations from their experiments alongside some of our key findings to provide needed context.

2.4.3 Network development

We first developed a generic network to represent human cell signaling networks using the BioGRID interaction database [37]. The BioGRID protein-protein interaction network can be downloaded as a non-linear data structure containing ordered pairs of proteins and all the other proteins with which they interact. This data structure can be represented as an undirected graph, with vertex set $\mathcal{V}$, where each vertex represents a protein, and edge set ($\mathcal{E}$) describes the interactions between proteins.

Using mRNA sequencing data from 6 Ewing sarcoma cell lines in triplicate, we then ascribed mRNA transcript concentration for each gene as an attribute to represent the protein concentration for each node in the graph. Through this process, we developed networks specific to each cell line and replicate in our study (18 total samples). We then repeated the same process to develop networks specifics to the 15 patient tumor samples.

2.4.4 Network potential calculation

Using the cell signaling network with attached cell line and replicate number specific normalized mRNA expression data, we defined a measure of tumor state following Reitman et al. [98], which we term network potential. Our Network potential metric was inspired by Gibbs free energy in physics or chemistry. We first calculate the network potential of the i-th node in the graph:

$$G_i = C_i \ln \left[\frac{C_i}{\sum C_j + C_i} \right]. \tag{2.1}$$

where G_i is equal to the network potential of an individual node of the graph, C_i is equal to the concentration of protein corresponding to node G_i, and C_j is the concentration of protein of the j-th neighbor of G_i. Because of the natural log transformation, G_i will always return a negative number. Total network potential (G) of the network can then be calculated as the sum over all nodes:

$$G = \sum_i G_i. \tag{2.2}$$

where G is equal to the total network potential for each biological replicate of a given cell line. We then compared total network potential across cell lines and biological replicates. More negative network potentials were interpreted as being "larger" in the absolute sense.

2.4.5 Ranking of protein targets

After calculating network potential for each node and the full network, we simulated "repression" of every node in each network by reducing their expression (computationally) to zero, individually [72]. Clinically, this would be akin to the application of a drug that perfectly inhibited the protein/mRNA of interest. Next, we re-calculated network potential for the full network and calculated the change in network potential (ΔG) by subtracting the new network potential value for the network potential value of the "unrepressed" network. Because the network potential of each node was negative, systematic repression of a given node always drove the total network potential to be less negative. As a result, this approach will always return a positive ΔG. We then ranked each node in the network according to ΔG for further analysis.

We also evaluated the top predicted genes by ΔG against a null model of ΔG to evaluate the likelihood that these observed disruptions were due to random chance. To construct our null distribution of ΔG, we repeated the following process 1000 times for each sample under study:

1. We constructed a random graph that preserves the original degree distribution for the underlying protein-protein interaction network by iteratively swapping edges. For each random graph, we performed $n \times 100$ swaps where n is the number of nodes in the original graph.

2. We then constructed a cell signaling network using the random graph and the mRNA expression data for that sample. mRNA expression was left unchanged for each random graph.

3. We computed the ΔG for the top 50 proteins under study on the new random graph.

We then calculated the average and standard deviation ΔG from all 1000 iterations of the above process to compute a bootstrapped null distribution of ΔG. We then computed confidence intervals for ΔG, employing the bonferroni correction to account for multiple hypothesis testing.

Our pipeline was designed to make use of parallel computing on the high-performance cluster (HPC) at Case Western Reserve University.

2.4.6 Identification of miRNA cocktails

To generate miRNA-mRNA mappings, we implemented a protocol described previously [125]. Briefly, we identified all predicted mRNA targets for each miRNA in our dataset using the miRNAtap database in R, version 1.18.0, as implemented through the Bioconductor targetscan org.Hs.eg.db package, version 3.8.2 [93]. We used all five possible databases (default settings): DIANA version 5.061 [95], Miranda 2010 release62 [126], PicTar 2005 release63 [127], TargetScan 7.164 [128] and miRDB 5.065 [94], with a minimum source number of 2, and the union of all targets found was taken as the set of targets for a given miRNA. Through this mapping, we identified a list of mRNA transcripts that are predicted to be repressed by a given miRNA. Our code and processed data files are available on Github at: https://github.com/DavisWeaver/MiR_Combo_Targeting/.

Using this mapping, as well as our ranked list of promising gene candidates for repression from our network analysis, we were able to identify a list of miRNA that we predict would maximally disrupt the Ewing sarcoma cell signaling network when introduced synthetically. To rank miRNA targets, we first identified all the genes on the full target list that a given miRNA was predicted to repress (described in Section 2.4.5). Next, we summed the predicted $\varDelta G$ when each of these genes was repressed *in silico* to generate the maximum potential disruption that could be achieved if a given miRNA were introduced synthetically into an Ewing sarcoma tumor. We then ranked miRNA candidates in descending order of the maximum predicted network disruption (Figure 2.4).

Given the documented cases of systemic toxicities associated with miRNA-based therapies, the miRNA that inhibits the most targets might not necessarily be the best drug target.

We therefore sought to identify combinations of miRNAs that individually repressed key drug targets, while avoiding repression of housekeeping genes that may lead to toxicity. We defined housekeeping genes using a previously described gene set [129]. In this study, housekeeping genes were identified by evaluating RNA sequencing data from a large number of normal tissue samples. Genes that are consistently expressed in all or nearly all tissue types were assumed to be so-called housekeeping genes. Our hypothesis is that by giving a cocktail of miRNAs with predicted activity against one or multiple identified drug targets, each individual miRNA could be given at a low dose such that only the mRNA transcripts that are targeted by multiple miRNAs in the cocktail are affected (Figure 2.1). Also with an eye towards limiting toxicity, we restrained our search to endogenous miRNAs rather than broadening to engineered exogenous miRNA mimics. To that end, we designed a loss function (see equation below) to balance the the effects of repressing the housekeeping gene set I as well as the target gene set J:

$$L(\mu) = \sum_{i,j} A(c)(G_i) - A(c)(G_j) \quad \text{for } i \in I \; and \; j \in J, \tag{2.3}$$

$$A(c) = \begin{cases} 0, & \text{if } c \leq 1 \\ 0.2c, & \text{if } 1 < c < 6 \\ 1, & \text{otherwise} \end{cases} \tag{2.4}$$

Where $A(c)$ determines the degree of repression as a function of the number of times, c, that a given gene, i, is targeted by a given miRNA cocktail, μ. We recognize that assuming each additional miRNA additively represses 20% of a given gene is somewhat arbitrary. Future work could improve on this miRNA cocktail optimization approach by more formally addressing miRNA repression in different contexts.

We first transformed the projected change in network potential for each gene such that housekeeping genes exerted a positive change in network potential and the top 10 predicted targets exerted a negative change in network potential. We then ranked 3-miRNA combinations according to their projected effect on network potential, where more negative changes in network potential were interpreted as most effective for maximizing on-target effects while minimizing off-target effects. As a further constraint, a gene had to be targeted by 2 or more miRNA in a given cocktail to be considered repressed. Each miRNA was assumed to downregulate a given gene by 20%, such that genes targeted by 2 miRs were assumed to have their expression decreased by 40%, and genes targeted by 3 miRs were assumed to have their expression decreased by 60%. We repeated our analysis, varying between 10% and 50% repression to assess the impact of this assumption on our predicted miRNA cocktails. Rather than evaluate every potential 3-miRNA combination, we limited our analysis to miRNA that target at least 2 of our 10 target genes. These constraints were defined *a priori*. We repeated this analysis to identify cocktails that target larger or smaller groups of mRNA (the top 5 or 15 mRNA targets) in order to assess the stability of the predicted cocktail to changing conditions.

Acknowledgements

We acknowledge experimental support from Julia Selich-Anderson. JGS thanks the NIH for their support through NIH R37CA244613, their Loan Repayment Program, and the Paul Calabresi Career Development Award for Clinical Oncology (NIH K12CA076917). KIP acknowledges financial support from the NHMRC CJ Martin Overseas Biomedical Fellowship (APP1111032) and Alex's Lemonade Stand Young Investigator Grant (APP37138). DTW acknowledges financial support from the Cancer Center Trainee Award for Cancer Research from the Case Comprehensive Cancer Center. This work made use of the High Performance Computing Resource in the Core Facility for Advanced Research Computing at Case Western Reserve University. JS acknowledges financial support from the NIH for their support through NIH T32GM007250.

Chapter 3

An open-source R package to facilitate drug target identification.

Pre-printed as: **Weaver, D. T.**, & Scott J. G. (2023, March). Crosstalkr: An open-source R package to facilitate drug target identification, *BioRxiv*, DOI: `https://doi.org/10.1101/2023.03.07.531526`

3.1 Introduction

In the last few decades, interest in graph-based analysis of biological networks has grown substantially. Researchers have leveraged gene regulatory networks, protein-protein interaction networks, and metabolic networks to make predictions about disease gene prioritization, drug target identification, drug repurposing, and patient outcomes [3, 39, 40, 98, 130, 131]. The explosion of methods and data repositories has led some to call this the "Biological Network Era" [4]. For the purpose of this chapter, we will focus on genome-scale protein association networks, often called protein-protein interaction networks (PPI networks). Genome scale PPI databases attempt to represent the relationships between every known protein and every other known protein. PPI networks can be inferred through expert curation of the literature to identify empirical interactions between proteins [12]. PPI networks can also be inferred using gene or protein co-expression, gene coevolution, and other methods to reduce the false negative rate associated with expert curated PPI networks [14, 15].

Researchers have found many uses for PPI networks. For example, they have applied graph search and graph clustering algorithms to biological networks in an effort to derive disease-specific subgraphs [55, 65, 132] or identify potential drug targets [11, 39, 40, 75]. Further, researchers have applied simple node ranking methods to prioritize genes during the generation of predictive biomarkers [133]. Many of the key bioinformatic steps are shared between these use cases. Researchers must interface with a publicly available genome-scale interaction database, reduce the size of the total network to a manageable subnetwork, and finally score nodes so that they can be ranked.

Graph filtration (some times called graph pruning or graph reduction) is a critical step in interactomic pipelines due to the continued expansion of known protein-protein interactions [12, 19]. The best available PPIs have catalogued more than 1 million individual interactions [15]. Network pruning is a key step due to the lack of context specificity in PPI repositories. All recorded interactions between proteins are made available to users, regardless of the cell cycle or phenotypic context in which they were measured. Therefore, it is critical to reduce the full PPI to a more phenotype-specific subgraph. Finally, manipulation of these data structures is extremely cumbersome and makes analysis slow. For example, the full PPI downloadable from StringDB contains 19,271 nodes and 5,934,147 edges. It requires 191.4 MB of memory.

PPI filtration is typically performed using a walk-based algorithm or through integration with additional phenotypic data such as gene expression [40, 55]. One of the most well-studied algorithms in this context is random walk with restarts. Random walk with restarts (RWR) has been used and adapted across disciplines and industries for applications ranging from internet search engines to drug target identification [70,134,135]. There are a growing suite of tools available in R for analyzing graph-structured data [136, 137], including a few R packages that implement RWR [138, 139]. These tools require some understanding of graph data structures and ask the user to find, download, and manipulate the relevant biological networks into adjacency matrices or igraph objects. In this chapter, we will describe crosstalkr, a free, open-source R software package. Crosstalkr provides a streamlined interface to facilitate all three of the most common steps of an interactomic analysis. Users can interface with online PPI repositories, prune or filter the resulting networks, and rank nodes based on a variety of graph-based scoring methods. In addition, crosstalkr is optimized to facilitate one-line implementation of an RWR-based algorithm designed to identify functional subgraphs of protein-protein interaction networks (PPI) [55]. In addition to providing a clean interface for non-graph theorists, crosstalkr improves upon existing tools by implementing methods to facilitate *in silico repression*.

3.2 Design and Data Sources

Crosstalkr facilitates all 3 of the most common steps in an interactomic pipeline; interfacing with online PPI repositories, pruning the resulting network, and ranking nodes to produce scored gene sets (**Figure 3.1**). We implemented this functionality in the “load_ppi”, “gfilter”, and “node_repression” methods (described below).

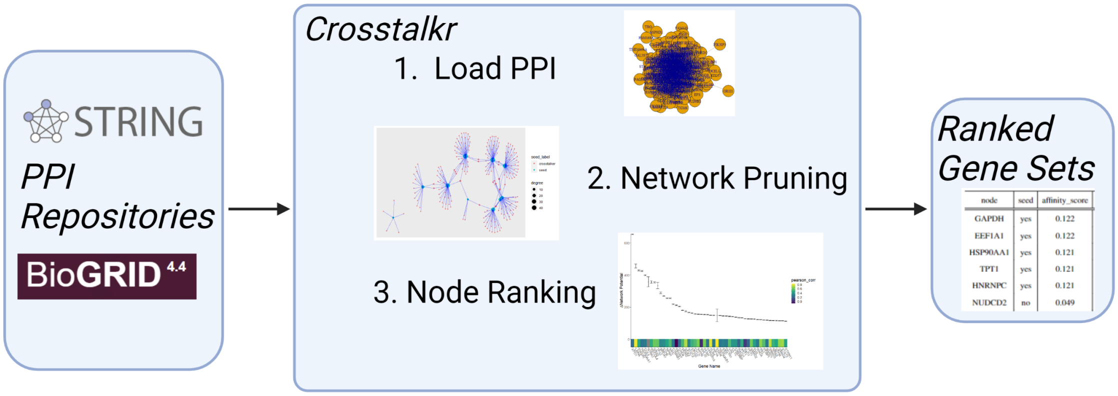

node	seed	affinity_score
GAPDH	yes	0.122
EEF1A1	yes	0.122
HSP90AA1	yes	0.121
TPT1	yes	0.121
HNRNPC	yes	0.121
NUDCD2	no	0.049

Figure 3.1: Example workflow for crosstalkr software package. crosstalkr supports all 3 of the most common steps in an interactomic pipeline; interfacing with online PPI repositories, pruning the resulting network, and ranking nodes to produce scored gene sets.

3.2.1 load_ppi

‘load_ppi’ allows users to interface directly with online protein interaction repositories. Currently, we support StringDB [15], and Biogrid [12]. If users specify StringDB as the ppi, they can customize the incoming PPI based on species code, interaction confidence score, and interaction type. Of note, StringDB incorporates Biogrid curated interactions in their network inference process [15]. Biogrid can therefore be thought of as a more strict subset of the PPI provided by StringDB. load_ppi will standardize the incoming data and return an igraph object where proteins are vertices and binary interactions are undirected edges. However, users do not necessarily need to interact with load_ppi directly. The methods below will call load_ppi if the ‘use_ppi’ bool is set to TRUE. Finally, users are encouraged to provide a file path to the ‘cache’ argument. The returned igraph object will then be stored at the local file path as a .Rda file rather than requiring users to repeatedly download resources from online repositories. By default, cached resources will be used when load_ppi is called subsequently.

3.2.2 gfilter

‘gfilter’ and related methods allow users to reduce large graphs to subgraphs based on a user-supplied method. Users can supply a graph as an igraph object or use a PPI network. All node scoring methods in the igraph package [136] are supported, in addition to simple node ranking based on a user-provided named numeric vector (i.e. gene expression). Node scoring methods also include one thermodynamic measure, network potential (also called Gibbs free energy) [74]. Finally, users can use an RWR-based method (described below). In all cases, users specify the number of nodes to be kept (n), and whether nodes should be ranked in descending order. After computing the node score based on the user-supplied method, a subgraph is creating using the ‘induced_subgraph’ igraph method.

3.2.3 compute_crosstalk

In the compute_crosstalk function, we implement a system for identification of phenotype-specific subnetworks. If users plan to search a supported protein protein interaction network, they are only required to provide a vector of seeds proteins. Seed proteins are used as the starting point for construction of the phenotype-specific sub-network, analogous to search terms in a web search. In this situation, compute_crosstalk will:

1. Download the requested PPI (or load it from the provided cache)
2. Process the requested PPI into a sparse adjacency matrix.
3. Perform a random walk with restart using the user provided seeds to generate affinity scores for every protein in the PPI.
4. Perform many random walks with restarts from n random seeds with a matching degree distribution to generate a null distribution of affinity score.
5. Compare the affinity scores to the null distribution to compute an adjusted p-value (using the method specified in p_adjust)
6. Remove proteins with an adjusted p-value < significance_level

Users can make use of caching to store processed PPIs and speed up future analyses substantially. Users can also make use of parallel computing by setting the ncores parameter > 1. For a formal definition of the algorithm, refer to **Section 3.5**.

3.2.4 node_repression

The node repression function provides support for *in silico* repression, a method to rank nodes for the identification of potential drug targets. *In silico* repression attempts to score the importance of a given node by computing some global measure of network state before and after node removal. A more formal definition of our *in silico* repression implementation can be found in **Section 3.5**. Currently, network potential (sometimes called Gibbs Free Energy) is the only state function available. Generalizing node_repression to any state function that scores nodes is an active area of development.

3.2.5 Data Sources

Users can leverage two high quality PPI networks through crosstalkr; StringDB and Biogrid [12, 13]. Biogrid is the most comprehensive available database of expert-curated empirically measured protein-protein interactions. StringDB expands the coverage of Biogrid by estimating interactions using gene expression, genomic context, and other evidence streams. While Biogrid supports dozens of species, we only provide direct integration for homo sapiens. Crosstalkr provides direct access to all core species support by String. For a more thorough discussion of these resources, refer to **Chapter 1**.

3.3 Example Uses

Here, We present two possible bioinformatic pipelines for drug target identification. Both rely on the integration of RNA or protein expression data from a given model system with interactomic data from protein-protein interaction networks. Integration of specific phenotypic data helps address the shortcoming that publicly available PPI repositories are not context-specific [13]. In the first, we execute a two-step pipeline, where the PPI network is filtered based on RNA expression data. We then rank the nodes again based on betweenness centrality to identify the most critical hub genes of the network. In the second, we use *in silico* repression to rank the nodes rather than a simple node scoring method after a filtration step.

For these examples, we will rely on previously published gene expression data from a Ewing Sarcoma cell line (A673)g (described in more detail here). Ewing Sarcoma is a rare bone malignancy primarily seen in pediatric and adolescent patients. It is thought to be driven by a fusion event that creates an abberant EWS-FLI1 transcription factor [80]. EWS-FLI1 then activates downstream effectors that drive cellular growth and proliferation.

The specific transcriptional effects have been studied extensively in the last 30 years [107, 140, 141]. Despite improved knowledge about the molecular mechanisms that cause Ewing Sarcoma, meaningful clinical improvements have not been made in decades [78]. While it is not our goal here to suggest specific targets for therapy, novel therapeutics with efficacy in Ewing Sarcoma are sorely needed. We will first isolate a single replicate for convenience and log2 normalize ($E = log_2(E)$) the gene expression values E. Then we will format the expression values into a named numeric vector, the preferred format for interacting with crosstalkr.

3.3.1 Pipeline 1

Graph Filtering

First, we will use the gene expression values to reduce the size of the PPI network we need to analyze (**Listing 3.1**). gfilter.value performs the following actions:

1. Download and modify PPI according to user-provided parameters (passed to load_ppi)
2. Take the maximum or minimum n genes according to the user provided named numeric vector (passed to the val parameter)
3. Call the igraph::induced_subgraph method to create a new igraph object that contains only the genes from step 2.
4. Add gene expression to the new graph object as an attribute to allow further manipulation.

```
g <- gfilter(method = "value", cache = NULL, val=exp,
             val_name = "expression", use_ppi = TRUE,
```

```
             min_score= 800,desc = TRUE,n=100)
length(igraph::V(g))
#> [1] 100
head(igraph::get.vertex.attribute(g, name = "expression"))
#> [1] 3.846040 3.961776 3.980389 3.843208 3.982157 3.955569
```

Listing 3.1: *gfilter* can reduce the size of a graph using a user-provided value to rank nodes.

Node Scoring

Next, we will use gfilter again to score this smaller subgraph according to the betweenness centrality and return only the top 5 proteins (**Listing 3.2**). It is worth noting that there are a number of potential measures of centrality that could be used here, as well as many other potentially useful node scoring methods provided in the igraph package [136].

```
g <- gfilter(g=g, igraph_method = "betweenness", n = 5, desc=
   TRUE,
             use_ppi=FALSE, val_name = "betweenness")
igraph::V(g)
#> + 5/5 vertices, named, from cbf2833:
#> [1] GAPDH    HSP90AA1 EEF1A1   HNRNPC   TPT1
```

Listing 3.2: *gfilter* can return the top n nodes ranked by centrality measures like betweenness

The top 5 proteins, ranked by betweenness centrality on this reduced subgraph were GAPDH, HSP90AA1, EEF1A1, HNRNPC, and TPT1. GAPDH is a glycolytic enzyme that has also been assigned a number of other functions, and is highly expressed in normal bone marrow. HSP90AA1 encodes for a heat shock protein and has been implicated in drug resistance development in cancer. EEF1A1 is involved in protein translation. HNRNPC is a ubiquitously expressed RNA binding protein. TPT1 is a regulator of cellular proliferation, and is causally implicated in cancer development [1]. Of our 5 candidate proteins, 4 are either ubiquitously expressed housekeeping genes or proteins involved in the generic cellular stress response. Our very simple pipeline (applied to just a single sample) identified a potentially promising candidate in TPT1. TPT1 was recently associated with Ewing Sarcoma drug resistance development in the literature [140]. TPT1 was found to be significantly upregulated in a doxorubicin-resistant Ewing Sarcoma cell line compared to an embryonic fibroblast cell line.

Disease-Specific Subnetwork Identification

Next, we use the compute_crosstalk function to identify a Ewing Sarcoma specific subnetwork from the full protein-protein interaction network (**Listing 3.3**). Here, we rely on the proteins identified above as seeds in the RWR-based algorithm.

```
out <- gfilter.ct(seeds = names(igraph::V(g)), use_ppi = TRUE,
                  min_score=800, return_df=TRUE, cache = NULL,
                  seed_name = "vignette")

plot_ct(out[[2]], out[[1]])
```

Listing 3.3: *gfilter* can reduce the size of a graph using the compute_crosstalk pipeline to rank nodes

We identified 341 proteins in the subnetwork defined by GAPDH, HSP90AA1, EEF1A1, HNRNPC, and TPT1. Most of these are neighbors of the highly connected HSP90AA1 (heat shock protein). The proteins with the highest affinity for the 5 seeds are labeled. ‘plot_ct‘ is a convenience function to quickly plot the returned subgraph (**Figure 3.2**).

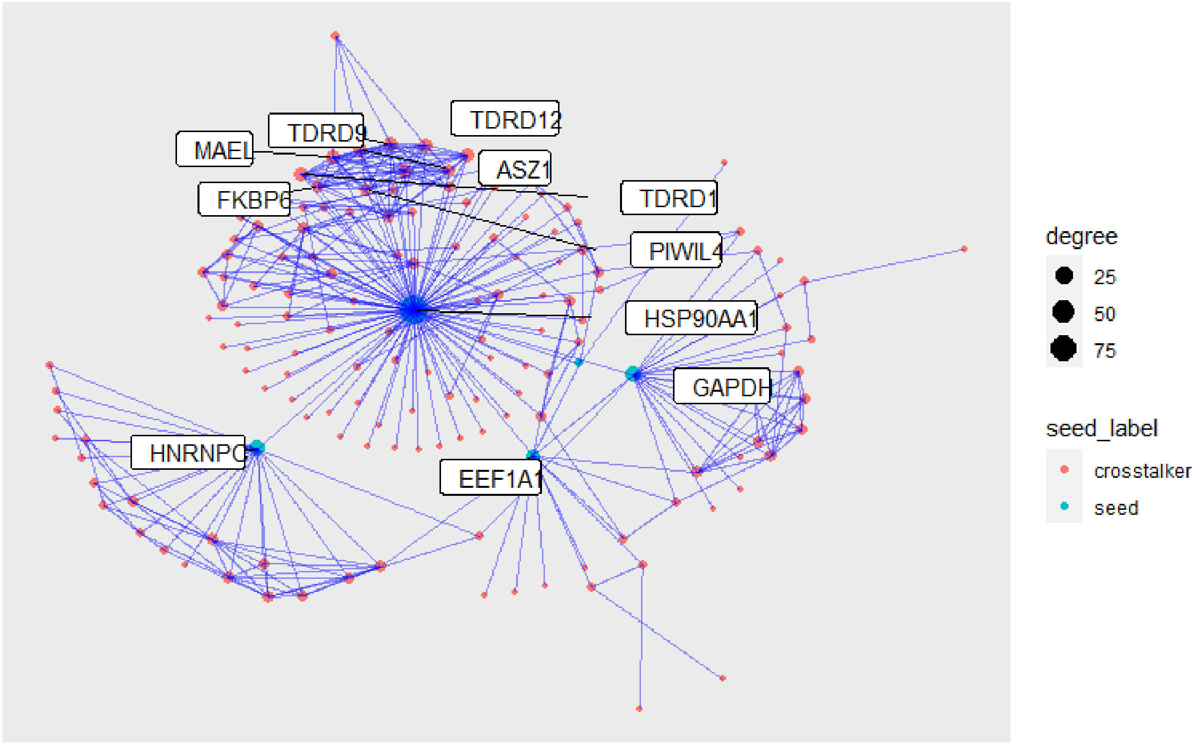

Figure 3.2: Protein-protein interaction subnetwork for GAPDH, HSP90AA1, EEF1A1,HNRNPC, and TPT1. HSP90AA1 was identified as a critical hub in the computed subnetwork.

Users can specify ‘prop_keep‘ to improve readability by only plotting the top x% of identified proteins, ranked by affinity score. gfilter.ct also returns additional information about the proteins in the identified subnetwork (**Table 3.1**).

Not including the initial seeds, NUDCD2, TSSK6m RALYL, POTEF, and TOMM34 showed the highest affinity for the provided seed proteins. Next, we will re-analyze these data using a different combination of analytic steps to illustrate more features of crosstalkr.

3.3.2 Pipeline Two

In this pipeline, we will reduce the graph two times, first using degree (number of neighbors for a given node), and then using gene expression (**Listing 3.4**). The resulting graph has

node	seed	affinity_score	adj_p_value
GAPDH	yes	0.122	0
EEF1A1	yes	0.122	0
HSP90AA1	yes	0.121	0
TPT1	yes	0.121	0
HNRNPC	yes	0.121	0
NUDCD2	no	0.049	0
TSSK6	no	0.049	0
RALYL	no	0.048	0
POTEF	no	0.025	0
TOMM34	no	0.024	0

Table 3.1: Top 10 proteins ranked by affinity to Ewing Sarcoma-related seeds.

500 nodes, all with both high connectivity in the PPI network and high expression in our Ewing Sarcoma cell line. There are 9624 edges (interactions) between these 500 nodes (proteins).

```
g <- gfilter(g = g_ppi, use_ppi=FALSE, cache = NULL, n = 2000,
   igraph_method = igraph::degree, val_name = "degree")
g <- gfilter(g=g, use_ppi = FALSE, n = 500, method = "val",
   val_name = "expression", val = exp)
igraph::gsize(g)
#> [1] 9624
```

Listing 3.4: *gfilter* can reduce the size of a graph using centrality measures like degree to rank nodes.

Node Scoring

Next, we will demonstrate a method for node scoring that uses *in silico* repression (**Listing 3.5**). For this example, we will use network potential as the state function. The resulting data structure is sparse matrix that shows the effect of removing the node in a given column on the state function value for the node represented in a given row. The global effect of removing a given node is then the column sum.

```
dnp <- node_repression(g=g, v_rm = names(igraph::V(g)), exp = exp)
dnp <- Matrix::colSums(dnp)
dnp <- sort(dnp, decreasing = TRUE)
names(dnp)[1:5]
#> [1] "HSP90AA1" "RPS27A"   "UBA52"    "CDK1"     "CCNB1"
```

Listing 3.5: *node_repression* ranks nodes through in-silico repression.

Using this modified pipeline, we identified HSP90AA1, RPS27A, UBA52, CDK1, and CCNB1 as the most critical proteins in this Ewing Sarcoma cell signaling network. HSP90AA1 was also identified in the first instance. RPS27A is a ribosomal protein that is involved in RNA binding. UBA52 is a paralog of RPS27A and appears to be an unlikely candidate for cancer therapy. CDK1 is a cyclin dependent kinase, involved in cellular proliferation and causally associated with cancer in the literature. CCNB1 is a cyclin, a cellular regulator of mitosis.

As with the method above, our bioinformatic pipeline identified a number of ubiquitiously expressed housekeeping genes. Interactomic analyses are often biased towards proteins with extremely high connectivity (i.e. high degree). For example, HSP90AA1 was also identified as a key hub in the Ewing Sarcoma network when we analyzed these data in **Chapter 2**. However, in that study, we performed a permutation test that showed the impact of HSP90AA1 was likely driven by the underlying network connectivity rather than by the actual phenotype under study. To do this, we computed a null distribution for our state function of interest by generating thousands of rewired graphs with preserved degree and then re-calculating our metric of interest. More details about this process can be found in **Chapter 2** [39]. In crosstalkr, We support this capability for the network potential example (see compute_null_dnp). Supporting this feature for additional state functions is an area of development.

3.4 Discussion

Despite the widespread integration of interactomics into bioinformatic workflows, no actively maintained R packages directly supported the core functions of an interactomic analysis. To remedy this, We developed crosstalkr, a free, open-source R package. Crosstalkr provides a toolkit for drug target identification, and supports users at every step in the pipeline. Crosstalkr ships with functions that helps users interface with centralized repositories of functional interactions, filter or prune the resulting PPI networks, and score nodes.

Many organizations have developed web applications to help industry and academic stakeholders perform basic interactomic analyses, including simple queries of protein neighborhood and random-walk inspired tests to identify functionally related proteins [13, 55, 138, 142, 143]. Unfortunately, closed-source web applications are often limited by the lack of programmatic access to the resources. We developed crosstalkr as an R package to enable rapid integration into modular, R-native bioinformatic pipelines. Further, users cannot rely on long-term maintenance of products in this sector. The vast majority of web applications (and software packages) that have been developed to assist with interactomic analyses fall into disrepair within a few years of their release [38, 144, 145]. Against that backdrop, open-source frameworks that can be extended and maintained by third-parties are far more likely to provide to provide long-term value to the academic community.

Overall, crosstalkr builds on a vibrant landscape of tools to help users interact with functional interaction networks, and provides a novel implementation of *in silico* repression. Crosstalkr has been downloaded more than 1700 times and has 13 stars on github, highlighting the potential impact of this work. Crosstalkr can be download directly from CRAN using “install.packages("crosstalkr")” from any R environment. The most recent developmental version can be found on github (https://github.com/DavisWeaver/crosstalkr).

3.5 Algorithm Definitions

3.5.1 Random Walks with Restarts

As mentioned above, compute_crosstalk is an implementation of the random walk with restarts, followed by a permutation test for signficance. The mathematical definition for a random-walk with restarts has been written many times, and can be found at [70]. Here, we provide pseudocode to describe the algorithm implementation in crosstalkr. w is a square adjacency matrix describing the graph. $seeds$ is a vector describing the indexes of the provided seeds. $gamma$ is a number describing the discount rate. eps describes the stop condition, when the change in p in a single iteration is less than eps, the algorithm will stop. $tmax$ describes the largest number of iterations that the algorithm will attempt. The ouput is p, a vector containing the affinity score of all nodes in w relative to the provided $seeds$.

```
Input: w, seeds, gamma, eps, tmax
DO
    Normalize w,
    Initialize p0
    Initialize p to be a vector of zeros with length nrow(w)
for t in 1:tmax, DO
    SET
        pold = p
```

```
        p = ((1-gamma) * w %*% pold + gamma * p0
    IF
        p-pold < eps
        STOP
    ELSE
        CONTINUE
output: vector of affinity scores
```

Listing 3.6: psuedocode for crosstalkr implementation of random walk with repeats

3.5.2 Bootstrapped Null Distributions

We computed bootstrapped null distributions using two related methods. In the algorithm implemented by compute_crosstalk, we bootsrapped a null by re-calculating affinity scores (through RWR) associated with a series of randomly selected seeds that shared a similar degree distribution compared to the user-provided seeds. For more details on this method, refer to [55]. In the algorithm implemented by compute_null_dnp, we bootstrapped a null by generating n completely new graphs with preserved degree distribution using the "keeping_degseq" method. We described this method in detail in **Chapter 2** and in [39].

3.5.3 Network Potential

Network Potential (i.e. Gibbs Free Energy) is a node-specific metric that depends on the local context of each node in a graph. We discuss network potential in detail in **Chapter 2**. For a given node V_i and node weight (C_i) in a graph $G(V, E)$, network potential (G_i) can be computed as follows:

$$G_i = C_i \ln \left[\frac{C_i}{\sum C_j + C_i} \right], \tag{3.1}$$

where C_j refers to the node weights for all neighbors of V_i.

3.5.4 Betweenness Centrality

Betweenness centrality measures the number of shortest paths that traverse a given node v. Betweenness can be computed as follows:

$$g(v) = \sum_{s \neq v \neq t} \sigma_{st}(v) / \sigma_{st} \tag{3.2}$$

where σ_{st} is the number of shortest paths from node s to node$t$ and $\sigma_{st}(v)$ is the number of shortest paths from node s to node t that traverse node v.

3.5.5 In silico repression

In silico repression attempts to quantify the importance of a given node v in some graph $G(V, E)$ by computing the a global measure of graph state S before and after the removal of node v from the graph. In crosstalkr, *in silico* repression is a 6-step process:

1. Calculation of a node score s_v for all nodes $v \in V$ in the graph using some state function $s(v)$.
2. Calculation of the total network state $S = \sum s_v$
3. Removal of a given node i from the graph.
4. Re-calculation of node score $s_i v$ for all nodes $v \in V$. (In practice, we only re-calculate the node score of those nodes affected by the removal of node i)

5. Re-calculation of total network state $S_i = \sum_{i \neq v} s_{iv}$

6. Calculation of $\mathit{\Delta} S_i = S - S_i$ where $\mathit{\Delta} S_i$ is used to score and rank nodes.

Part II

AI-enabled treatment of evolving cell populations

Chapter 4

Reinforcement Learning informs optimal treatment strategies to limit antibiotic resistance

Pre-printed as: **Weaver, D. T.**, Maltas, J, & Scott J. G. (2023, January). Reinforcement Learning informs optimal treatment strategies to limit antibiotic resistance, *BioRxiv*, DOI: `https://doi.org/10.1101/2023.01.12.523765`

4.1 Introduction

Drug resistant pathogens are a wide-spread and deadly phenomenon that were responsible for nearly 5 million deaths worldwide in 2019 [146]. In the US alone, 3 million cases of antimicrobial resistant infections are observed each year [147]. The increasing prevalence of pan-drug resistance has prompted the CDC to declare that we have entered a "post-antibiotic era" [147]. Despite this evident public health crisis, development of novel antibiotics has all but ceased due to the poor return on investment currently associated with this class of drugs [148]. Novel approaches to designing therapies that explicitly take into account the adaptive nature of microbial cell populations while leveraging existing treatment options are desperately needed.

Evolutionary medicine is a rapidly growing discipline that aims to develop treatment strategies that explicitly account for the capacity of pathogens and cancer to evolve [149–155]. Such treatment strategies, termed "evolutionary therapies", typically cycle between drugs or drug doses to take advantage of predictable patterns of disease evolution. Evolutionary therapies are typically developed by applying optimization methods to a mathematical or simulation-based model of the evolving system under study [156–166]. For example, in castrate-resistant prostate cancer, researchers developed an on-off drug cycling drug protocol that allows drug-sensitive cancer cells to regrow following a course of treatment [167, 168]. Clinical trials have shown this therapy prevents the emergence of a resistant phenotype and enables superior long-term tumor control and patient survival compared to conventional strategies [167, 168].

Current methods for the development of evolutionary therapies require an enormous amount of data on the evolving system. For example, many researchers have optimized treatment by using genotype-phenotype maps to define evolutionary dynamics and model the evolving cell population [160, 169–177]. However, most methods for optimization of these models requires a complete understanding of the underlying system dynamics [159, 160, 178, 179]. Such detailed knowledge is currently unobtainable in the clinical setting. Approaches that can approximate these optimal policies given only a fraction of the available information would fill a key unmet need in evolutionary medicine.

We hypothesize that reinforcement learning algorithms can develop effective drug cycling policies given only experimentally measurable information about the evolving pathogen. Reinforcement learning (RL) is a well-studied subfield of machine learning that has been successfully used in applications ranging from board games and video games to manufacturing automation [178, 180–182]. Broadly, RL methods train artificial intelligence agents to select actions that maximize a reward function. Importantly, RL methods are particularly suited for optimization problems where little is known about the dynamics of the underlying system. Further, RL and related optimal control methods have been previously applied for the development of clinical optimization protocols in onocology and anesthesiology [165, 183–188].

In this study, we develop of a novel approach to discovering evolutionary therapies, using a well studied set of empirical fitness landscapes as a model system. We will explore "perfect information" optimization methods such as dynamic programming in addition to RL methods that can learn policies given only limited information about a system. We show that it is possible to learn effective drug cycling treatments given extremely limited information about the evolving population, even in situations where the measurements reaching the RL agent is extremely noisy and the information content is low.

index	drug code	drug
1	AMP	Ampicillin
2	AM	Amoxicillin
3	CEC	Cefaclor
4	CTX	Cefotaxime
5	ZOX	Ceftizoxime
6	CXM	Cefuroxime
7	CRO	Ceftriaxone
8	AMC	Amoxicillin + Clavulanic acid
9	CAZ	Ceftazidime
10	CTT	Cefotetan
11	SAM	Ampicillin + Sulbactam
12	CPR	Cefprozil
13	CPD	Cefpodoxime
14	TZP	Pipercillin + Tazobactam
15	FEP	Cefepime

Table 4.1: Reference codes for drugs under study

4.2 Methods

As a model system, we simulated an evolving population of *Escherichia coli (E. coli)* using the well-studied fitness landscape paradigm, where each genotype is associated with a certain fitness under selection [160,170,173]. We relied on a previously described 4 allele landscape of the *E. Coli* β-lacatamase gene where each mutation has a measured impact on the sensitivity of an *E. coli* population to one of 15 β-lactam antibiotics [170, 173]. We then defined 15 different fitness regimes on the same underlying genotype space, each representing the selective effect of one of 15 β-lactam antibiotics (Table 4.1) [170]. We used this well-studied *E. coli* model system because it is one of the few microbial cell populations for which a combinatorially complete genotype-phenotype mapping has been measured [170,173]. By simulating an evolving *E. coli* cell population using the described fitness landscape paradigm, we were able to define an optimization problem on which to train RL agents (**Fig 4.1**).

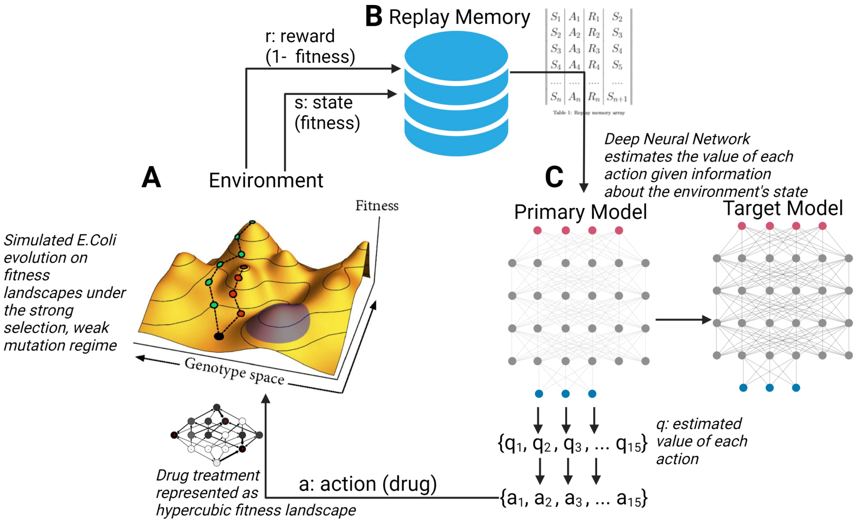

Figure 4.1: Schematic of artificial intelligence system for controlling evolving cell populations. A: *E. coli* population evolving on fitness landscapes under the strong selection, weak mutation evolutionary regime. At each time step, a reward signal r and a measure of system state s are sent to the replay memory structure. **B:** Replay memory array stores (s, a, r, s') tuples where s' is state s+1. These are then used to batch train the neural network. **C:** Deep Neural network estimates the value of each action given information about the environment's state. The action with the largest estimated value is then applied to the evolving cell population.

4.2.1 Simulation of Evolution Using Fitness Landscapes

We use a previously described fitness-landscape based model of evolution [160, 171]. In brief, we begin by modeling an evolving asexual haploid population with N mutational sites. Each site can have one of two alleles (0 or 1). We can therefore represent the genotype of a population using an N-length binary sequence, for a total of 2^N possible genotypes. We can model theoretical drug interventions by defining fitness as a function of genotype. These "drugs" can then be represented using N-dimensional hyper-cubic graphs (**Fig 4.1A**). Further, if we assume that drug evolution under drug treatment follows the strong selection and weak mutation paradigm, we can then compute the probability of mutation between adjacent genotypes and represent each landscape as a Markov chain as described by Nichol et al. [160]. At each time step, we sampled from the probability distribution defined by the Markov chain to simulate the evolutionary course of a single population.

4.2.2 Optimization approaches

We applied two related optimization approaches to identify effective drug cycling policies in this setting. First, we extended the Markov chain framework to formulate a complete Markov Decision Process (MDP). An MDP is a discrete-time framework for modeling optimal decision-making [178]. Critically, the system under study must be partially under the control of the decision-making agent. MDPs can be solved using dynamic programming to generate optimal policies for the defined control problem [178]. The dynamic programming algorithm requires perfect information (e.g. the complete transition matrix and instantaneous state from the MDP) in order to yield optimal policies. Next, we trained agents with imperfect information using reinforcement learning to approximate a clinical scenario where perfect information will never be available. Notably, the state set, action set, and reward assignment were shared between the perfect and imperfect information conditions. The action set corresponded to the drugs available to the optimization process. We considered this system to have a finite time horizon (20 evolutionary steps in the base case). We chose a finite time horizon rather than an infinite time horizon assumption in order to more faithfully represent clinical disease courses. For our purposes, we assume that one evolutionary time step is the equivalent of one day of evolution.

Perfect Information

The state set S represents all potential genotypes (16 total in our base case) that the evolving population can explore. The action set A corresponds to the 15 available β-lactam antibiotics. Finally, we define the reward set (R) and the set of transition probabilities (P)

as a function of the current genotype s as well as the chosen action, a (eq 4.1):

$$\begin{gathered} R = 1 - f(s|a) \text{ for } s \in S \text{ and } a \in A, \text{and,} \\ P = f(s_{t+1}|s_t, a_t) \text{ for } s \in S \text{ and } a \in A. \end{gathered} \tag{4.1}$$

We solved the defined MDP using backwards induction, a dynamic programming approach designed to solve MDPs with finite time horizons [189], to generate an optimal drug cycling policy for each evolutionary episode. Backwards induction is used to estimate a value function $V(s)$ which estimates the discounted reward of being in each state s. Optimal policies $\mathit{\Pi}(s)$ are then inferred from the value function. Throughout the remainder of the paper we will refer to this optimal drug cycling policy as the "MDP" condition.

Imperfect Information

In order to assess the viability of developing optimal drug therapies from potentially clinically available information, we trained a Deep Q learner to interact with the evolving *E. Coli* system described above. Deep Q learning is an extremely well-studied and characterized method of reinforcement learning, and is particularly suited to situations where very little *a priori* knowledge about the environment is available [178, 190]. We used two different training inputs to model a gradient of information loss. In the first condition, termed RL-genotype, the instantaneous genotype of the population was provided as the key training input at each time point. For this condition, the neural architecture was composed of an input layer, two 1d convolutional layers, a max pooling layer, a dense layer with 28 neurons, and an output layer with a linear activation function.

In the second condition, termed RL-fit, instantaneous population fitness of the population was provided as the key training input at each time point. The neural architecture of RL-fit was composed of a neural network with an input layer, two dense hidden layers with 64 and 28 neurons, and an output layer with a linear activation function. RL-fit takes population fitness at time t and one-hot encoded action at time $t-1$ as inputs and outputs Q-values. Q-values are estimates of the future value of a given action. Q-value estimates are improved by minimizing the temporal difference between Q-values computed by the current model and a target model, which has weights and biases that are only updated rarely. We used mean squared error (MSE) as the loss function.

We further explored the effect of information content on learned policy effectiveness by introducing a noise parameter. With noise active, fitness values $s \in S$ that were used as training inputs were first adjusted according to:

$$s_t = s_t + w \ \in \ W \sim \mathcal{N}(\mu, 0.05 \times \sigma^2). \tag{4.2}$$

For the noise experiment, μ was set to 0 such that $\sigma^2 = 0$ would introduce no noise. We then varied σ^2 (referred to as ’noise parameter’) from 0 (no noise) to 100 (profound loss of signal fidelity).

All code and data needed to define and implement the evolutionary simulation and reinforcement learning framework can be found at https://github.com/DavisWeaver/evo_dm. The software can be installed in your local python environment using ’pip install git+https://github.com/Da

4.3 Results

In this study, we explored the viability of developing effective drug cycling policies for antibiotic treatment given less and less information about the evolving system. To this end,

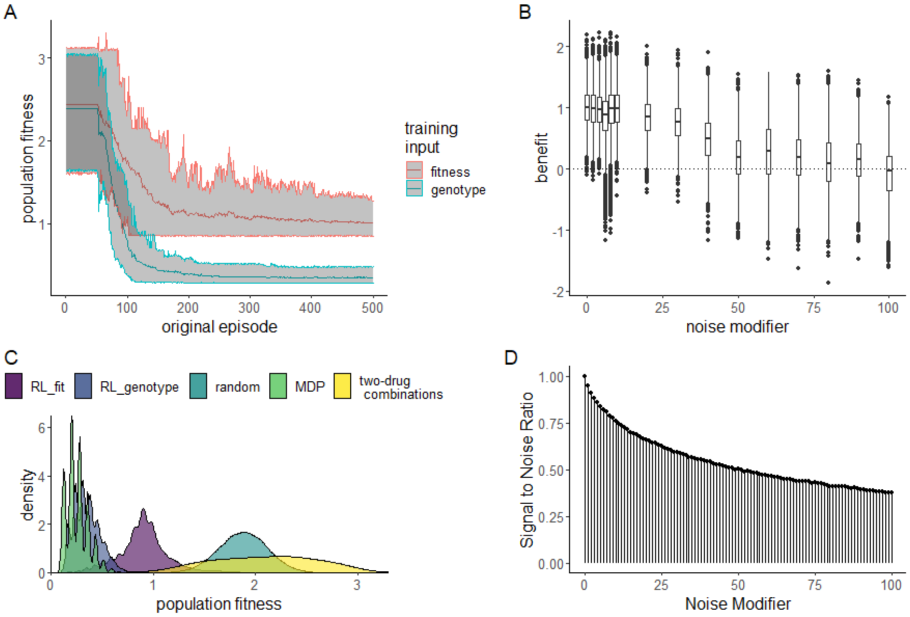

Figure 4.2: Performance of RL agents in a simulated *E. coli* system. **A:** Line plot showing the effectiveness of the average learned policy as training time increases on the x-axis for RL agents trained using fitness (red) or genotype (blue). **B:** Boxplot showing the effectiveness of 10 fully trained RL-fit replicates as a function of noise. Each data point corresponds to a single episode. The width of the distribution provides information about the episode by episode variability in RL-fit performance. **C:** Density plot summarizing the performance of the two experimental conditions (measured by average population fitness) relative to the three control conditions. **D:** Signal to noise ratio associated with different noise parameters. Increasing noise parameter decreases the fidelity of the signal that reaches the reinforcement learner.

we developed a reinforcement learning framework to design policies that limit the growth of an evolving *E. Coli* population *in silico*. We evaluated this system in a well-studied *E. coli* system for which empirical fitness landscapes for 15 antibiotics are available in the literature [170]. A given RL agent could select from any of these 15 drugs when designing a policy to minimize population fitness. We defined three experimental conditions. In the first, we solved a Markov decision process formulation of the optimization problem under study. In doing so, we generated true optimal drug cycling policies given perfect information of the underlying system (described in **Section 4.2.1**). In the second, RL agents were trained using the current genotype of the simulated *E. Coli* population under selection (RL-genotype). Stepping further down the information gradient, RL agents were trained

drug sequence	replicate	condition
CTX,AMC,CTX,CPR,CTX,CPR,CTX,CPR,CTX,CPR	53	RL-fit
CTX,CPR,CPR,CPR,CTX,CPR,CPR,CPR,CTX,SAM	53	RL-genotype
CTX,AMC,CTX,AMC,CTX,AMC,CTX,AMC,CTX,CPR	23	RL-fit
CTX,AMC,CTX,AMC,CTX,AMC,CTX,AMC,CTX,AMC	23	RL-genotype
CTX,AMC,CTX,AMC,CTX,CPR,CTX,AMC,CTX,CPR	96	RL-fit
CTX,SAM,CTX,SAM,CTX,CPR,CTX,CPR,CTX,CPR	96	RL-genotype

Table 4.2: Example drug sequences. Here, we show the first 10 selected drugs for representative episodes of the three top-performing replicates.

using only observed fitness of the *E. Coli* population (RL-fit). Finally, we introduced noise into these measures of observed fitness to simulate real-world conditions where only imprecise proxy measures of the true underlying state may be available. Each experimental condition was evaluated based on its ability to minimize the fitness of the population under study. We compared these conditions to two negative controls; a drug cycling policy that selects drugs completely at random (which we will refer to as "random"), and all possible two-drug cycles (i.e AMP-AM-AMP-AM-AMP). We tested 100 replicates of RL-fit and RL-genotype against each of these conditions. Each replicate was trained for 500 episodes of 20 evolutionary steps (10000 total observations of system behavior). We chose 500 episodes as the training time after extensive hyper-parameter tuning showed decreased or equal effectiveness with additional training.

Comparison of RL drug cycling policies to negative controls.

We found that both RL conditions dramatically reduced fitness relative to the random policy. In both cases, the RL conditions learned effective drug cycling policies after about 100 episodes of training and then fine-tuned them with minimal improvement through episode 500 (**Fig 4.2A**). As expected, RL-genotype learned a more effective drug cycling policy on average compared to RL-fit. RL-genotype had access to the instaneous state (genotype) of the evolving population, while RL-fit was only trained using a proxy measure (population fitness). In 98/100 replicates, we observed a measurable decrease in population fitness under the learned RL-fit policy versus a random drug cycling policy (**Fig 4.7A**). Further, we found that the average RL-fit replicate outperformed all possible two-drug cycling policies (**Fig 4.2C**). RL-genotype outperformed both negative controls in all 100 replicates (**Fig 4.2C**). In some replicates, RL-genotype achieved similar performance compared to the MDP policy (**Fig 4.7D**). In addition, the distribution of performance for RL-genotype policies nearly overlapped with MDP performance (**Fig 4.2C**). Introduction of additional noise to the training process for RL-fit led to degraded performance. However, even with a large noise modifier, RL-fit still outperformed the random drug cycling condition. With a noise modifier of 40 (fitness $+\mathcal{N}(\mu = 0, \sigma^2 = 0.05 \times 40)$), RL-fit achieved an average population fitness of 1.41 compared to 1.88 for the random drug cycling condition (**Fig 4.2D**).

Overview of learned drug cycling policies for RL-fit and RL-genotype.

We evaluated the learned drug cycling policies of RL-fit, RL-genotype for the 15 β-lactam antibiotics under study. We compared these to the true optimal drug cycling policy as a reference. For this system, we show that the optimal drug cycling policy relies heavily on Cefotaxime, Ampicillin + Sulbactam, and Ampicillin (**Fig 4.3A**). Cefotaxime was used as treatment in more than 50% of time-steps, with Ampicillin + Sulbactam and Ampicillin used next most frequently. The optimal drug cycling policy used Cefprozil, Pipercillin + Tazobactam, and Cefaclor infrequently. The remaining drugs were not used at all. The different RL-fit replicates largely converged on a similar policy. They relied heavily on Cefotaxime and Amoxicillin + Clavulanic acid. However, they relied infrequently on Cefprozil. RL-genotype replicates also converged on a relatively conserved policy. Further, RL-genotype replicates showed a much more consistent mapping of state to action

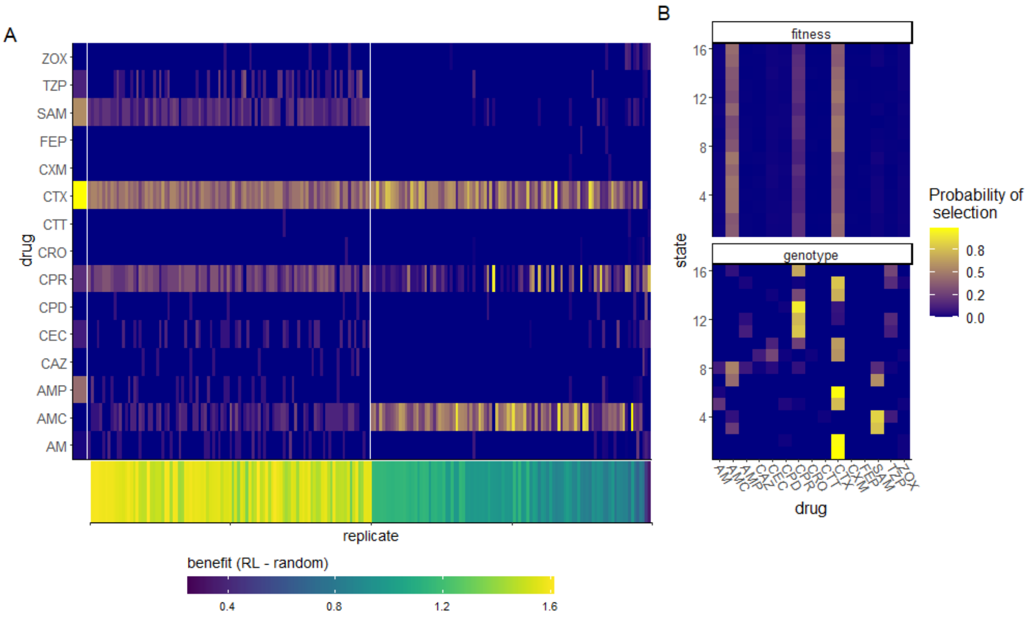

Figure 4.3: Drug cycling policies learned by RL-genotype and RL-fit. A: Heatmap depicting the learned policy for 100 replicates (on the x-axis) of the RL-genotype and 100 replicates of RL-fit. Far left column (enlarged) corresponds to the optimal policy derived from the MDP condition. The Y-axis describes the β-lactam antibiotics each RL agent could choose from while the color corresponds to the probability that the learned policy selected a given antibiotic. Bottom heatmap shows the median fitness benefit observed under the policy learned by a given replicate. **B:** Heatmap showing the average learned policy for RL-fit and RL-genotype. RL-genotype learns a more consistent mapping of state to action compared to RL-fit.

compared to RL-fit (**Fig 4.3B**). All optimization paradigms identified complex drug cycles that use 3 or more drugs to treat the evolving cell population. None of the tested two-drug combinations compete with policies learned by RL-genotype, and are generally outperformed by RL-fit. We show that policies that do not rely on Cefotaxime are suboptimal in this system. The three replicates that showed the least benefit compared to the random drug cycling case did not use Cefotaxime at all (**Fig 4.3B**). The importance of Cefotaxime is likely explained by the topography of the CTX drug landscape (**Fig 4.12**). More than half of the available genotypes in the CTX landscape lie in fitness valleys, providing ample opportunities to combine CTX with other drugs and "trap" the evolving population in low-fitness genotypes.

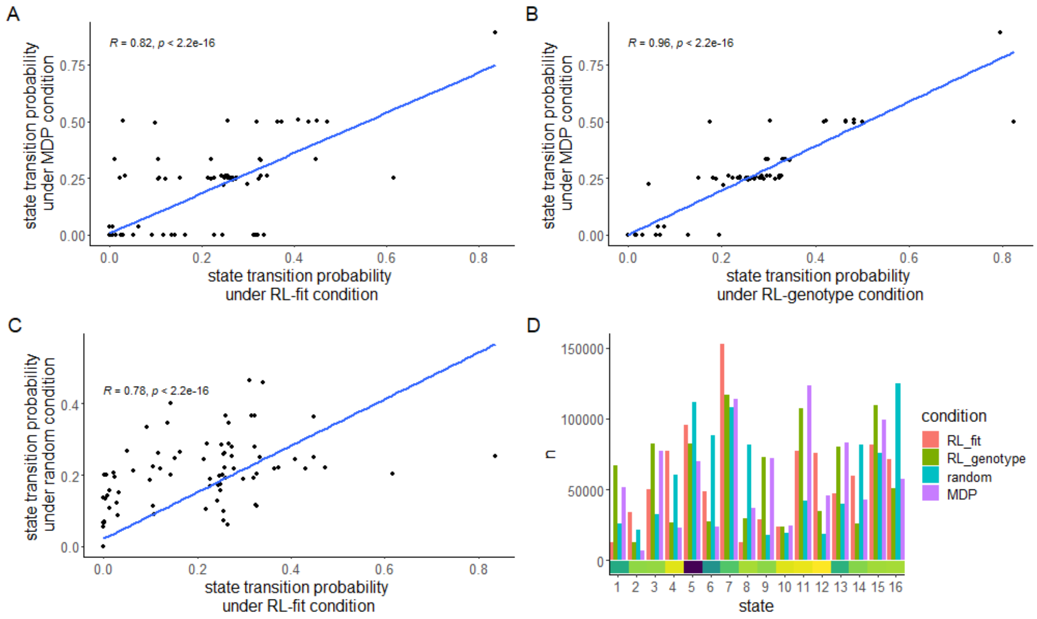

Figure 4.4: Comparison of evolutionary trajectories seen under different regimes A-C: Selected Pairwise comparisons of state transition frequency under different experimental conditions. State transition frequency is nearly identical for the RL-genotype and MDP conditions (R=0.97). In contrast, state-transition-frequency for the RL-fit and MDP conditions are related but less strongly correlated (R=0.75). As expected, state transition frequency is not very similar between the RL-fit and random conditions (R=0.62). **D:** Bar chart comparing the frequency that states are observed under different experimental conditions. The value of each state (to the learner) is highlighted for each state by the bottom heatmap. High value states are observed more frequently in RL-fit, RL-genotype, and MDP conditions compared to the random condition.

Evolutionary trajectories observed under RL-Genotype, RL-fit, and MDP drug policies.

Next, we compared the evolutionary paths taken by the simulated *E. coli* population under the MDP, RL-fit, RL-genotype, and random policy paradigms. The edge weights (corresponding to the probability of observed state transitions) of the RL-genotype and MDP landscapes show a 0.96 pearson correlation (**Fig 4.4**). In contrast, the edge weights of the RL-fit and MDP landscapes show a 0.82 pearson correlation (**Fig 4.4**). During the course of training the MDP condition, the backwards induction algorithm generated a value function $V(s, a)$ for all $s \in S$ and $a \in A$. In **Figure 4.4D**, we use this value function to show that certain genotypes (namely 1,5,6, and 13) were more advantageous to the evolving population than to the learner. These states were frequented much more often under the random drug cycling condition compared to any of the experimental conditions (**Fig 4.4D**). We also show that other genotypes (namely 12 and 11) were particularly advantageous for the learner compared to the evolving population. These states were frequented much more often under the experimental conditions compared to the random drug cycling condition (**Fig 4.4D**).

We also show that certain state transitions occur more frequently than others, independent of experimental conditions. For example, the population nearly always transitioned from genotype 5 to genotype 7 (**Fig 4.5**). This transition highlights the way these learned policies use drug landscapes to guide evolution. Genotype 5 (0100) is a fitness peak in most of the drug landscapes used in the learned policies, and is therefore a very disadvantageous state for the controlling agent. CTX, the most commonly used drug in all effective policies, has a slightly higher peak at genotype 7 (0110), which forces the population away from genotype 5 (**Fig 4.10**). As another example, the evolving population very rarely transitioned from state 1 to state 9 in the RL-fit condition. This state transition occurred commonly in the MDP and RL-genotype conditions (**Fig 4.5**). This difference is explained by the policies shown in **Figure 4.3B**. Under the RL-genotype policy, CTX was selected every time the population was in state 1 (the initial condition). The CTX landscape topography allows transition to 3 of the 4 single mutants, including state 9 (1000) (**Fig 4.12**). Under the RL-fit policy, CTX and AMC were used in about equal proportion when the population is in state 1. Unlike the CTX landscape, the AMC landscape topography does not permit evolution from state 1 to state 9 (**Fig 4.12**)).

Characteristics of selected drug policies

To better understand why certain drugs were used so frequently by RL-genotype, RL-fit, and the MDP policies, we developed the concept of an “opportunity landscape”. An

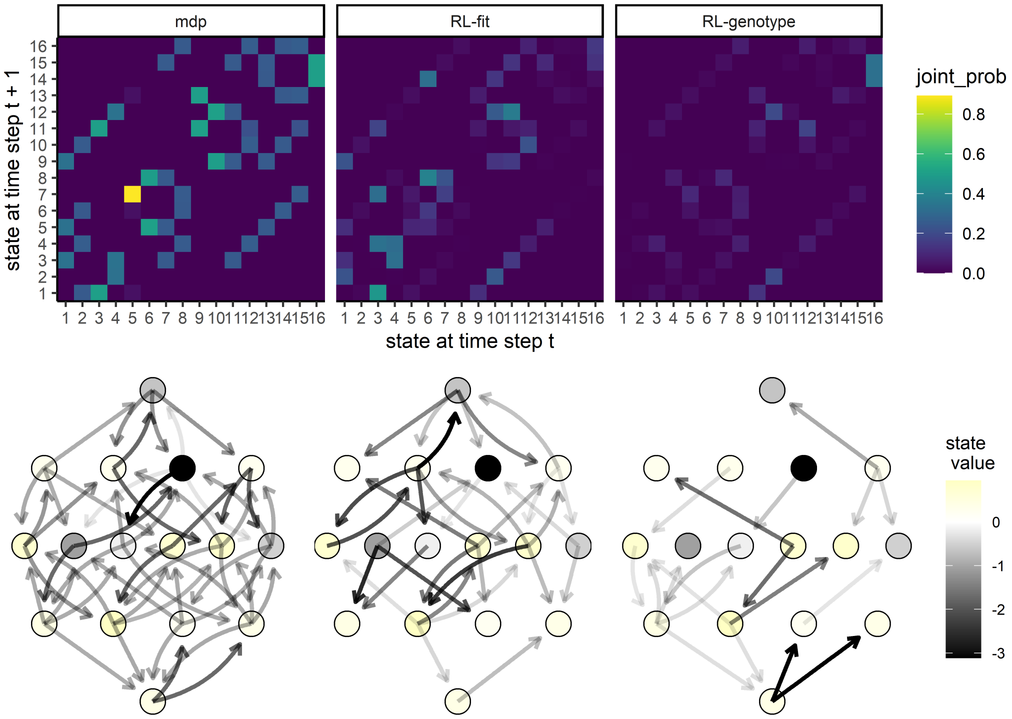

Figure 4.5: Movement of simulated *E. Coli* population through the genomic landscape. Top row: Heatmap depicting the joint probability distribution for each state transition under the different experimental conditions. The second two show the difference in state transition probability compared to the MDP condition. **Bottom row:** Graph depicting the fitness landscape, beginning with the wild type (bottom) all the way to the quadruple mutant (top). Size of arrow depicts the frequency with which a state transition was observed under the labeled experimental condition. The color of each node corresponds with the expected value (to the learner) of being in that state. As above, the second two arrows correspond to the observed difference between RL-Fit or RL-genotype and the MDP condition.

opportunity landscape is an optimistic summation of n fitness landscapes. We computed each opportunity landscape by taking the minimum fitness value for each genotype from a given set of fitness landscapes. This simplified framework gives a sense of a potential best case scenario if the drugs in a given combination are used optimally. For example, the MDP policy relied heavily on CTX, CPR, AMP, SAM, and TZP to control the simulated *E. Coli* population. The resultant opportunity landscape (**Fig 4.6A**) contains only a single fitness peak, with 15/16 of the genotypes in or near fitness valleys. In **Figure 4.6B**, we show the actual state transitions observed during evolution under the MDP policy. We also color the nodes based on the value function estimated by solving the MDP. As expected, the value function estimated by the MDP aligns closely with the topography of the opportunity landscape. There is only one genotype that the value function scores as being very poor for the learner, corresponding to the single peak in the opportunity fitness landscape (**Fig 4.6**). Interestingly, the opportunity landscape predicted that the population would evolve to the single fitness peak and fix. In contrast, the observed state transitions suggest that the MDP policy was able to guide the population away from that single fitness peak. A more detailed discussion of opportunity landscapes can be found in the supplemental materials.

We also show that both the MDP and RL-genotype conditions select the drug with the lowest fitness for most genotypes (**Fig 4.6C**). There are a few notable exceptions to this rule, which highlight RL-genotype's capacity for treatment planning. A greedy policy that selects the lowest drug-fitness combination for every genotype would select Amoxicillin (AM) when the population is identified as being in genotype 5. The AM drug landscape then strongly favors transition back to the wild-type genotype (state 1). From state 1, most available drugs encourage evolution back to the genotype 5 fitness peak. As we see in **Figure 4.6B**, state 5 is by far the least advantageous for the learner. The greedy policy therefore creates an extremely disadvantageous cycle of evolution. In fact, none of the tested policies rely heavily on AM in state 5 (**Fig 4.3B**), instead taking a fitness penalty to select Cefotaxime (CTX). The CTX drug landscape encourages evolution to the double mutant, which has access to the highest value areas of the landscape. Finally, we rank drug landscapes based on the number of genotypes with a fitness value < 1 (**Fig 4.6D**). Based on the defined reward function, these genotypes would be considered advantageous to the learner. We show that drugs identified as useful by the optimal policy or RL-genotype tend to have more advantageous genotypes in their drug landscape. The only two highly permissive landscapes (CPD, CPR) that aren't used have extremely similar topography to

CTX, which most policies were built around.

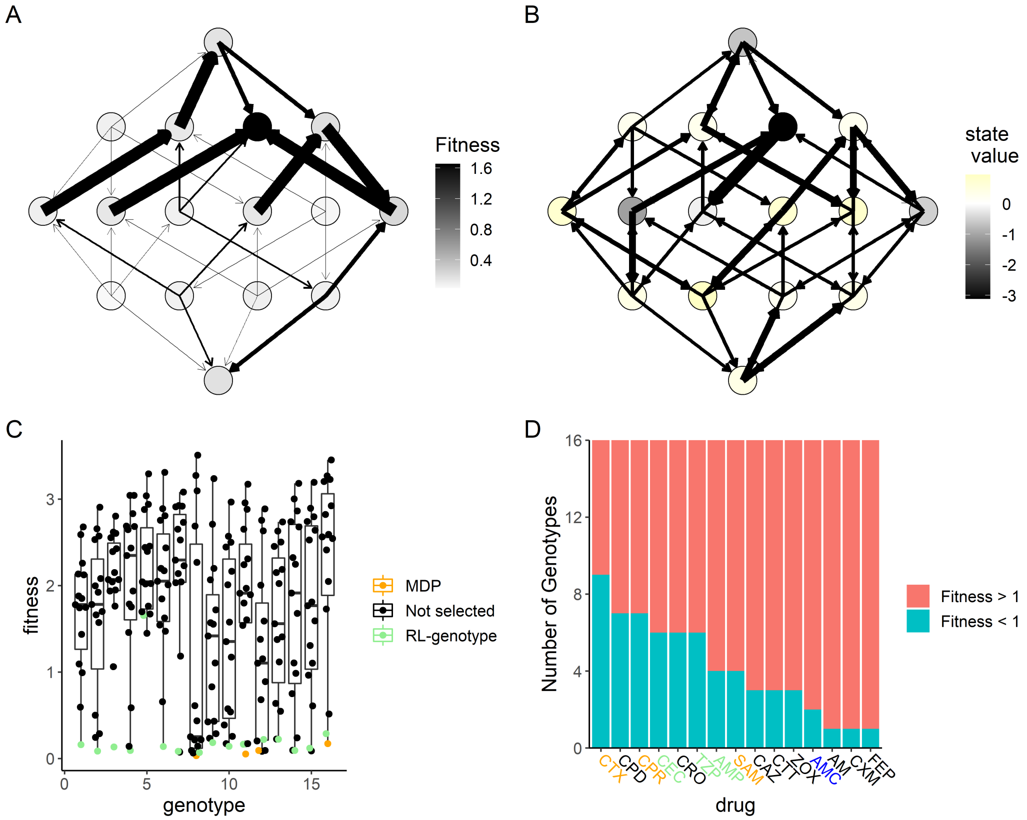

Figure 4.6: MDP value function closely matches opportunity landscape for drugs commonly used under MDP policy. Panels A and B show the 16-genotype fitness landscape under study, starting with the wild type at the top, progressing throw the single mutants, double mutants, triple mutants, and finally the quadruple mutant at the bottom. **A:** Opportunity landscape for the 5 drugs most commonly used under the MDP policy (CTX, CPR, AMP, SAM, and TZP). **B:** Observed state transitions under the MDP policy. The node color corresponds to the value function estimated by solving the MDP. Lower values correspond to states the MDP policy attempts to avoid while higher values correspond to states the MDP policy attempts to steer the population. **C:** Scatter Plot showing the distribution of fitness with respect to genotype for the 15 β -lactam antibiotics under study. The drug selected by RL-genotype in a given genotype is highlighted in light blue. In cases where the MDP selected a different drug than RL-genotype, that drug is highlighted in orange. **D:** Number of genotypes with fitness above or below 1 for each drug under study. Drugs that are used by both the MDP and RL-genotype are highlighted in orange. Drugs that are used by only the MDP are highlighted in green. Drugs that are used by only RL-genotype are highlighted in blue.

4.3.1 Additional Analyses

Hyperparameter tuning

We varied key hyperparameters one at a time in order to identify optimal values to promote learning in this setting. Parameter ranges and the selected value are shown in **Table 4.3**. Due to the long run-times of the training process, we were unable to make use of more formal hyper-parameter optimization approaches. Future work will increase the efficiency of training reinforcement learners in this setting, opening up a number of interesting follow-on studies.

Additional performance data for RL agents

As mentioned in the main text, we tested both the RL-fit and RL-genotype conditions 100 times each. In **Figure 4.7**, we show the perfomance of all 100 RL-fit and RL-genotype replicates. In 98/100 replicates, RL-fit outperformed the random drug cycling case (**Fig 4.7A**). The very best RL-fit replicates still fell short of the MDP-derived optimal policy (**Fig 4.7B**). In all 100 replicates, RL-genotype outperformed the random drug cycling case ((**Fig 4.7C**). RL-genotype performance approached the performance of the optimal policy (**Fig 4.7D**).

Table 4.3: Key Hyperparameters for reinforcement learner

Parameter	Value	Range
gamma	0.99	0-1
learning rate	0.0001	0.000001-0.1
minibatch size	60	20-500
update target model frequency	310	100-1000

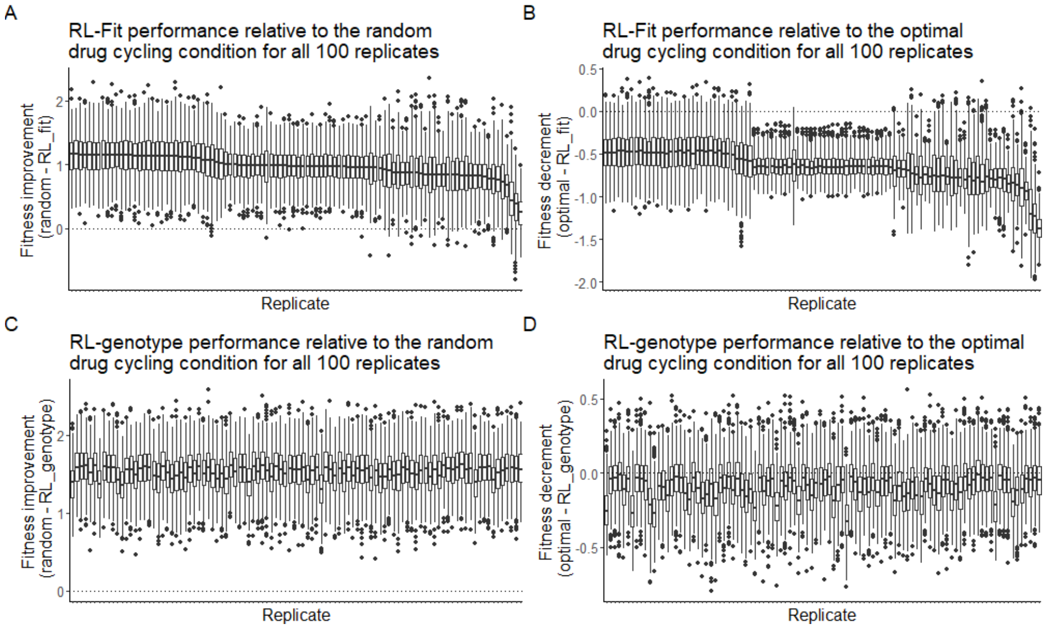

Figure 4.7: Performance of RL-fit and RL-genotype for each replicate. A: Fitness observed under RL-fit policy compared to random drug cycling condition. **B:** Fitness observed under RL-fit policy compared to fitness observed under optimal policy. **C:** Fitness observed under RL-genotype policy compared to random drug cycling condition. **D:** Fitness observed under RL-genotype policy compared to fitness observed under optimal policy.

Policy Clustering

We performed PCA on the policies for all 100 replicates from the RL-genotype and RL-fit conditions. We then used the silhouette method, implemented in the factoextra package, to estimate the appropriate number of clusters. We found that either 2 or 5 clusters would be optimal. As two clusters would only recapitulate our orignal RL-genotype and RL-fit conditions, we performed kmeans clustering with 5 centers. Results are shown in **Fig 4.8**.

We plotted the top 3 pairwise comparisons of principal components, which together account for about 56.4% of the variance in this dataset. The RL-fit and RL-genotype policies were clearly separated by this method, with RL-genotype being split between clusters 1,2, and 4. RL-fit was split between clusters 3 and 5 (Fig 4.8). Clusters 3 and 5 represent meaningfully different policy motifs that RL-fit found frequently over the course

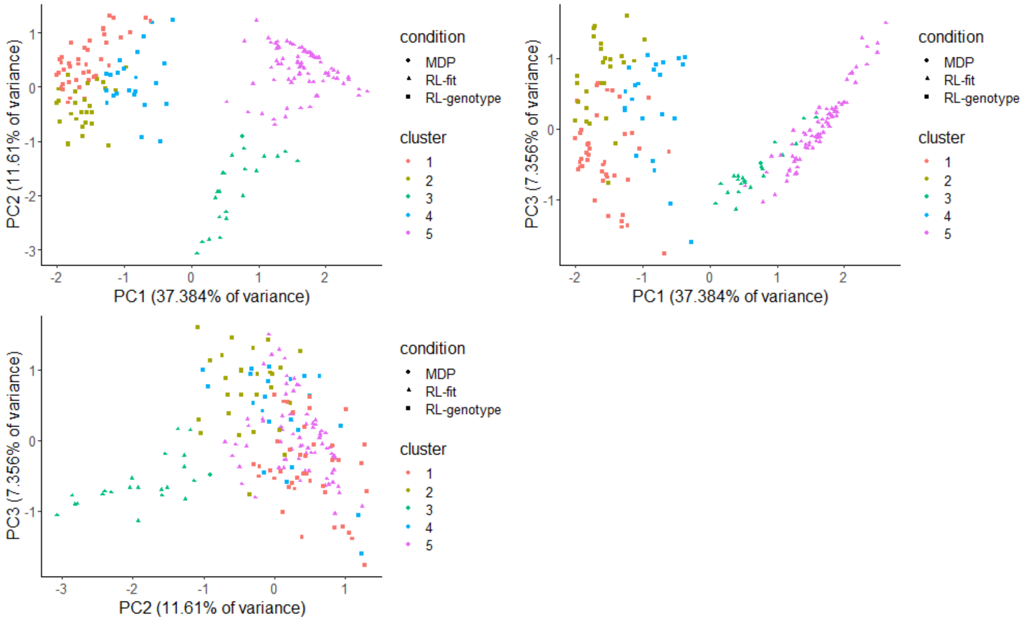

Figure 4.8: Identification of policy groups using PCA and kmeans clustering. RL-fit replicates were separated into 2 distinct replicates. RL-genotype replicates were separated into 3 distinct groups but it is unclear if these groups are meaningfully different.

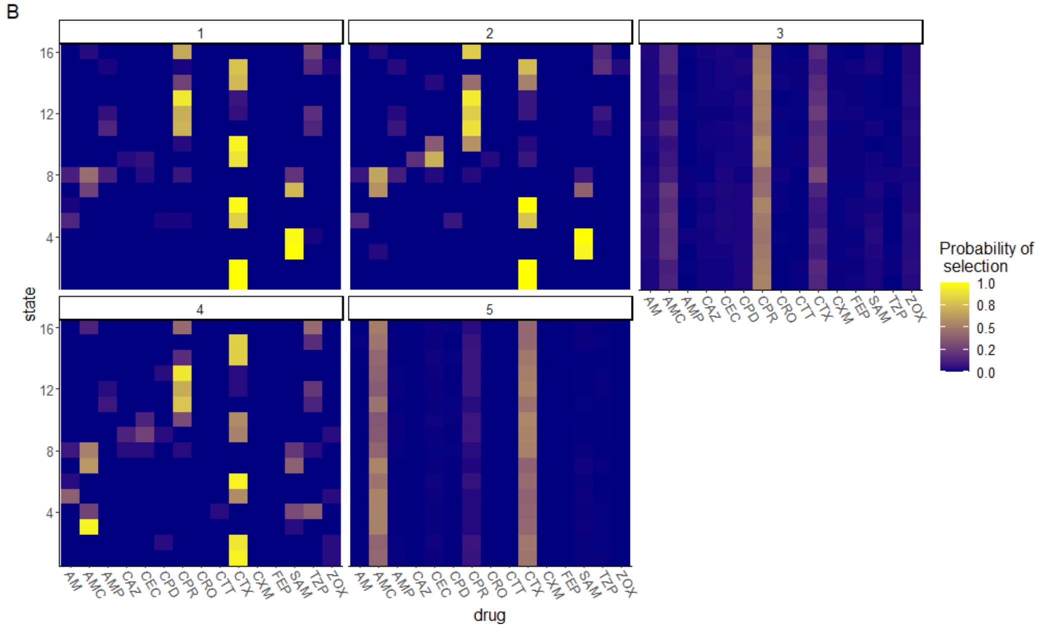

Figure 4.9: policy heatmaps for groups identified using PCA and kmeans clustering. Groups 1,2, and 4 correspond to RL-genotype policies. Groups 3 and 5 correspond to RL-fit policies. color gradient represents probability that a given drug (x-axis) will be selected when population is in a given state (y-axis).

of 100 replicates. Cluster 3 replicates were Cefprozil dominant, and used Cefotaxime and Amoxicillin + Clavulanic acid infrequently. Cluster 5 replicates were Cefotaxime and Amoxicillin + Clavulanic acid dominant, and used Cefprozil infrequently (**Fig 4.9**). Notably, cluster 3 replicates tended to have worse performance compared to cluster 5 replicates (**Fig 4.3**).

Opportunity Landscapes

We define an opportunity landscape to be the most optimistic combination of n landscapes, formed by taking the minimum possible fitness at each genotypic position. This construct can help us better understand how the learner uses different combinations of drugs to maintain the evolving population at extremely low fitness values. Figure 4.10 describes the opportunity landscape discovered by the MDP condition. As noted in the main text, the MDP primarily uses 5 drugs (CTX, CPR, AMP, SAM, and TZP) in combination to trap the evolving population of E. Coli at extremely low fitness genotypes. In the combined opportunity landscape, just one genotype (0100) had a high fitness in all 5 drugs. As expected, the opportunity landscape closely matches the value function estimated by the MDP (**Fig 4.5**)

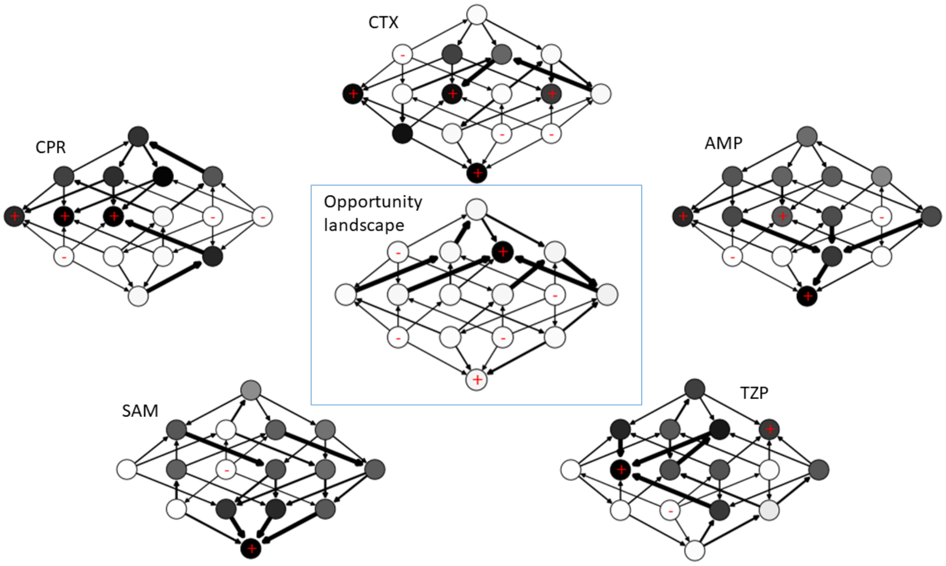

Figure 4.10: Opportunity Landscape for MDP-derived policy. Opportunity landscape is an optimistic combination of 5 empirically measured drug landscapes. Just 1/16 genotypes is near a fitness peak on the opportunity landscape, helping to explain the extremely low fitness observed in the simulated E.Coli population when the MDP-derived policy is applied.

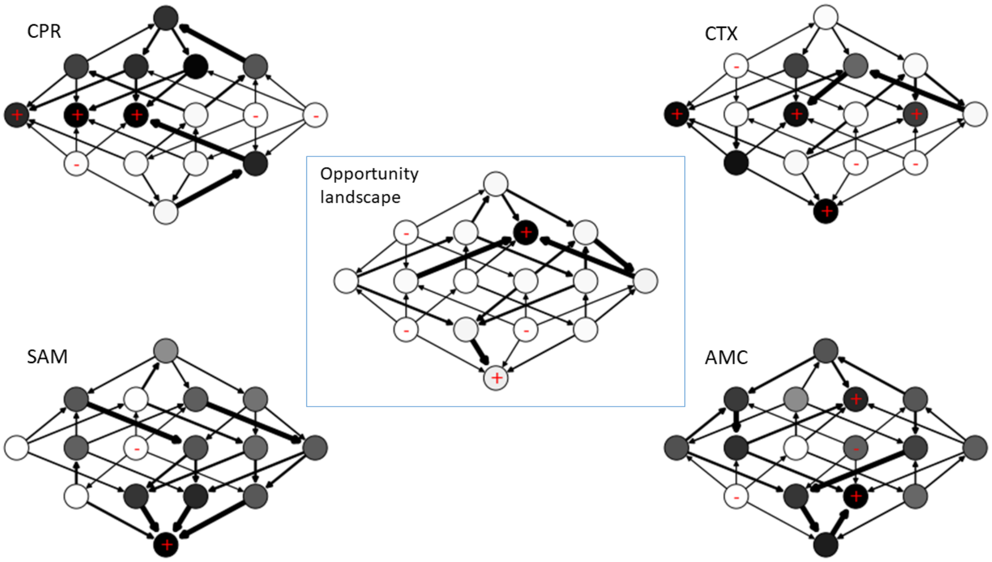

Figure 4.11: Opportunity landscape for most common policy identified in the RL-genotype condition. As in the MDP-derived policy, just 1/16 genotypes is near a fitness peak in the opportunity landscape.

The opportunity landscape for the RL-genotype is almost identical to the opportunity landscape observed for the MDP policy (**Fig 4.11**). Interestingly, RL-genotype only uses 3 of the 5 drugs in the MDP policy; CPR, CTX, and SAM. RL-fit discovered policies that typically only used two drugs. The most effective RL-fit policies relied heavily on AMC and CTX. We present the resulting opportunity landscape in **Figure 4.12**. As expected, there are more genotypes with high fitness values under this two-drug paradigm compared to the 4 or 5 drug policies discovered by RL-genotype and the MDP, respectively.

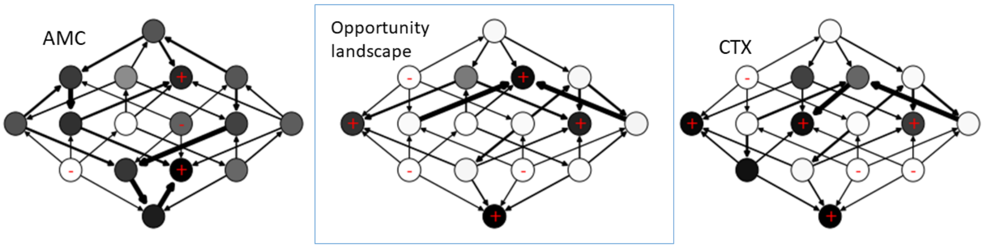

Figure 4.12: Opportunity landscape for the most common policy identified in the RL-fit condition. The most common RL-fit policy relies on AMC and CTX to control the E. Coli population. Assuming the most optimistic combination of these two drug landscapes, 4/16 genotypes are near a fitness peak.

MDP policy

As mentioned in the main text, we computed the MDP policy by formulating a markov decision process of the strong selection, weak mutation model of evolution under study. We then solved the MDP using backward induction, an algorithm designed to identify an optimal policy for a finite time discrete MDP. The identified policy is a function of current state and current time step, making it even more specific than the policies identified by the reinforcement learning conditions. We show the time and state-specific MDP policy in **Figure 4.13**. Near the end of an episode (steps 19 and 20), we see a switch to a greedy policy that simply selects the drug with the minimum fitness for a given genotype. We

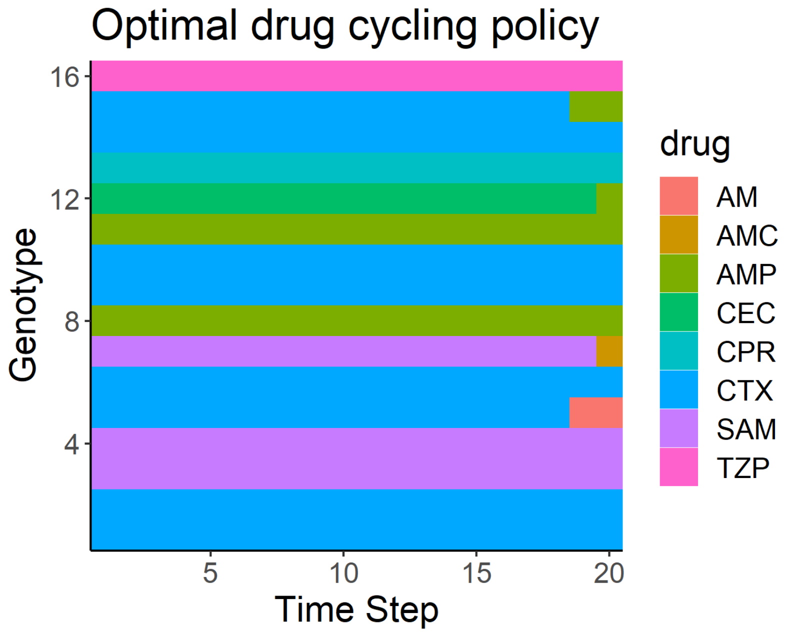

Figure 4.13: MDP-derived optimal policy.

also varied the discount rate (gamma), between 0 and 1 during the hyperparameter tuning process. In **Figure 4.14**, we show the effect of gamma on the avergage fitness achieved by the MDP policy. While gamma didn't have a large effect, likely due to the relatively short length (20 time steps) of each episode, we show that increasing gamma led to increased performance of the computed MDP policy. We also show that increasing gamma led to increased use of CTX (drug 4).

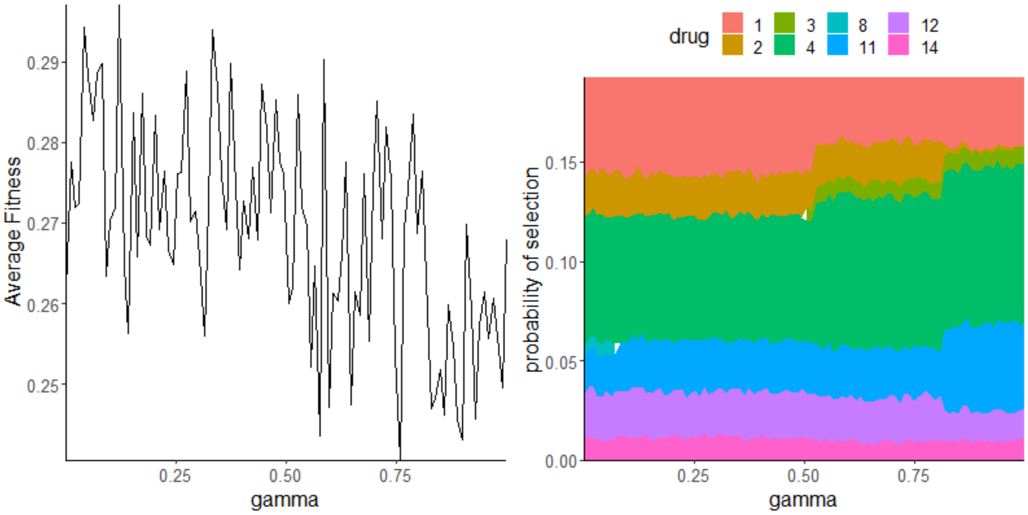

Figure 4.14: Effect of variation in gamma on optimal policy performance and composition

Two-Drug Combinations

As noted in **Figure 4.2C** in the main text, we evaluated the performance of all A-B-A-B two-drug cycles to use as a comparison group for RL-fit and RL-genotype. In **Figure 4.15A**, we examine these combinations in greater depth. We also show the landscape correlation between the two drugs in every combination. We show that anti-correlated landscapes tend to make more effective combinations, likely due to collateral sensitivity. Highly correlated landscapes tend to make ineffective drug combinations, likely due to collateral resistance.

Finally, we evaluated the effect of starting population genotype on the performance of each two-drug combination. We found that the starting genotype of the population had no effect on the overall distribution of performance for these two-drug combinations (**Fig 4.15**).

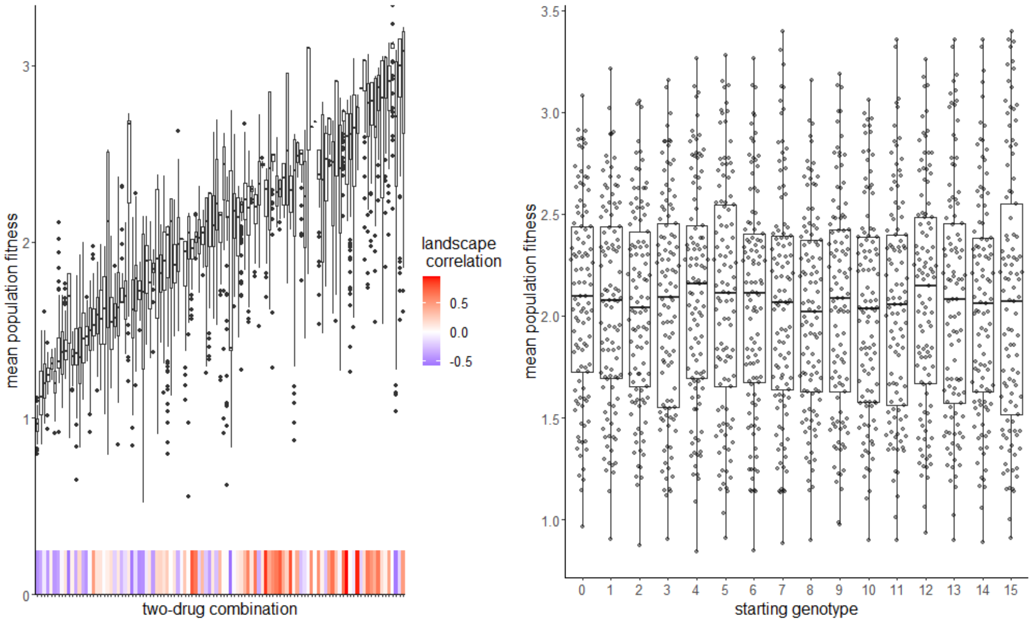

Figure 4.15: Two-drug cycling policies.

4.4 Discussion

The evolution of widespread microbial drug resistance is driving a growing public health crisis around the world. In this study, we show a proof of concept for how existing drugs could be leveraged to control microbial populations without increasing drug resistance. To that end, we tested optimization approaches given decreasing amounts of information about an evolving system of *E.Coli*, and showed that it is possible to learn highly effective drug cycling policies given only empirically measurable information. To accomplish this, we developed a novel reinforcement learning approach to control an evolving population of *E. Coli in silico*. We focused on 15 empirically measured fitness landscapes pertaining to different clinically available β-lactam antibiotics (**Table 4.1**). In this setting, RL agents selected treatments that, on average, controlled population fitness much more effectively than either of the two negative controls. We showed that RL agents with access to the instantaneous genotype of the population over time approach the MDP-derived optimal policy for these landscapes. Critically, we showed that RL agents were capable of developing effective drug cycling protocols even when the measures of fitness used for training were first adjusted by a noise parameter. This suggests that even imperfect measurements of an imperfect measure of population state (the kind of measurements we are able to make in clinical settings) may be sufficient to develop effective control policies. We also show that RL or MDP-derived policies consistently outperform simple alternating drug cycling policies. Finally, we introduced the concept of the "Opportunity Landscape" which can provide powerful intuition into the viability of various drug combinations.

Our work expands a rich literature on the subject of evolutionary control through formal optimization approaches. Our group and others have developed and optimized perfect information systems to generate effective drug cycling policies [156, 157, 159, 161, 162].

Further, a limited number of studies have used RL-based methods for the development of clinical optimization protocols [165, 183–186, 188]. These studies have been limited so far to contrived simulated systems, including a recent study that introduced Celludose, a RL framework capable of controlling evolving bacterial populations in a stochastic simulated system [187].

Much like the studies noted above, we show that AI or MDP-based policies for drug selection or drug dosing dramatically outperform sensible controls in the treatment of an evolving cell population. We also extend this literature in two key ways. To our knowledge, ours is the first optimization protocol capable of learning effective drug cycling policies using only observed population fitness (a clinically tractable measure) as the key training input. Second, we grounded our work with empirically measured fitness landscapes which will facilitate more natural extension to the bench.

There are several limitations to this work which bear mention. We assume that selection under drug therapy represents a strong-selection and weak mutation regime in order to compute transition matrices for our models. While this is likely true in most cases, it is possible that other selection regimes emerge in cases of real world pharmacokinetics where the drug concentration fluctuates dramatically. In addition, we chose to keep drug concentration constant throughout are analysis, largely owing to the lack of robust empirical data linking genotype to phenotype under dose varying conditions (sometimes called a fitness seascape) [191]. As more empirical fitness seascape data becomes available, a natural extension would be to explore the efficacy of the RL system in controlling a population by varying both drug and dose.

While we present the most extensive genotype-phenotype modeling work to date on this subject, we still only modeled the effect of mutations at 4 genotypic positions. The real *E. Coli* genome is approximately 5×10^6 base pairs [192]. The evolutionary landscape for living organisms is staggeringly large, and not tractable to model *in silico*. It is possible that empirical measures of fitness like growth rate or cell count may not provide a robust enough signal of the underlying evolutionary state on real genomes. *In vitro* implementations of reinforcement learning-based drug cycle optimization systems are needed to address this potential shortcoming. Another potential alternative would be to use the comparatively low-dimensional phenotype landscape of drug resistance [193].

In this work, we present a novel reinforcement-learning framework capable of controlling an evolving population of *E. Coli in silico*. We show that RL agents stably learn multi-drug combinations that were state specific and reliably out-performed a random drug cycling policy as well as all possible two-drug cycling policies. We also highlight key features of the types of drug landscapes that are useful for the design of evolutionary control policies. Our work represents an important proof-of-concept for AI-based evolutionary control, an emerging field with the potential to revolutionize clinical medicine.

Acknowledgements

This work was made possible by the National Institute of Health (5R37CA244613-03, 5T32GM007250-46, and T32CA094186) and American Cancer Society (RSG-20-096-01). Figure 1 was created with BioRender.com.

Part III

Discussion and Conclusion

Chapter 5

Other contributions to clinical optimization

Published as:

Imaging Follow-up of Low-Risk Incidental Pancreas and Kidney Findings: Effects of Patient Age and Comorbidity on Projected Life Expectancy, Tiana J. Raphel, **Davis T. Weaver**, Lincoln L. Berland, Brian R. Herts, Alec J. Megibow, Amy B. Knudsen, and Pari V. Pandharipande Radiology 2018 287:2, 504-514

Davis T. Weaver, Tiana J. Raphel, Alexander Melamed, J. Alejandro Rauh-Hain, John O. Schorge, Amy B. Knudsen, Pari V. Pandharipande, Modeling treatment outcomes for patients with advanced ovarian cancer: Projected benefits of a test to optimize treatment selection, Gynecologic Oncology, Volume 149, Issue 2, 2018.

Testing for Verification Bias in Reported Malignancy Risks for Side-Branch Intraductal Papillary Mucinous Neoplasms: A Simulation Modeling Approach, **Davis T. Weaver**, Anna P. Lietz, Sarah F. Mercaldo, Mary Linton B. Peters, Chin Hur, Chung Yin Kong, Brian M. Wolpin, Alec J. Megibow, Lincoln L. Berland, Amy B. Knudsen, and Pari V. Pandharipande American Journal of Roentgenology 2019 212:3, 596-601

Chen Q, Larochelle MR, **Weaver DT**, et al. Prevention of Prescription Opioid Misuse and Projected Overdose Deaths in the United States. JAMA Netw Open. 2019;2(2):e187621. doi:10.1001/jamanetworkopen.2018.7621

5.1 Introduction

The field of medical decision-making strives to improve patient outcomes and increase efficiency in healthcare by applying mathematical and simulation models. For disease areas or patient populations that lack robust clinical trials, medical decisions are often made based on a combination of retrospective studies, expert guidelines, and physician intuition. Clinical trial data may be unavailable due to lack of financial incentives for drug companies, newness of a disease area or treatment, or other factors. Cancer screening in particular is often extremely under-served by clinical trials due to the massive accruals and long follow-up required to uncover a robust signal. Mathematical models and simulations can fill the void left by missing clinical trial data and help physicians better understand the real-world costs and benefits of potential treatment or screening options. Training in a medicial decision-making lab helped prepare me for the modeling and bioinformatic work that I undertook during my PhD. Many of the stochastic and deterministic frameworks we employed are similar to models used in evolutionary biology. Further, model parameterization required me to parse large databases and apply statistical models. My work focused on ovarian and pancreatic cancer, two deadly diseases where outcomes haven't significantly improved in decades.

5.2 Modeling outcomes for patients with co-morbid conditions

In this study, we set out to determine the effects of patient age and comorbidity level on life expectancy (LE) benefits associated with imaging follow-up of Bosniak IIF renal cysts and pancreatic side-branch (SB) intraductal papillary mucinous neoplasms (IPMNs) [194]. We were motivated by the lack of consideration of patients with significant age and comorbidity in other studies on the subject. We developed a decision-analytic Markov model to evaluate LE benefits. Hypothetical cohorts with varied age (60–80 years) and comorbidities (none, mild, moderate, or severe) were evaluated. For each finding, LE projections from two strategies were compared: imaging follow-up and no imaging follow-up. Under follow-up, it was assumed that cancers associated with the incidental finding were successfully treated before they spread. For patients without follow-up, mortality risks from Bosniak IIF cysts (renal cell carcinoma) and SBIPMNs (pancreatic ductal adenocarcinoma) were incorporated. Model assumptions and parameter uncertainty were evaluated in sensitivity analysis.

In the youngest, healthiest cohorts (age, 60 years; no comorbidities), projected LE benefits from follow-up were as follows: Bosniak IIF cyst, 6.5 months (women) and 5.8 months (men); SBIPMN, 6.4 months (women) and 5.3 months (men). Follow-up of Bosniak IIF cysts in 60-year-old women with severe comorbidities yielded a LE benefit of 3.9 months; in 80-year-old women with no comorbidities, the benefit was 2.8 months, and with severe comorbidities the benefit was 1.5 months (**Figure 5.1**). Similar trends were observed in men and for SBIPMN. Results were sensitive to the performance of follow-up for cancer detection; malignancy risks; and stage at presentation of malignant, unfollowed Bosniak IIF cysts. With progression of age and comorbidity level, follow-up of low-risk

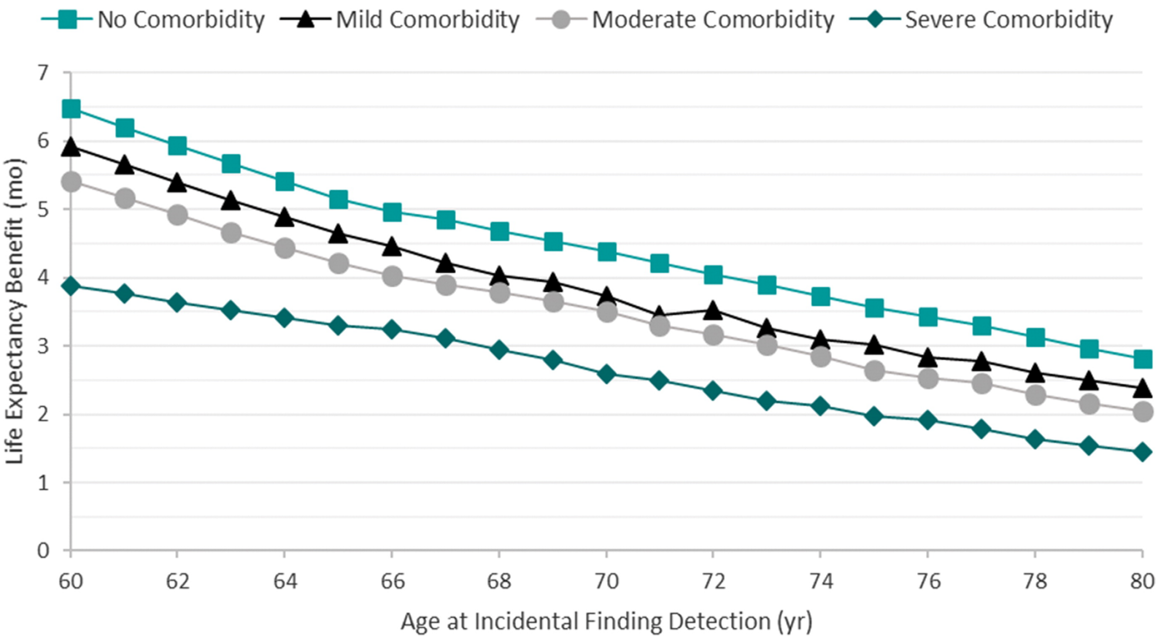

Figure 5.1: Projected benefit in life years as a function of patient comorbidity and age.

incidental findings yields increasingly limited benefits for patients.

5.3 Modeling Treatment Outcomes in Ovarian Cancer

For patients with advanced stage epithelial ovarian cancer (EOC), substantial emphasis has been placed on diagnostic tests that can discern which of two treatment options - primary cytoreductive surgery (PCS) or neoadjuvant chemotherapy followed by interval cytoreductive surgery (NACT+ICS) - optimizes patient-level outcomes [195]. Our goal was to project potential life expectancy (LE) gains that could be achieved by use of such a test. To address this question, we developed a microsimulation model to project LE for patients with stage IIIC EOC. We compared: a "standard-of-care" strategy, in which patients were triaged to PCS vs. NACT+ICS based on current clinical practice; and a "test" strategy, in which patients were triaged based on results of a hypothetical test. We identified those test performance characteristics for which the test strategy outperformed the standard-of-care strategy, from a LE standpoint. Effects of parameter uncertainty were evaluated in sensitivity analysis.

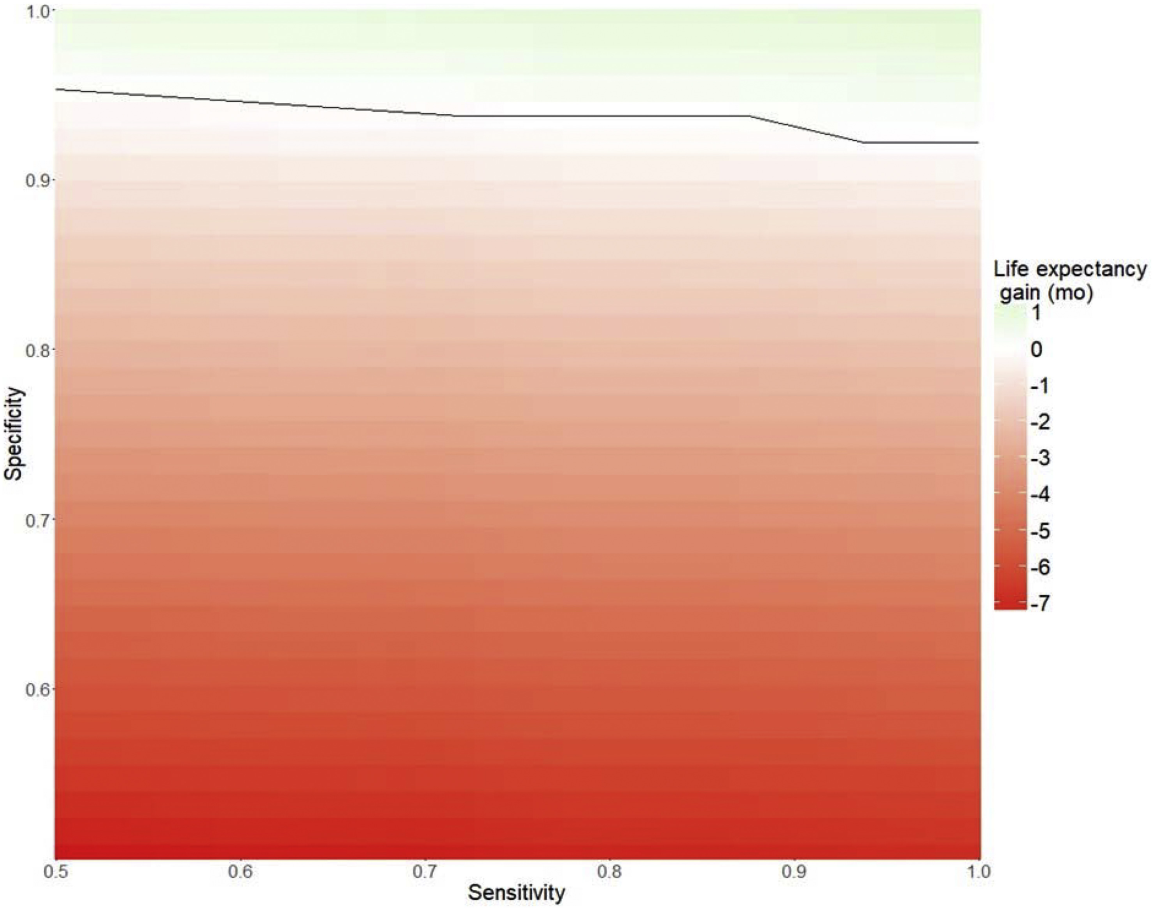

Figure 5.2: Life expectancy (LE) differences (in months) between two strategies ("test" – "standard-of-care") for treatment selection in patients with stage IIIC ovarian cancer, across varied performance assumptions for a hypothetical test. The black line represents test characteristics where the life expectancy gain = 0. High test specificity (correct triage of patients with resectable disease to primary cytoreductive surgery) is more important than high test sensitivity (correct triage of patients with unresectable disease to neoadjuvant chemotherapy plus interval cytoreductive surgery) for achieving life expectancy gains. Even with a perfect test (top right), the life expectancy gain is modest (1.2 months)

Even with a perfect test, We found that the LE gain was modest (LE with test vs. standard-of-care strategy = 67.6 vs. 66.4 months; LE gain = 1.2 months) (**Figure 5.2**). In order to outperform the standard-of-care, the test had to have a high probability of correctly identifying "resectable" patients at PCS (i.e. those for whom complete or optimal cytoreduction would be possible); this test property was more important than correct triage of unresectable patients to NACT+ICS. Results were sensitive to the proportion of patients whose underlying disease was resectable at PCS. In conclusion, we showed that diagnostic tests that are designed to triage patients with advanced stage EOC will likely have only a modest effect on LE.

5.4 Computing the rate of progression for pre-malignant pancreatic lesions

The objective of our study was to test for the possibility that published malignancy risks for side-branch intraductal papillary mucinous neoplasms (IPMNs) are overestimates, likely due to verification bias [196]. We tested for possible verification bias using simulation modeling techniques. First, in age-defined hypothetical cohorts of 10 million persons, we projected the frequency of pancreatic ductal adenocarcinoma (PDAC) arising from side-branch IPMNs over 5 years using published estimates of their prevalence (4.4%) and rate of malignant transformation (1.9%). Second, we projected the total number of PDAC cases in corresponding cohorts over the same time horizon using national cancer registry data. For each cohort, we determined whether the percentage of all PDAC cases that arose from side-branch IPMNs (i.e., side-branch IPMN–associated PDAC cases) was clinically plausible using an upper limit of 10% to define plausibility, as estimated from the literature. Model assumptions and parameter uncertainty were evaluated in sensitivity analysis.

Across all cohorts, percentages of side-branch IPMN–associated PDACs greatly exceeded 10%. In the base case (mean age = 55.7 years), 80% of PDAC cases arose from side-branch IPMNs (7877/9786). In the oldest cohort evaluated (mean age = 75 years), this estimate was 76% (14,227/18,714). In a secondary analysis, we found that if an upper limit threshold of 10% for side-branch IPMN–associated PDAC was imposed, the model-predicted rate of malignancy for side-branch IPMNs would be less than 0.24% over a 5-year time horizon, substantially lower than most literature-based estimates. Our results suggest that reported malignancy risks associated with side-branch IPMNs are likely to be overestimates and imply the presence of verification bias.

5.5 Forecasting the trajectory of the opioid pandemic

Deaths due to opioid overdose tripled between 2008-2018. Efforts to curb this trend have focused on restricting the prescription opioid supply; however, the near-term effects of such efforts are unknown. Our objective was to project effects of interventions to lower prescription opioid misuse on opioid overdose deaths from 2016 to 2025 [197]. We defined a mathematical model using of the US opioid epidemic using ordinary differential equations. We then projected outcomes of simulated individuals who engage in nonmedical prescription or illicit opioid use from 2016 to 2025. The analysis was performed in 2018 by retrospectively calibrating the model from 2002 to 2015 data from the National Survey on Drug Use and Health and the Centers for Disease Control and Prevention [198]. We compared interventions that would lower the incidence of prescription opioid misuse from 2016 to 2025 based on historical trends (a 7.5% reduction per year) and 50% faster than historical trends (an 11.3% reduction per year), vs a circumstance in which the incidence of misuse remained constant after 2015.

Under the status quo, the annual number of opioid overdose deaths is projected to increase from 33,100 in 2015 to 81,700 (95% uncertainty interval: 63,600 - 101,700) in 2025 (a 147% increase from 2015). From 2016 to 2025, 700,400 (95% UI: 590,200-

817,100) individuals in the United States are projected to die from opioid overdose, with 80% of the deaths attributable to illicit opioids. The number of individuals using illicit opioids is projected to increase by 61% - from 0.93 million (95% UI, 0.83-1.03 million) in 2015 to 1.50 million (95% UI, 0.98-2.22 million) by 2025. Across all interventions tested, further lowering the incidence of prescription opioid misuse from 2015 levels is projected to decrease overdose deaths by only 3.0% to 5.3%. Our findings suggest that interventions targeting prescription opioid misuse such as prescription monitoring programs may have a modest effect, at best, on the number of opioid overdose deaths in the near future. Additional policy interventions are urgently needed to change the course of the epidemic.

Chapter 6

Discussion

When I was young, I learned that my family suffers from a hereditary cancer syndrome. My grandmother was diagnosed with cancer at a young age, followed by an aunt, my mother, an uncle, and one of my first cousins. Due to advances in cancer therapy beginning in the 1970s, my grandmother, mother, and cousin were able to survive their disease. My grandmother's story in particular inspired me to study cancer. She was diagnosed with stage III breast cancer when she was in her early 40s in the 1970s. Her oncologist at Duke University told her that most patients with her disease don't survive more than a year. They did, however, offer her a place in a clinical trial for a new hormonal therapy called tamoxifen. She joined the clinical trial and lived another 40 years, getting to know all 10 of her grandchildren in the process. When my mother (also alive thanks to modern cancer care) told me that story, I knew that I would spend my career studying or treating cancer.

My first foray into cancer research was in a quantitative patient outcomes institute in Boston. While working in Boston, I became acutely aware that recent advances had not translated particularly well to improved patient outcomes. One of our functions was to assess new drugs based on their cost-effectiveness from the perspective of a health system or payer. We defined reasonable ranges of what drugs should cost based on standard willingness-to-pay thresholds. Put another way, we tried to estimate how much a drug should cost to deliver value to both patients and society. We then compared those estimates to what drug companies were actually charging. The results for new cancer drugs were not particularly encouraging. They typically extended life by a few months at great financial cost to the patient and society. I found myself extremely discouraged and jaded by the time I started thinking about pursuing research as a career prior to medical school. When people would ask me "Do you think we are close to curing cancer?" (a surprisingly common question when people find out you research cancer), I would respond with some variation of "Don't count on it". It was in this state of mind that I first encountered Professor Scott.

To my great surprise, Professor Scott and his lab members agreed with my skepticism about recent cancer advances. Unlike me, they had an explanation for why the precision medicine revolution hasn't cured cancer: evolution. While working with Professor Scott, I learned that the current maximum tolerated dose treatment paradigm for cancer fails to account for the capacity of these disease agents to evolve. When treatment fails to achieve rapid extinction, drug resistant clones rapidly proliferate into an uncontrollable tumor. I also adopted the lab's philosophy: To make significant progress for cancer patients, we need to better understand the evolutionary processes that drive cancer, and design treatments that explicitly account for them. The majority of my PhD has been dedicated to the final clause of that philosophy, the design of so-called "evolutionary therapies" that explicitly account for the capacity of cancer (and bacteria) to evolve resistance in response to drug therapy.

6.1 Summary of Contributions

In **Chapter 2**, we developed two novel methods to support precision targeting of tumors; a novel leave-one-out style method for node ranking and a novel algorithm for ranking miRNA combinations that maximizes tumor disruption while minimizing toxicity. As a proof of concept, we applied these methods to a set of Ewing Sarcoma cell lines. We identified candidate genes and miRNA cocktails that are predicted to be effective therapies in this setting.

In **Chapter 3**, we described crosstalkr, an open-source software package to facilitate interactomic analyses. In late 2020, I began to work on a few follow-up ideas from the miRNA study. We were (and are) particularly interested in adding a graph reduction step to the pipeline to make it more computationally tractable. When I couldn't find a suitable third-party R package to help facilitate disease-specific subnetwork identification, I decided to write one. Over time, crosstalkr grew to encompass all 3 of the most common steps in an interactomic pipeline; interaction with genome-scale PPI repositories, graph reduction, and node ranking.

In **Chapter 4**, we described a novel approach for designing evolutionary therapies that leverages reinforcement learning. We developed software that allows us to define and simulate model systems of evolution, apply any number of selective pressures (i.e. drugs), and finally solve for optimal drug cycling policies to minimize growth of the evolving population. Using this framework, we found that reinforcement learning agents were capable of superior control of a simulated population of E. coli given only information about instantaneous growth rate over time. We further introduced noise into the signal reaching the reinforcement learner, and show that the performance benefits do not decay until the signal to noise ratio is less than 0.5.

6.2 Next Steps

We are excited to validate the methods from **Chapter 2** further computationally and *in vitro*. I am working with Kristi Lin, a PhD student in Dr. Scott's lab to validate the node ranking method using a large drug sensitivity dataset that she is generating for the Ewing Sarcoma cell lines we studied. In addition, I am collaborating with Dr. Andrew Dhawan to design validation experiments for our miRNA combination therapy algorithm, using glioma as a model system. We plan to develop a pre-clinical model system of toxicity to assess our ability to identify non-toxic or less toxic miRNA combinations that still effectively inhibit glioma cell lines. I am also interested in improving the algorithm by using evolutionary algorithms rather than direct search. This change would allow us to search a much larger space of miRNA combinations, generalizing to any number of miRNA rather than being constrained to 3-miRNA combinations.

Crosstalkr (**Chapter 3**) has drawn significant interest from the bioinformatic community. In just a few months after we published the beta version of the software on CRAN (the central repository for R packages), crosstalkr has been downloaded more than 1,700 times. Crosstalkr is also up to 13 stars on github. I have no specific plans to extend the software, but am looking forward to maintaining it and collaborating with users to identify the most useful improvements.

In **Chapter 4**, we used reinforcement learning to explore how well an agent could learn drug cycling policies given less and less information about the system it is controlling. We have planned at least three studies that build on these foundations. While therapy optimization has received significant attention in the literature, most studies limit to so-called "perfect information" conditions where the optimizer understands complete system dynamics. Further, many studies constrain themselves to analytically tractable model systems. By instead relying on algorithms from the field of computer science (i.e. reinforcement learning), we can relax both constraints substantially.

In one future study, we will modify this framework to instead explore how well agents can learn to "speed up" evolution and drive populations to specific steady states more quickly. In another study, we will test the hypothesis that horizontal gene transfer makes drug cycling policies less effective. To do so, we will develop a stochastic system of evolution based on the Wright-Fisher model with additional terms to account for frequency-dependent recombination of alleles. We will further model an arbitrary number of alleles using a fitness-landscape based framework [199, 200]. While our model is expected to be analytically intractable, we can still solve for near-optimal drug cycling policies using methods to solve Markov decision-processes with continuous state space. Deep Q learning is one such method. Finally, we plan to validate the proposed methods *in vitro* using a high-throughput, massively parallel system of experimental evolution. In brief, we will evolve at least 1,000 replicates (10 or more 96-well plates) of E. coli under drug selection in parallel. After several two-week "episodes", we will batch train a reinforcement learner and then begin a period of on-policy learning where the reinforcement learner assumes control of the liquid handling robot and determines which drug should be applied to which replicate at each time point.

6.3 Conclusions

In 7 years in academics, I have been privileged to have exceptional mentors. In my post-baccalaureate work and then throughout my PhD, I was given nearly complete academic freedom to pursue my own ideas, make mistakes, and grow as a scientist and human. Through the completion of my PhD, I developed a passion for science that I know will drive me to conduct research throughout my career. I hope to become the PI of an NIH-funded laboratory focused on the design and validation of more intelligent approaches to cancer therapy.

Part IV

Bibliography

Bibliography

[1] J. G. Tate, S. Bamford, H. C. Jubb, Z. Sondka, D. M. Beare, N. Bindal, H. Boutselakis, C. G. Cole, C. Creatore, E. Dawson, P. Fish, B. Harsha, C. Hathaway, S. C. Jupe, C. Y. Kok, K. Noble, L. Ponting, C. C. Ramshaw, C. E. Rye, H. E. Speedy, R. Stefancsik, S. L. Thompson, S. Wang, S. Ward, P. J. Campbell, and S. A. Forbes, “COSMIC: the Catalogue Of Somatic Mutations In Cancer,” *Nucleic Acids Research*, vol. 47, pp. D941–D947, 1 2019.

[2] D. P. Nusinow, J. Szpyt, M. Ghandi, C. M. Rose, E. R. McDonald, M. Kalocsay, J. Jané-Valbuena, E. Gelfand, D. K. Schweppe, M. Jedrychowski, J. Golji, D. A. Porter, T. Rejtar, Y. K. Wang, G. V. Kryukov, F. Stegmeier, B. K. Erickson, L. A. Garraway, W. R. Sellers, and S. P. Gygi, “Quantitative Proteomics of the Cancer Cell Line Encyclopedia,” *Cell*, vol. 180, pp. 387–402.e16, jan 2020.

[3] D. K. Arrell and A. Terzic, “Network Systems Biology for Drug Discovery,” *Clinical Pharmacology & Therapeutics*, vol. 88, no. 1, pp. 120–125, 2010. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1038/clpt.2010.91.

[4] M. Koutrouli, E. Karatzas, D. Paez-Espino, and G. A. Pavlopoulos, “A Guide to Conquer the Biological Network Era Using Graph Theory,” *Frontiers in Bioengineering and Biotechnology*, vol. 8, 2020.

[5] R. L. Costa, M. Boroni, and M. A. Soares, “Distinct co-expression networks using multi-omic data reveal novel interventional targets in HPV-positive and negative head-and-neck squamous cell cancer,” *Scientific Reports*, vol. 8, p. 15254, Oct. 2018. Number: 1 Publisher: Nature Publishing Group.

[6] D. Vella, I. Zoppis, G. Mauri, P. Mauri, and D. Di Silvestre, “From protein-protein interactions to protein co-expression networks: a new perspective to evaluate large-scale proteomic data,” *EURASIP Journal on Bioinformatics and Systems Biology*, vol. 2017, no. 1, p. 6, 2017.

[7] J. A. Scarborough, S. A. Eschrich, J. Torres-Roca, A. Dhawan, and J. G. Scott, “Exploiting convergent evolution to derive a pan-cancer cisplatin response gene expression signature,” *medRxiv*, 2022.

[8] F. M. Buffa, A. L. Harris, C. M. West, and C. J. Miller, “Large meta-analysis of multiple cancers reveals a common, compact and highly prognostic hypoxia metagene,” *British Journal of Cancer*, vol. 102, pp. 428–435, Jan. 2010. Number: 2 Publisher: Nature Publishing Group.

[9] A. K. Chavali, K. M. D’Auria, E. L. Hewlett, R. D. Pearson, and J. A. Papin, “A metabolic network approach for the identification and prioritization of antimicrobial drug targets,” *Trends in Microbiology*, vol. 20, pp. 113–123, Mar. 2012.

[10] R. Caspi, R. Billington, I. M. Keseler, A. Kothari, M. Krummenacker, P. E. Midford, W. K. Ong, S. Paley, P. Subhraveti, and P. D. Karp, "The MetaCyc database of metabolic pathways and enzymes—a 2019 update," *Nucleic Acids Research*, vol. 48, no. D1, pp. d445–d453, 2019.

[11] V. Martínez, C. Navarro, C. Cano, W. Fajardo, and A. Blanco, "DrugNet: Network-based drug–disease prioritization by integrating heterogeneous data," *Artificial Intelligence in Medicine*, vol. 63, pp. 41–49, Jan. 2015.

[12] R. Oughtred, J. Rust, C. Chang, B. Breitkreutz, C. Stark, A. Willems, L. Boucher, G. Leung, N. Kolas, F. Zhang, S. Dolma, J. Coulombe-Huntington, A. Chatr-aryamontri, K. Dolinski, and M. Tyers, "The BioGRID database: A comprehensive biomedical resource of curated protein, genetic, and chemical interactions," *Protein Science : A Publication of the Protein Society*, vol. 30, pp. 187–200, Jan. 2021.

[13] D. Szklarczyk, A. L. Gable, K. C. Nastou, D. Lyon, R. Kirsch, S. Pyysalo, N. T. Doncheva, M. Legeay, T. Fang, P. Bork, L. J. Jensen, and C. von Mering, "The STRING database in 2021: customizable protein-protein networks, and functional characterization of user-uploaded gene/measurement sets," *Nucleic Acids Research*, vol. 49, pp. D605–D612, Jan. 2021.

[14] C. Y. Kim, S. Baek, J. Cha, S. Yang, E. Kim, E. M. Marcotte, T. Hart, and I. Lee, "HumanNet v3: an improved database of human gene networks for disease research," *Nucleic Acids Research*, vol. 50, no. D1, pp. d632–d639, 2021.

[15] D. Szklarczyk, A. L. Gable, D. Lyon, A. Junge, S. Wyder, J. Huerta-Cepas, M. Simonovic, N. T. Doncheva, J. H. Morris, P. Bork, L. J. Jensen, and C. Von Mering, "STRING v11: Protein-protein association networks with increased coverage, supporting functional discovery in genome-wide experimental datasets," *Nucleic Acids Research*, vol. 47, pp. D607–D613, jan 2019.

[16] R. J. Wilson, *Introduction to graph theory*. Harlow Munich: Prentice Hall, 4. ed., [nachdr.] ed., 2009.

[17] L. Salwinski, C. S. Miller, A. J. Smith, F. K. Pettit, J. U. Bowie, and D. Eisenberg, "The Database of Interacting Proteins: 2004 update," *Nucleic Acids Research*, vol. 32, pp. D449–D451, Jan. 2004.

[18] P. Pagel, S. Kovac, M. Oesterheld, B. Brauner, I. Dunger-Kaltenbach, G. Frishman, C. Montrone, P. Mark, V. Stümpflen, H.-W. Mewes, A. Ruepp, and D. Frishman, "The MIPS mammalian protein–protein interaction database," *Bioinformatics*, vol. 21, pp. 832–834, Mar. 2005.

[19] S. Razick, G. Magklaras, and I. M. Donaldson, "iRefIndex: A consolidated protein interaction database with provenance," *BMC Bioinformatics*, vol. 9, p. 405, Sept. 2008.

[20] D. M. Camacho, K. M. Collins, R. K. Powers, J. C. Costello, and J. J. Collins, "Next-Generation Machine Learning for Biological Networks," *Cell*, vol. 173, no. 7, pp. 1581–1592, 2018.

[21] S. Orchard, S. Kerrien, S. Abbani, B. Aranda, J. Bhate, S. Bidwell, A. Bridge, L. Briganti, F. S. L. Brinkman, G. Cesareni, A. Chatr-aryamontri, E. Chautard, C. Chen, M. Dumousseau, J. Goll, R. E. W. Hancock, L. I. Hannick, I. Jurisica, J. Khadake, D. J. Lynn, U. Mahadevan, L. Perfetto, A. Raghunath, S. Ricard-Blum, B. Roechert, L. Salwinski, V. Stümpflen, M. Tyers, P. Uetz, I. Xenarios, and H. Hermjakob, "Protein interaction data curation: the International Molecular Exchange (IMEx) consortium," *Nature Methods*, vol. 9, no. 4, pp. 345–350, 2012.

[22] S. Orchard, M. Ammari, B. Aranda, L. Breuza, L. Briganti, F. Broackes-Carter, N. H. Campbell, G. Chavali, C. Chen, N. del Toro, M. Duesbury, M. Dumousseau, E. Galeota, U. Hinz, M. Iannuccelli, S. Jagannathan, R. Jimenez, J. Khadake, A. Lagreid, L. Licata, R. C. Lovering, B. Meldal, A. N. Melidoni, M. Milagros, D. Peluso, L. Perfetto, P. Porras, A. Raghunath, S. Ricard-Blum, B. Roechert, A. Stutz, M. Tognolli, K. van Roey, G. Cesareni, and H. Hermjakob, "The MIntAct project–IntAct as a common curation platform for 11 molecular interaction databases," *Nucleic Acids Research*, vol. 42, pp. D358–363, Jan. 2014.

[23] J. Snider, M. Kotlyar, P. Saraon, Z. Yao, I. Jurisica, and I. Stagljar, "Fundamentals of protein interaction network mapping," *Molecular Systems Biology*, vol. 11, p. 848, Dec. 2015. Publisher: John Wiley & Sons, Ltd.

[24] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, and J. Kang, "BioBERT: a pre-trained biomedical language representation model for biomedical text mining," *Bioinformatics*, vol. 36, no. 4, pp. 1234–1240, 2019.

[25] S. Van Landeghem, J. Björne, C.-H. Wei, K. Hakala, S. Pyysalo, S. Ananiadou, H.-Y. Kao, Z. Lu, T. Salakoski, Y. Van de Peer, and F. Ginter, "Large-Scale Event Extraction from Literature with Multi-Level Gene Normalization," *PLOS ONE*, vol. 8, no. 4, p. e55814, 2013.

[26] C.-H. Wei, H.-Y. Kao, and Z. Lu, "PubTator: a web-based text mining tool for assisting biocuration," *Nucleic Acids Research*, vol. 41, no. W1, pp. w518–w522, 2013.

[27] E. Persson, M. Castresana-Aguirre, D. Buzzao, D. Guala, and E. L. L. Sonnhammer, "FunCoup 5: Functional Association Networks in All Domains of Life, Supporting

Directed Links and Tissue-Specificity," *Journal of Molecular Biology*, vol. 433, p. 166835, May 2021.

[28] Q. C. Zhang, D. Petrey, L. Deng, L. Qiang, Y. Shi, C. A. Thu, B. Bisikirska, C. Lefebvre, D. Accili, T. Hunter, T. Maniatis, A. Califano, and B. Honig, "Structure-based prediction of protein–protein interactions on a genome-wide scale," *Nature*, vol. 490, pp. 556–560, Oct. 2012. Number: 7421 Publisher: Nature Publishing Group.

[29] C. Chen, Q. Zhang, Q. Ma, and B. Yu, "LightGBM-PPI: Predicting protein-protein interactions through LightGBM with multi-information fusion," *Chemometrics and Intelligent Laboratory Systems*, vol. 191, pp. 54–64, Aug. 2019.

[30] C. Chen, Q. Zhang, B. Yu, Z. Yu, P. J. Lawrence, Q. Ma, and Y. Zhang, "Improving protein-protein interactions prediction accuracy using XGBoost feature selection and stacked ensemble classifier," *Computers in Biology and Medicine*, vol. 123, p. 103899, Aug. 2020.

[31] B. Tian, X. Wu, C. Chen, W. Qiu, Q. Ma, and B. Yu, "Predicting protein–protein interactions by fusing various Chou's pseudo components and using wavelet denoising approach," *Journal of Theoretical Biology*, vol. 462, pp. 329–346, Feb. 2019.

[32] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back, S. Petersen, D. Reiman, E. Clancy, M. Zielinski, M. Steinegger, M. Pacholska, T. Berghammer, S. Bodenstein, D. Silver, O. Vinyals, A. W. Senior, K. Kavukcuoglu, P. Kohli, and D. Hassabis, "Highly accurate protein structure prediction with AlphaFold," *Nature*, vol. 596, pp. 583–589, Aug. 2021. Number: 7873 Publisher: Nature Publishing Group.

[33] P. Bryant, G. Pozzati, and A. Elofsson, "Improved prediction of protein-protein interactions using AlphaFold2," *Nature Communications*, vol. 13, p. 1265, Mar. 2022. Number: 1 Publisher: Nature Publishing Group.

[34] "Pathguide: the pathway resource list."

[35] T. S. Keshava Prasad, R. Goel, K. Kandasamy, S. Keerthikumar, S. Kumar, S. Mathivanan, D. Telikicherla, R. Raju, B. Shafreen, A. Venugopal, L. Balakrishnan, A. Marimuthu, S. Banerjee, D. S. Somanathan, A. Sebastian, S. Rani, S. Ray, C. J. Harrys Kishore, S. Kanth, M. Ahmed, M. K. Kashyap, R. Mohmood, Y. L. Ramachandra, V. Krishna, B. A. Rahiman, S. Mohan, P. Ranganathan,

S. Ramabadran, R. Chaerkady, and A. Pandey, “Human Protein Reference Database–2009 update,” *Nucleic Acids Research*, vol. 37, pp. D767–772, Jan. 2009.

[36] A. Chatr-aryamontri, A. Ceol, L. M. Palazzi, G. Nardelli, M. V. Schneider, L. Castagnoli, and G. Cesareni, “MINT: the Molecular INTeraction database,” *Nucleic Acids Research*, vol. 35, pp. D572–D574, Jan. 2007.

[37] R. Oughtred, C. Stark, B.-J. Breitkreutz, J. Rust, L. Boucher, C. Chang, N. Kolas, L. O’Donnell, G. Leung, R. McAdam, F. Zhang, S. Dolma, A. Willems, J. Coulombe-Huntington, A. Chatr-aryamontri, K. Dolinski, and M. Tyers, “The BioGRID interaction database: 2019 update,” *Nucleic Acids Research*, vol. 47, pp. D529–D541, 1 2019.

[38] S. A. Kennedy, M.-A. Jarboui, S. Srihari, C. Raso, K. Bryan, L. Dernayka, T. Charitou, M. Bernal-Llinares, C. Herrera-Montavez, A. Krstic, D. Matallanas, M. Kotlyar, I. Jurisica, J. Curak, V. Wong, I. Stagljar, T. LeBihan, L. Imrie, P. Pillai, M. A. Lynn, E. Fasterius, C. Al-Khalili Szigyarto, J. Breen, C. Kiel, L. Serrano, N. Rauch, O. Rukhlenko, B. N. Kholodenko, L. F. Iglesias-Martinez, C. J. Ryan, R. Pilkington, P. Cammareri, O. Sansom, S. Shave, M. Auer, N. Horn, F. Klose, M. Ueffing, K. Boldt, D. J. Lynn, and W. Kolch, “Extensive rewiring of the EGFR network in colorectal cancer cells expressing transforming levels of KRASG13D,” *Nature Communications*, vol. 11, p. 499, Jan. 2020. Number: 1 Publisher: Nature Publishing Group.

[39] D. T. Weaver, K. I. Pishas, D. Williamson, J. Scarborough, S. L. Lessnick, A. Dhawan, and J. G. Scott, “Network potential identifies therapeutic miRNA cocktails in Ewing sarcoma,” *PLOS Computational Biology*, vol. 17, p. e1008755, Oct. 2021. Publisher: Public Library of Science.

[40] E. A. Rietman, J. G. Scott, J. A. Tuszynski, and G. L. Klement, “Personalized anticancer therapy selection using molecular landscape topology and thermodynamics,” *Oncotarget*, vol. 8, pp. 18735–18745, Oct. 2016.

[41] B. Pfeifer, H. Baniecki, A. Saranti, P. Biecek, and A. Holzinger, “Multi-omics disease module detection with an explainable Greedy Decision Forest,” *Scientific Reports*, vol. 12, p. 16857, Oct. 2022. Number: 1 Publisher: Nature Publishing Group.

[42] C. Durón, Y. Pan, D. H. Gutmann, J. Hardin, and A. Radunskaya, “Variability of Betweenness Centrality and Its Effect on Identifying Essential Genes,” *Bulletin of Mathematical Biology*, vol. 81, pp. 3655–3673, Sept. 2019.

[43] H. Keane, B. J. Ryan, B. Jackson, A. Whitmore, and R. Wade-Martins, “Protein-protein interaction networks identify targets which rescue the MPP+ cellular model

of Parkinson's disease," *Scientific Reports*, vol. 5, p. 17004, Nov. 2015. Number: 1 Publisher: Nature Publishing Group.

[44] H. Azevedo and C. A. Moreira-Filho, "Topological robustness analysis of protein interaction networks reveals key targets for overcoming chemotherapy resistance in glioma," *Scientific Reports*, vol. 5, p. 16830, Nov. 2015. Number: 1 Publisher: Nature Publishing Group.

[45] R. Uddin and F. Jamil, "Prioritization of potential drug targets against P. aeruginosa by core proteomic analysis using computational subtractive genomics and Protein-Protein interaction network," *Computational Biology and Chemistry*, vol. 74, pp. 115–122, June 2018.

[46] X. Tang, J. Wang, J. Zhong, and Y. Pan, "Predicting Essential Proteins Based on Weighted Degree Centrality," *IEEE/ACM transactions on computational biology and bioinformatics*, vol. 11, no. 2, pp. 407–418, 2014.

[47] S. M. Almasi and T. Hu, "Measuring the importance of vertices in the weighted human disease network," *PloS One*, vol. 14, no. 3, p. e0205936, 2019.

[48] H. Jeong, S. P. Mason, A.-L. Barabási, and Z. N. Oltvai, "Lethality and centrality in protein networks," *Nature*, vol. 411, pp. 41–42, May 2001. Number: 6833 Publisher: Nature Publishing Group.

[49] T.-D. Tran and Y.-K. Kwon, "Hierarchical closeness efficiently predicts disease genes in a directed signaling network," *Computational Biology and Chemistry*, vol. 53, pp. 191–197, Dec. 2014.

[50] D. Mistry, R. P. Wise, and J. A. Dickerson, "DiffSLC: A graph centrality method to detect essential proteins of a protein-protein interaction network," *PLOS ONE*, vol. 12, p. e0187091, Nov. 2017. Publisher: Public Library of Science.

[51] M. Li, J. Wang, X. Chen, H. Wang, and Y. Pan, "A local average connectivity-based method for identifying essential proteins from the network level," *Computational Biology and Chemistry*, vol. 35, pp. 143–150, June 2011.

[52] K. A. Ahmed, P. Chinnaiyan, W. J. Fulp, S. Eschrich, J. F. Torres-Roca, and J. J. Caudell, "The radiosensitivity index predicts for overall survival in glioblastoma," *Oncotarget*, vol. 6, pp. 34414–34422, Oct. 2015.

[53] S.-H. Yeh, H.-Y. Yeh, and V.-W. Soo, "A network flow approach to predict drug targets from microarray data, disease genes and interactome network - case study on prostate cancer," *Journal of Clinical Bioinformatics*, vol. 2, p. 1, Jan. 2012.

[54] S. Köhler, S. Bauer, D. Horn, and P. N. Robinson, "Walking the Interactome for Prioritization of Candidate Disease Genes," *The American Journal of Human Genetics*, vol. 82, pp. 949–958, Apr. 2008.

[55] R. K. Nibbe, M. Koyutürk, and M. R. Chance, "An Integrative -omics Approach to Identify Functional Sub-Networks in Human Colorectal Cancer," *PLOS Computational Biology*, vol. 6, p. e1000639, Jan. 2010. Publisher: Public Library of Science.

[56] O. Vanunu and R. Sharan, *A Propagation-based Algorithm for Inferring Gene-Disease Associations*. Gesellschaft für Informatik e. V., 2008. Accepted: 2019-04-03T12:07:34Z.

[57] V. Martínez, C. Cano, and A. Blanco, "ProphNet: a generic prioritization method through propagation of information," *BMC bioinformatics*, vol. 15 Suppl 1, p. S5, 2014.

[58] N. Sambaturu, V. Pusadkar, S. Hannenhalli, and N. Chandra, "PathExt: a general framework for path-based mining of omics-integrated biological networks," *Bioinformatics*, vol. 37, pp. 1254–1262, May 2021.

[59] A. Sadeghi and H. Fröhlich, "Steiner tree methods for optimal sub-network identification: an empirical study," *BMC Bioinformatics*, vol. 14, p. 144, Apr. 2013.

[60] N. Carels, "Optimization of combination chemotherapy based on the calculation of network entropy for protein-protein interactions in breast cancer cell lines," 2015.

[61] D. Breitkreutz, L. Hlatky, E. Rietman, and J. A. Tuszynski, "Molecular signaling network complexity is correlated with cancer patient survivability," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 109, pp. 9209–9212, June 2012.

[62] S. Benzekry, J. A. Tuszynski, E. A. Rietman, and G. Lakka Klement, "Design principles for cancer therapy guided by changes in complexity of protein-protein interaction networks," *Biology Direct*, vol. 10, p. 32, May 2015.

[63] C. Yu, E. A. Rietman, H. T. Siegelmann, M. Cavaglia, and J. A. Tuszynski, "Application of Thermodynamics and Protein–Protein Interaction Network Topology for Discovery of Potential New Treatments for Temporal Lobe Epilepsy," *Applied Sciences*, vol. 11, p. 8059, Jan. 2021. Number: 17 Publisher: Multidisciplinary Digital Publishing Institute.

[64] T. Zeng, D. C. Wang, X. Wang, F. Xu, and L. Chen, "Prediction of dynamical drug sensitivity and resistance by module network rewiring-analysis based on transcriptional profiling," *Drug Resistance Updates*, vol. 17, pp. 64–76, July 2014.

[65] B. Pfeifer, A. Saranti, and A. Holzinger, "GNN-SubNet: disease subnetwork detection with explainable graph neural networks," *Bioinformatics*, vol. 38, pp. ii120–ii126, Sept. 2022.

[66] R. Schulte-Sasse, S. Budach, D. Hnisz, and A. Marsico, "Integration of multiomics data with graph convolutional networks to identify new cancer genes and their associated molecular mechanisms," *Nature Machine Intelligence*, vol. 3, pp. 513–526, June 2021. Number: 6 Publisher: Nature Publishing Group.

[67] Q. Liu, Z. Hu, R. Jiang, and M. Zhou, "DeepCDR: a hybrid graph convolutional network for predicting cancer drug response," *Bioinformatics*, vol. 36, pp. i911–i918, Dec. 2020.

[68] S. Choobdar, M. E. Ahsen, J. Crawford, M. Tomasoni, T. Fang, D. Lamparter, J. Lin, B. Hescott, X. Hu, J. Mercer, T. Natoli, R. Narayan, A. Subramanian, J. D. Zhang, G. Stolovitzky, Z. Kutalik, K. Lage, D. K. Slonim, J. Saez-Rodriguez, L. J. Cowen, S. Bergmann, and D. Marbach, "Assessment of network module identification across complex diseases," *Nature Methods*, vol. 16, no. 9, pp. 843–852, 2019.

[69] A. N. Langville and C. D. Meyer, "Deeper Inside PageRank," *Internet Mathematics*, vol. 1, pp. 335–380, Jan. 2004. Publisher: Taylor & Francis _eprint: https://www.tandfonline.com/doi/pdf/10.1080/15427951.2004.10129091.

[70] H. Tong, C. Faloutsos, and J.-y. Pan, "Fast Random Walk with Restart and Its Applications," in *Sixth International Conference on Data Mining (ICDM'06)*, pp. 613–622, Dec. 2006. ISSN: 2374-8486.

[71] M.-N. I. of Genetic Medicine, "Online Mendelian Inheritance in Man, OMIM®."

[72] R. Albert, H. Jeong, and A. L. Barabási, "Error and attack tolerance of complex networks," *Nature*, vol. 406, pp. 378–382, 7 2000.

[73] J. G. Scott, G. Sedor, P. Ellsworth, J. A. Scarborough, K. A. Ahmed, D. E. Oliver, S. A. Eschrich, M. W. Kattan, and J. F. Torres-Roca, "Pan-cancer prediction of radiotherapy benefit using genomic-adjusted radiation dose (GARD): a cohort-based pooled analysis," *The Lancet Oncology*, vol. 22, pp. 1221–1229, Sept. 2021. Publisher: Elsevier.

[74] E. Rietman and J. A. Tuszynski, "Using Thermodynamic Functions as an Organizing Principle in Cancer Biology," in *Theoretical and Applied Aspects of Systems Biology*, pp. 139–157, Springer, Cham, 2018.

[75] E. A. Rietman, S. Taylor, H. T. Siegelmann, M. A. Deriu, M. Cavaglia, and J. A. Tuszynski, "Using the Gibbs Function as a Measure of Human Brain Development

Trends from Fetal Stage to Advanced Age," *International Journal of Molecular Sciences*, vol. 21, p. 1116, Jan. 2020. Number: 3 Publisher: Multidisciplinary Digital Publishing Institute.

[76] M. Kotlyar, C. Pastrello, Z. Ahmed, J. Chee, Z. Varyova, and I. Jurisica, "IID 2021: towards context-specific protein interaction analyses by increased coverage, enhanced annotation and enrichment analysis," *Nucleic Acids Research*, vol. 50, no. D1, pp. d640–d647, 2021.

[77] C. Mackintosh, J. Madoz-Gúrpide, J. L. Ordóñez, D. Osuna, and D. Herrero-Martín, "The molecular pathogenesis of Ewing's sarcoma.," *Cancer biology & therapy*, vol. 9, pp. 655–67, 5 2010.

[78] N. Esiashvili, M. Goodman, and R. B. Marcus, "Changes in Incidence and Survival of Ewing Sarcoma Patients Over the Past 3 Decades," *Journal of Pediatric Hematology/Oncology*, vol. 30, pp. 425–430, 6 2008.

[79] E. R. Lawlor and P. H. Sorensen, "Twenty Years On – What Do We Really Know About Ewing Sarcoma And What Is The Path Forward?," *Critical reviews in oncogenesis*, vol. 20, no. 0, p. 155, 2015.

[80] N. Riggi, M.-L. Suvà, C. De Vito, P. Provero, J.-C. Stehle, K. Baumer, L. Cironi, M. Janiszewska, T. Petricevic, D. Suvà, S. Tercier, J.-M. Joseph, L. Guillou, and I. Stamenkovic, "EWS-FLI-1 modulates miRNA145 and SOX2 expression to initiate mesenchymal stem cell reprogramming toward Ewing sarcoma cancer stem cells.," *Genes & development*, vol. 24, pp. 916–32, 5 2010.

[81] R. Rupaimoole and F. J. Slack, "MicroRNA therapeutics: towards a new era for the management of cancer and other diseases," *Nature Reviews Drug Discovery*, vol. 16, pp. 203–222, 3 2017.

[82] A. Dhawan, J. G. Scott, A. L. Harris, and F. M. Buffa, "Pan-cancer characterisation of microRNA across cancer hallmarks reveals microRNA-mediated downregulation of tumour suppressors," *Nature Communications*, vol. 9, p. 5228, 12 2018.

[83] D. Hanahan and R. A. Weinberg, "The hallmarks of cancer," *cell*, vol. 100, no. 1, pp. 57–70, 2000.

[84] X. LIN, Z. YANG, P. ZHANG, and G. SHAO, "miR-154 suppresses non-small cell lung cancer growth in vitro and in vivo," *Oncology Reports*, vol. 33, pp. 3053–3060, 6 2015.

[85] K. Lin, M. Farahani, Y. Yang, G. G. Johnson, M. Oates, M. Atherton, A. Douglas, N. Kalakonda, and A. R. Pettitt, "Loss of <i>MIR15A</i> and <i>MIR16-1</i> at 13q14 is associated with increased <i>TP53</i> mRNA, de-repression

of <i>BCL2</i> and adverse outcome in chronic lymphocytic leukaemia," *British Journal of Haematology*, vol. 167, pp. 346–355, 11 2014.

[86] M. J. Reilley, P. McCoon, C. Cook, P. Lyne, R. Kurzrock, Y. Kim, R. Woessner, A. Younes, J. Nemunaitis, N. Fowler, M. Curran, Q. Liu, T. Zhou, J. Schmidt, M. Jo, S. J. Lee, M. Yamashita, S. G. Hughes, L. Fayad, S. Piha-Paul, M. V. P. Nadella, X. Xiao, J. Hsu, A. Revenko, B. P. Monia, A. R. MacLeod, and D. S. Hong, "STAT3 antisense oligonucleotide AZD9150 in a subset of patients with heavily pretreated lymphoma: results of a phase 1b trial," *Journal for ImmunoTherapy of Cancer*, vol. 6, p. 119, 12 2018.

[87] D. Hong, R. Kurzrock, Y. Kim, R. Woessner, A. Younes, J. Nemunaitis, N. Fowler, T. Zhou, J. Schmidt, M. Jo, S. J. Lee, M. Yamashita, S. G. Hughes, L. Fayad, S. Piha-Paul, M. V. P. Nadella, M. Mohseni, D. Lawson, C. Reimer, D. C. Blakey, X. Xiao, J. Hsu, A. Revenko, B. P. Monia, and A. R. MacLeod, "AZD9150, a next-generation antisense oligonucleotide inhibitor of <i>STAT3</i> with early evidence of clinical activity in lymphoma and lung cancer," *Science Translational Medicine*, vol. 7, pp. 185–314, 11 2015.

[88] S. Odate, V. Veschi, S. Yan, N. Lam, R. Woessner, and C. J. Thiele, "Inhibition of <i>STAT3</i> with the Generation 2.5 Antisense Oligonucleotide, AZD9150, Decreases Neuroblastoma Tumorigenicity and Increases Chemosensitivity," *Clinical Cancer Research*, vol. 23, pp. 1771–1784, 4 2017.

[89] A. L. Kasinski and F. J. Slack, "miRNA-34 prevents cancer initiation and progression in a therapeutically resistant K-ras and p53-induced mouse model of lung adenocarcinoma.," *Cancer research*, vol. 72, pp. 5576–87, 11 2012.

[90] J. F. Wiggins, L. Ruffino, K. Kelnar, M. Omotola, L. Patrawala, D. Brown, and A. G. Bader, "Development of a Lung Cancer Therapeutic Based on the Tumor Suppressor MicroRNA-34," *Cancer Research*, vol. 70, pp. 5923–5930, 7 2010.

[91] C. Stahlhut and F. J. Slack, "Combinatorial Action of MicroRNAs <i>let-7</i> and miR-34 Effectively Synergizes with Erlotinib to Suppress Non-small Cell Lung Cancer Cell Proliferation," *Cell Cycle*, vol. 14, pp. 2171–2180, 7 2015.

[92] C. Liu, K. Kelnar, B. Liu, X. Chen, T. Calhoun-Davis, H. Li, L. Patrawala, H. Yan, C. Jeter, S. Honorio, J. F. Wiggins, A. G. Bader, R. Fagin, D. Brown, and D. G. Tang, "The microRNA miR-34a inhibits prostate cancer stem cells and metastasis by directly repressing CD44.," *Nature medicine*, vol. 17, pp. 211–5, 2 2011.

[93] T. I. S. Maciej Pajak, "miRNAtap: miRNAtap: microRNA Targets— Aggregated Predictions.," 2019.

[94] N. Wong and X. Wang, "miRDB: an online resource for microRNA target prediction and functional annotations," *Nucleic Acids Research*, vol. 43, pp. D146–D152, 1 2015.

[95] M. Maragkakis, M. Reczko, V. A. Simossis, P. Alexiou, G. L. Papadopoulos, T. Dalamagas, G. Giannopoulos, G. Goumas, E. Koukis, K. Kourtis, T. Vergoulis, N. Koziris, T. Sellis, P. Tsanakas, and A. G. Hatzigeorgiou, "DIANA-microT web server: elucidating microRNA functions through target prediction," *Nucleic Acids Research*, vol. 37, pp. W273–W276, 7 2009.

[96] V. O'Neill, "A Multicenter Phase I Study of MRX34, MicroRNA miR-RX34 Liposomal Injection - Full Text View - ClinicalTrials.gov."

[97] X. Lai, M. Eberhardt, U. Schmitz, and J. Vera, "Systems biology-based investigation of cooperating microRNAs as monotherapy or adjuvant therapy in cancer," 2019.

[98] E. A. Rietman, J. G. Scott, J. A. Tuszynski, and G. L. Klement, "Personalized anti-cancer therapy selection using molecular landscape topology and thermodynamics.," *Oncotarget*, vol. 8, pp. 18735–18745, 3 2017.

[99] E. A. Rietman, J. Platig, J. A. Tuszynski, and G. Lakka Klement, "Thermodynamic measures of cancer: Gibbs free energy and entropy of protein-protein interactions.," *Journal of biological physics*, vol. 42, no. 3, pp. 339–50, 2016.

[100] G. Stelzer, N. Rosen, I. Plaschkes, S. Zimmerman, M. Twik, S. Fishilevich, T. I. Stein, R. Nudel, I. Lieder, Y. Mazor, S. Kaplan, D. Dahary, D. Warshawsky, Y. Guan-Golan, A. Kohn, N. Rappaport, M. Safran, and D. Lancet, "The GeneCards Suite: From Gene Data Mining to Disease Genome Sequence Analyses," in *Current Protocols in Bioinformatics*, vol. 54, pp. 1–1, Hoboken, NJ, USA: John Wiley & Sons, Inc., 6 2016.

[101] A. Liberzon, A. Subramanian, R. Pinchback, H. Thorvaldsdottir, P. Tamayo, and J. P. Mesirov, "Molecular signatures database (MSigDB) 3.0," *Bioinformatics*, vol. 27, pp. 1739–1740, 6 2011.

[102] A. Liberzon, C. Birger, H. Thorvaldsdóttir, M. Ghandi, J. P. Mesirov, and P. Tamayo, "The Molecular Signatures Database Hallmark Gene Set Collection," *Cell Systems*, vol. 1, pp. 417–425, 12 2015.

[103] G. Korotkevich, V. Sukhov, and A. A. Sergushichev, "An algorithm for fast preranked gene set enrichment analysis using cumulative statistic calculation," *bioRxiv*, p. 060012, 2016.

[104] Y. Benjamini, Yoav ; Hochberg, "Controlling the False Discovery Rate - a Practical and Powerful Approach to Multiple Testing. Journal of the Royal Statistical Society Series B-Methodological 1995.pdf," *Journal of the Royal Statistical Society Series B (Methological)*, vol. 57, no. 1, pp. 289–300, 1995.

[105] C. McLeod, A. M. Gout, X. Zhou, A. Thrasher, D. Rahbarinia, S. W. Brady, M. Macias, K. Birch, D. Finkelstein, J. Sunny, R. Mudunuri, B. A. Orr, M. Treadway, B. Davidson, T. K. Ard, A. Chiao, A. Swistak, S. Wiggins, S. Foy, J. Wang, E. Sioson, S. Wang, J. R. Michael, Y. Liu, X. Ma, A. Patel, M. N. Edmonson, M. R. Wilkinson, A. M. Frantz, T.-C. Chang, L. Tian, S. Lei, S. A. Islam, C. Meyer, N. Thangaraj, P. Tater, V. Kandali, S. Ma, T. Nguyen, O. Serang, I. McGuire, N. Robison, D. Gentry, X. Tang, L. E. Palmer, G. Wu, E. Suh, L. Tanner, J. McMurry, M. Lear, A. S. Pappo, Z. Wang, C. L. Wilson, Y. Cheng, S. Meshinchi, L. B. Alexandrov, M. J. Weiss, G. T. Armstrong, L. L. Robison, Y. Yasui, K. E. Nichols, D. W. Ellison, C. Bangur, C. G. Mullighan, S. J. Baker, M. A. Dyer, G. Miller, S. Newman, M. Rusch, R. Daly, K. Perry, J. R. Downing, and J. Zhang, "St. Jude Cloud: A Pediatric Cancer Genomic Data-Sharing Ecosystem," *Cancer Discovery*, vol. 11, pp. 1082–1099, may 2021.

[106] J. L. Goodier, L. E. Cheung, and H. H. Kazazian, "MOV10 RNA Helicase Is a Potent Inhibitor of Retrotransposition in Cells," *PLoS Genetics*, vol. 8, p. 1002941, oct 2012.

[107] M. Hameiri-Grossman, A. Porat-Klein, I. Yaniv, S. Ash, I. J. Cohen, Y. Kodman, R. Haklai, G. Elad-Sfadia, Y. Kloog, E. Chepurko, M. Feinmesser, J. Issakov, O. Sher, D. Luria, Y. Kollender, A. Weizman, and S. Avigad, "The association between let-7, RAS and HIF-1 in Ewing Sarcoma tumor growth," *Oncotarget*, vol. 6, pp. 33834–48, 10 2015.

[108] M. Guo, X. Zhao, X. Yuan, J. Jiang, and P. Li, "MiR-let-7a inhibits cell proliferation, migration, and invasion by down-regulating PKM2 in cervical cancer," *Oncotarget*, vol. 8, pp. 28226–28236, 4 2017.

[109] R. Tang, C. Yang, X. Ma, Y. Wang, D. Luo, C. Huang, Z. Xu, P. Liu, and L. Yang, "MiR-let-7a inhibits cell proliferation, migration, and invasion by down-regulating PKM2 in gastric cancer.," *Oncotarget*, vol. 7, pp. 5972–84, 2 2016.

[110] X.-M. Pan, R.-F. Sun, Z.-H. Li, X.-M. Guo, Z. Zhang, H.-J. Qin, G.-H. Xu, and L.-B. Gao, "A let-7 KRAS rs712 polymorphism increases colorectal cancer risk," *Tumor Biology*, vol. 35, pp. 831–835, 1 2014.

[111] R. Wang, Z. Yu, F. Chen, H. Xu, S. Shen, W. Chen, L. Chen, Q. Su, L. Zhang, J. Bi, W. Zeng, W. Li, X. Huang, and Q. Wang, "miR-300 regulates the epithelial-mesenchymal transition and invasion of hepatocellular carcinoma by targeting the

FAK/PI3K/AKT signaling pathway," *Biomedicine & Pharmacotherapy*, vol. 103, pp. 1632–1642, 7 2018.

[112] H.-q. Liang, R.-j. Wang, C.-f. Diao, J.-w. Li, J.-l. Su, and S. Zhang, "The PTTG1-targeting miRNAs miR-329, miR-300, miR-381, and miR-655 inhibit pituitary tumor cell tumorigenesis and are involved in a p53/PTTG1 regulation feedback loop.," *Oncotarget*, vol. 6, pp. 29413–27, 10 2015.

[113] J. Wang, S. Wang, J. Zhou, and Q. Qian, "miR-424-5p regulates cell proliferation, migration and invasion by targeting doublecortin-like kinase 1 in basal-like breast cancer," *Biomedicine & Pharmacotherapy*, vol. 102, pp. 147–152, 6 2018.

[114] G.-Q. J. J. Zhao, "MiR-4282 inhibits proliferation, invasion and metastasis of human breast cancer by targeting Myc," *European review for medical and pharmacological sciences*, 2018.

[115] C. Alderman and Y. Yang, "The anti-melanoma activity and oncogenic targets of hsa-miR-15a-5p," *RNA & disease (Houston, Tex.)*, vol. 3, no. 4, 2016.

[116] M. Rohini, M. Gokulnath, P. Miranda, and N. Selvamurugan, "miR-590–3p inhibits proliferation and promotes apoptosis by targeting activating transcription factor 3 in human breast cancer cells," *Biochimie*, vol. 154, pp. 10–18, 11 2018.

[117] X. Ge and L. Gong, "MiR-590-3p suppresses hepatocellular carcinoma growth by targeting TEAD1," *Tumor Biology*, vol. 39, p. 101042831769594, 3 2017.

[118] Y. Liu, A. Beyer, and R. Aebersold, "On the Dependency of Cellular Protein Levels on mRNA Abundance.," *Cell*, vol. 165, pp. 535–50, 4 2016.

[119] C. Vogel and E. M. Marcotte, "Insights into the regulation of protein abundance from proteomic and transcriptomic analyses," *Nature Reviews Genetics*, vol. 13, pp. 227–232, 4 2012.

[120] J. J. Li, P. J. Bickel, and M. D. Biggin, "System wide analyses have underestimated protein abundances and the importance of transcription in mammals," *PeerJ*, vol. 2, p. e270, 2 2014.

[121] M. Jovanovic, M. S. Rooney, P. Mertins, D. Przybylski, N. Chevrier, R. Satija, E. H. Rodriguez, A. P. Fields, S. Schwartz, R. Raychowdhury, M. R. Mumbach, T. Eisenhaure, M. Rabani, D. Gennert, D. Lu, T. Delorey, J. S. Weissman, S. A. Carr, N. Hacohen, and A. Regev, "Dynamic profiling of the protein life cycle in response to pathogens," *Science*, vol. 347, pp. 1259038–1259038, 3 2015.

[122] K. I. Pishas, C. D. Drenberg, C. Taslim, E. R. Theisen, K. M. Johnson, R. S. Saund, I. L. Pop, B. D. Crompton, E. R. Lawlor, F. Tirode, J. Mora, O. Delattre, M. C. Beckerle, D. F. Callen, S. Sharma, and S. L. Lessnick, "Therapeutic targeting of KDM1A/LSD1 in ewing sarcoma with SP-2509 engages the endoplasmic reticulum stress response," *Molecular Cancer Therapeutics*, vol. 17, pp. 1902–1916, 9 2018.

[123] I. Zwiener, B. Frisch, and H. Binder, "Transforming RNA-Seq Data to Improve the Performance of Prognostic Gene Signatures," *PLoS ONE*, vol. 9, p. e85150, 1 2014.

[124] M. I. Love, W. Huber, and S. Anders, "Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2," *Genome Biology*, vol. 15, p. 550, 12 2014.

[125] A. Dhawan, D. Nichol, F. Kinose, M. E. Abazeed, A. Marusyk, E. B. Haura, and J. G. Scott, "Collateral sensitivity networks reveal evolutionary instability and novel treatment strategies in ALK mutated non-small cell lung cancer," *Scientific Reports*, vol. 7, p. 1232, 12 2017.

[126] A. J. Enright, B. John, U. Gaul, T. Tuschl, C. Sander, and D. S. Marks, "MicroRNA targets in Drosophila.," *Genome Biology*, vol. 5, no. 1, p. R1, 2003.

[127] S. Lall, D. Grün, A. Krek, K. Chen, Y.-L. Wang, C. N. Dewey, P. Sood, T. Colombo, N. Bray, P. MacMenamin, H.-L. Kao, K. C. Gunsalus, L. Pachter, F. Piano, and N. Rajewsky, "A Genome-Wide Map of Conserved MicroRNA Targets in C. elegans," *Current Biology*, vol. 16, pp. 460–471, 3 2006.

[128] R. C. Friedman, K. K.-H. Farh, C. B. Burge, and D. P. Bartel, "Most mammalian mRNAs are conserved targets of microRNAs," *Genome Research*, vol. 19, pp. 92–105, 10 2008.

[129] E. Eisenberg and E. Y. Levanon, "Human housekeeping genes, revisited," oct 2013.

[130] M. L. Shahreza, N. Ghadiri, S. R. Mousavi, J. Varshosaz, and J. R. Green, "A review of network-based approaches to drug repositioning," *Briefings in Bioinformatics*, vol. 19, no. 5, pp. 878–892, 2017.

[131] D. Chasman, A. Fotuhi Siahpirani, and S. Roy, "Network-based approaches for analysis of complex biological systems," *Current Opinion in Biotechnology*, vol. 39, pp. 157–166, June 2016.

[132] U. Chitra, T. Y. Park, and B. J. Raphael, "NetMix2: Unifying Network Propagation and Altered Subnetworks," in *Research in Computational Molecular Biology*, pp. 193–208, Springer, Cham, 2022.

[133] J. G. Scott, A. Berglund, M. J. Schell, I. Mihaylov, W. J. Fulp, B. Yue, E. Welsh, J. J. Caudell, K. Ahmed, T. S. Strom, E. Mellon, P. Venkat, P. Johnstone, J. Foekens, J. Lee, E. Moros, W. S. Dalton, S. A. Eschrich, H. McLeod, L. B. Harrison, and J. F. Torres-Roca, "A genome-based model for adjusting radiotherapy dose (GARD): a retrospective, cohort-based study," *The Lancet Oncology*, vol. 18, pp. 202–211, 2 2017.

[134] M. Bianchini, M. Gori, and F. Scarselli, "Inside PageRank," *ACM Transactions on Internet Technology*, vol. 5, pp. 92–128, Feb. 2005.

[135] C. Navarro, V. Martínez, A. Blanco, and C. Cano, "ProphTools: general prioritization tools for heterogeneous biological networks," *GigaScience*, vol. 6, pp. 1–8, Dec. 2017.

[136] S. A. f. i. a. details, "igraph: Network Analysis and Visualization," Sept. 2022.

[137] L. Gatto and A. Christoforou, "Using R and Bioconductor for proteomics data analysis," *Biochimica et Biophysica Acta (BBA) - Proteins and Proteomics*, vol. 1844, pp. 42–51, Jan. 2014. arXiv:1305.6559 [q-bio].

[138] H. Fang and J. Gough, "The 'dnet' approach promotes emerging research on cancer patient survival," *Genome Medicine*, vol. 6, p. 64, Aug. 2014.

[139] G. Valentini, "RANKS: Ranking of Nodes with Kernelized Score Functions," Sept. 2022.

[140] S. Yakushov, M. Menyailo, E. Denisov, I. Karlina, V. Zainullina, K. Kirgizov, O. Romantsova, P. Timashev, and I. Ulasov, "Identification of Factors Driving Doxorubicin-Resistant Ewing Tumor Cells to Survival," *Cancers*, vol. 14, p. 5498, Nov. 2022.

[141] I. M. Bennani-Baiti, I. Machado, A. Llombart-Bosch, and H. Kovar, "Lysine-specific demethylase 1 (LSD1/KDM1A/AOF2/BHC110) is expressed and is an epigenetic drug target in chondrosarcoma, Ewing's sarcoma, osteosarcoma, and rhabdomyosarcoma.," *Human pathology*, vol. 43, pp. 1300–7, 8 2012.

[142] D. Vella, S. Marini, F. Vitali, D. Di Silvestre, G. Mauri, and R. Bellazzi, "MTGO: PPI Network Analysis Via Topological and Functional Module Identification," *Scientific Reports*, vol. 8, no. 1, p. 5499, 2018.

[143] D. Warde-Farley, S. L. Donaldson, O. Comes, K. Zuberi, R. Badrawi, P. Chao, M. Franz, C. Grouios, F. Kazi, C. T. Lopes, A. Maitland, S. Mostafavi, J. Montojo, Q. Shao, G. Wright, G. D. Bader, and Q. Morris, "The GeneMANIA prediction server: biological network integration for gene prioritization and predicting gene function," *Nucleic Acids Research*, vol. 38, no. suppl_2, pp. w214–w220, 2010.

[144] A. Kamburov, U. Stelzl, H. Lehrach, and R. Herwig, “The ConsensusPathDB interaction database: 2013 update,” *Nucleic Acids Research*, vol. 41, no. D1, pp. d793–d800, 2012.

[145] A. Vinayagam, T. E. Gibson, H.-J. Lee, B. Yilmazel, C. Roesel, Y. Hu, Y. Kwon, A. Sharma, Y.-Y. Liu, N. Perrimon, and A.-L. Barabási, “Controllability analysis of the directed human protein interaction network identifies disease genes and drug targets,” *Proceedings of the National Academy of Sciences*, vol. 113, pp. 4976–4981, May 2016. Publisher: Proceedings of the National Academy of Sciences.

[146] C. J. Murray, K. S. Ikuta, F. Sharara, L. Swetschinski, G. R. Aguilar, A. Gray, C. Han, C. Bisignano, P. Rao, E. Wool, S. C. Johnson, A. J. Browne, M. G. Chipeta, F. Fell, S. Hackett, G. Haines-Woodhouse, B. H. K. Hamadani, E. A. P. Kumaran, B. McManigal, R. Agarwal, S. Akech, S. Albertson, J. Amuasi, J. Andrews, A. Aravkin, E. Ashley, F. Bailey, S. Baker, B. Basnyat, A. Bekker, R. Bender, A. Bethou, J. Bielicki, S. Boonkasidecha, J. Bukosia, C. Carvalheiro, C. Castañeda-Orjuela, V. Chansamouth, S. Chaurasia, S. Chiurchiù, F. Chowdhury, A. J. Cook, B. Cooper, T. R. Cressey, E. Criollo-Mora, M. Cunningham, S. Darboe, N. P. J. Day, M. D. Luca, K. Dokova, A. Dramowski, S. J. Dunachie, T. Eckmanns, D. Eibach, A. Emami, N. Feasey, N. Fisher-Pearson, K. Forrest, D. Garrett, P. Gastmeier, A. Z. Giref, R. C. Greer, V. Gupta, S. Haller, A. Haselbeck, S. I. Hay, M. Holm, S. Hopkins, K. C. Iregbu, J. Jacobs, D. Jarovsky, F. Javanmardi, M. Khorana, N. Kissoon, E. Kobeissi, T. Kostyanev, F. Krapp, R. Krumkamp, A. Kumar, H. H. Kyu, C. Lim, D. Limmathurotsakul, M. J. Loftus, M. Lunn, J. Ma, N. Mturi, T. Munera-Huertas, P. Musicha, M. M. Mussi-Pinhata, T. Nakamura, R. Nanavati, S. Nangia, P. Newton, C. Ngoun, A. Novotney, D. Nwakanma, C. W. Obiero, A. Olivas-Martinez, P. Olliaro, E. Ooko, E. Ortiz-Brizuela, A. Y. Peleg, C. Perrone, N. Plakkal, A. Ponce-de Leon, M. Raad, T. Ramdin, A. Riddell, T. Roberts, J. V. Robotham, A. Roca, K. E. Rudd, N. Russell, J. Schnall, J. A. G. Scott, M. Shivamallappa, J. Sifuentes-Osornio, N. Steenkeste, A. J. Stewardson, T. Stoeva, N. Tasak, A. Thaiprakong, G. Thwaites, C. Turner, P. Turner, H. R. v. Doorn, S. Velaphi, A. Vongpradith, H. Vu, T. Walsh, S. Waner, T. Wangrangsimakul, T. Wozniak, P. Zheng, B. Sartorius, A. D. Lopez, A. Stergachis, C. Moore, C. Dolecek, and M. Naghavi, “Global burden of bacterial antimicrobial resistance in 2019: a systematic analysis,” *The Lancet*, vol. 399, pp. 629–655, Feb. 2022. Publisher: Elsevier.

[147] Centers for Disease Control and Prevention (U.S.), “Antibiotic resistance threats in the United States, 2019,” tech. rep., Centers for Disease Control and Prevention (U.S.), Nov. 2019.

[148] B. Plackett, “Why big pharma has abandoned antibiotics,” *Nature*, vol. 586, Oct. 2020.

[149] S. C. Stearns, “Evolutionary medicine: Its scope, interest and potential,” *Proceedings of the Royal Society B: Biological Sciences*, vol. 279, pp. 4305–4321, Nov. 2012.

[150] D. Z. Grunspan, R. M. Nesse, M. E. Barnes, and S. E. Brownell, “Core principles of evolutionary medicine,” *Evolution, Medicine, and Public Health*, vol. 2018, pp. 13–23, Jan. 2018.

[151] G. H. Perry, “Evolutionary medicine,” *eLife*, vol. 10, p. e69398, July 2021.

[152] D. I. Andersson, N. Q. Balaban, F. Baquero, P. Courvalin, P. Glaser, U. Gophna, R. Kishony, S. Molin, and T. Tønjum, “Antibiotic resistance: turning evolutionary principles into clinical reality,” *FEMS Microbiology Reviews*, vol. 44, no. 2, pp. 171–188, 2020. Publisher: Oxford University Press.

[153] S. Manrubia, J. A. Cuesta, J. Aguirre, S. E. Ahnert, L. Altenberg, A. V. Cano, P. Catalán, R. Diaz-Uriarte, S. F. Elena, J. A. García-Martín, P. Hogeweg, B. S. Khatri, J. Krug, A. A. Louis, N. S. Martin, J. L. Payne, M. J. Tarnowski, and M. Weiß, “From genotypes to organisms: State-of-the-art and perspectives of a cornerstone in evolutionary dynamics,” *Physics of Life Reviews*, vol. 38, pp. 55–106, Sept. 2021.

[154] M. Stracy, O. Snitser, I. Yelin, Y. Amer, M. Parizade, R. Katz, G. Rimler, T. Wolf, E. Herzel, G. Koren, J. Kuint, B. Foxman, G. Chodick, V. Shalev, and R. Kishony, “Minimizing treatment-induced emergence of antibiotic resistance in bacterial infections,” *Science*, vol. 375, pp. 889–894, Feb. 2022. Publisher: American Association for the Advancement of Science.

[155] M. Baym, L. K. Stone, and R. Kishony, “Multidrug evolutionary strategies to reverse antibiotic resistance,” *Science*, vol. 351, no. 6268, p. aad3292, 2016. Publisher: American Association for the Advancement of Science.

[156] N. Yoon, R. Vander Velde, A. Marusyk, and J. G. Scott, “Optimal Therapy Scheduling Based on a Pair of Collaterally Sensitive Drugs,” *Bulletin of Mathematical Biology*, vol. 80, pp. 1776–1809, July 2018.

[157] J. Maltas and K. B. Wood, “Pervasive and diverse collateral sensitivity profiles inform optimal strategies to limit antibiotic resistance,” *PLoS biology*, vol. 17, p. e3000515, Oct. 2019.

[158] J. Maltas and K. B. Wood, “Dynamic collateral sensitivity profiles highlight challenges and opportunities for optimizing antibiotic sequences,” *bioRxiv*, 2021.

[159] M. Gluzman, J. G. Scott, and A. Vladimirsky, "Optimizing adaptive cancer therapy: dynamic programming and evolutionary game theory," *Proceedings of the Royal Society B: Biological Sciences*, vol. 287, p. 20192454, Apr. 2020. Publisher: Royal Society.

[160] D. Nichol, P. Jeavons, A. G. Fletcher, R. A. Bonomo, P. K. Maini, J. L. Paul, R. A. Gatenby, A. R. A. Anderson, and J. G. Scott, "Steering Evolution with Sequential Therapy to Prevent the Emergence of Bacterial Antibiotic Resistance," *PLOS Computational Biology*, vol. 11, p. e1004493, Sept. 2015. Publisher: Public Library of Science.

[161] N. Yoon, N. Krishnan, and J. Scott, "Theoretical modeling of collaterally sensitive drug cycles: shaping heterogeneity to allow adaptive therapy," *Journal of Mathematical Biology*, vol. 83, p. 47, Oct. 2021.

[162] S. Iram, E. Dolson, J. Chiel, J. Pelesko, N. Krishnan, O. Güngör, B. Kuznets-Speck, S. Deffner, E. Ilker, J. G. Scott, and M. Hinczewski, "Controlling the speed and trajectory of evolution with counterdiabatic driving," *Nature Physics*, vol. 17, pp. 135–142, Jan. 2021. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 1 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Biophysics;Statistical physics;Theoretical physics Subject_term_id: biophysics;statistical-physics;theoretical-physics.

[163] J. Maltas, K. R. Singleton, K. C. Wood, and K. B. Wood, "Drug dependence in cancer is exploitable by optimally constructed treatment holidays," *bioRxiv*, 2022.

[164] S. Chakrabarti and F. Michor, "Pharmacokinetics and drug interactions determine optimum combination strategies in computational models of cancer evolution," *Cancer Research*, vol. 77, p. 3908–3921, Jul 2017.

[165] P. K. Newton and Y. Ma, "Nonlinear adaptive control of competitive release and chemotherapeutic resistance," *Physical Review. E*, vol. 99, p. 022404, Feb. 2019.

[166] S. Kim, T. D. Lieberman, and R. Kishony, "Alternating antibiotic treatments constrain evolutionary paths to multidrug resistance," *Proceedings of the National Academy of Sciences*, vol. 111, pp. 14494–14499, Oct. 2014. Publisher: Proceedings of the National Academy of Sciences.

[167] J. Zhang, J. J. Cunningham, J. S. Brown, and R. A. Gatenby, "Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer," *Nature Communications*, vol. 8, Nov. 2017.

[168] J. J. Cunningham, J. S. Brown, R. A. Gatenby, and K. Staňková, "Optimal control to develop therapeutic strategies for metastatic castrate resistant prostate cancer," *Journal of Theoretical Biology*, vol. 459, pp. 67–78, Dec. 2018.

[169] D. M. Weinreich, R. A. Watson, and L. Chao, “Perspective: Sign Epistasis and Genetic Costraint on Evolutionary Trajectories,” *Evolution*, vol. 59, no. 6, pp. 1165–1174, 2005. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.0014-3820.2005.tb01768.x.

[170] P. M. Mira, K. Crona, D. Greene, J. C. Meza, B. Sturmfels, and M. Barlow, “Rational Design of Antibiotic Treatment Plans: A Treatment Strategy for Managing Evolution and Reversing Resistance,” *PLOS ONE*, vol. 10, p. e0122283, May 2015. Publisher: Public Library of Science.

[171] J. Maltas, D. M. McNally, and K. B. Wood, “Evolution in alternating environments with tunable inter-landscape correlations,” *Evolution; international journal of organic evolution*, vol. 75, pp. 10–24, Jan. 2021.

[172] J. A. G. M. de Visser and J. Krug, “Empirical fitness landscapes and the predictability of evolution,” *Nature Reviews Genetics*, vol. 15, pp. 480–490, July 2014. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 7 Primary_atype: Reviews Publisher: Nature Publishing Group Subject_term: Epistasis;Evolutionary genetics;Experimental evolution Subject_term_id: epistasis;evolutionary-genetics;experimental-evolution.

[173] D. M. Weinreich, N. F. Delaney, M. A. Depristo, and D. L. Hartl, “Darwinian evolution can follow only very few mutational paths to fitter proteins,” *Science (New York, N.Y.)*, vol. 312, pp. 111–114, Apr. 2006.

[174] C. B. Ogbunugafor, C. S. Wylie, I. Diakite, D. M. Weinreich, and D. L. Hartl, “Adaptive Landscape by Environment Interactions Dictate Evolutionary Dynamics in Models of Drug Resistance,” *PLoS Comput. Biol.*, vol. 12, no. 1, pp. 1–20, 2016.

[175] E. Toprak, A. Veres, J.-B. Michel, R. Chait, D. L. Hartl, and R. Kishony, “Evolutionary paths to antibiotic resistance under dynamically sustained drug selection,” *Nature genetics*, vol. 44, no. 1, pp. 101–105, 2012. Publisher: Nature Publishing Group.

[176] S. F. Greenbury, A. A. Louis, and S. E. Ahnert, “The structure of genotype-phenotype maps makes fitness landscapes navigable,” *Nature Ecology & Evolution*, pp. 1–11, Sept. 2022. Publisher: Nature Publishing Group.

[177] M. Baym, T. D. Lieberman, E. D. Kelsic, R. Chait, R. Gross, I. Yelin, and R. Kishony, “Spatiotemporal microbial evolution on antibiotic landscapes,” *Science (New York, N.Y.)*, vol. 353, pp. 1147–1151, Sept. 2016.

[178] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[179] L. Du Plessis, G. E. Leventhal, and S. Bonhoeffer, “How good are statistical models at approximating complex fitness landscapes?,” *Molecular biology and evolution*, vol. 33, no. 9, pp. 2454–2468, 2016. Publisher: Oxford University Press.

[180] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, “Mastering the game of Go with deep neural networks and tree search,” *Nature*, vol. 529, Jan. 2016.

[181] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, “Human-level control through deep reinforcement learning,” *Nature*, vol. 518, pp. 529–533, Feb. 2015. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 7540 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Computer science Subject_term_id: computer-science.

[182] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, J. Oh, D. Horgan, M. Kroiss, I. Danihelka, A. Huang, L. Sifre, T. Cai, J. P. Agapiou, M. Jaderberg, A. S. Vezhnevets, R. Leblond, T. Pohlen, V. Dalibard, D. Budden, Y. Sulsky, J. Molloy, T. L. Paine, C. Gulcehre, Z. Wang, T. Pfaff, Y. Wu, R. Ring, D. Yogatama, D. Wünsch, K. McKinney, O. Smith, T. Schaul, T. Lillicrap, K. Kavukcuoglu, D. Hassabis, C. Apps, and D. Silver, “Grandmaster level in StarCraft II using multi-agent reinforcement learning,” *Nature*, vol. 575, pp. 350–354, Nov. 2019. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 7782 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Computer science;Statistics Subject_term_id: computer-science;statistics.

[183] B. L. Moore, L. D. Pyeatt, Vivekan, Kulkarni, P. Panousis, K. Padrez, and A. G. Doufas, “Reinforcement Learning for Closed-Loop Propofol Anesthesia: A Study in Human Volunteers,” *Journal of Machine Learning Research*, vol. 15, no. 21, pp. 655–696, 2014.

[184] B. K. Petersen, J. Yang, W. S. Grathwohl, C. Cockrell, C. Santiago, G. An, and D. M. Faissol, “Deep Reinforcement Learning and Simulation as a Path Toward Precision Medicine,” *Journal of Computational Biology: A Journal of Computational Molecular Cell Biology*, vol. 26, pp. 597–604, June 2019.

[185] R. Padmanabhan, N. Meskin, and W. M. Haddad, “Reinforcement learning-based control of drug dosing for cancer chemotherapy treatment,” *Mathematical Biosciences*, vol. 293, pp. 11–20, Nov. 2017.

[186] I. Ahn and J. Park, “Drug scheduling of cancer chemotherapy based on natural actor-critic approach,” *Bio Systems*, vol. 106, pp. 121–129, Nov. 2011.

[187] D. Engelhardt, “Dynamic Control of Stochastic Evolution: A Deep Reinforcement Learning Approach to Adaptively Targeting Emergent Drug Resistance,” *Journal of Machine Learning Research*, vol. 21, no. 203, pp. 1–30, 2020.

[188] R. B. Martin, “Optimal control drug scheduling of cancer chemotherapy,” *Automatica*, vol. 28, pp. 1113–1123, Nov. 1992.

[189] L. Kallenberg, “Lecture Notes Markov Decision Problems - version 2020,” Feb. 2020.

[190] K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, “Deep reinforcement learning: A brief survey,” *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 26–38, 2017.

[191] E. S. King, J. Pelesko, J. Maltas, S. J. Owen, E. Dolson, and J. G. Scott, “Fitness seascapes facilitate the prediction of therapy resistance under time-varying selection,” June 2022. Pages: 2022.06.10.495696 Section: New Results.

[192] C. K. Rode, L. J. Melkerson-Watson, A. T. Johnson, and C. A. Bloch, “Type-Specific Contributions to Chromosome Size Differences in Escherichia coli,” *Infection and Immunity*, vol. 67, pp. 230–236, Jan. 1999.

[193] J. Iwasawa, T. Maeda, A. Shibai, H. Kotani, M. Kawada, and C. Furusawa, “Analysis of the evolution of resistance to multiple antibiotics enables prediction of the Escherichia coli phenotype-based fitness landscape,” *PLOS Biology*, vol. 20, p. e3001920, Dec. 2022. Publisher: Public Library of Science.

[194] T. J. Raphel, D. T. Weaver, L. L. Berland, B. R. Herts, A. J. Megibow, A. B. Knudsen, and P. V. Pandharipande, “Imaging Follow-up of Low-Risk Incidental Pancreas and Kidney Findings: Effects of Patient Age and Comorbidity on Projected Life Expectancy,” *Radiology*, vol. 287, pp. 504–514, May 2018. Publisher: Radiological Society of North America.

[195] D. T. Weaver, T. J. Raphel, A. Melamed, J. A. Rauh-Hain, J. O. Schorge, A. B. Knudsen, and P. V. Pandharipande, “Modeling treatment outcomes for patients with advanced ovarian cancer: Projected benefits of a test to optimize treatment selection,” *Gynecologic Oncology*, vol. 149, pp. 256–262, May 2018.

[196] D. T. Weaver, A. P. Lietz, S. F. Mercaldo, M. L. B. Peters, C. Hur, C. Y. Kong, B. M. Wolpin, A. J. Megibow, L. L. Berland, and A. B. Knudsen, “Testing for Verification Bias in Reported Malignancy Risks for Side-Branch Intraductal

Papillary Mucinous Neoplasms: A Simulation Modeling Approach," *AJR. American journal of roentgenology*, vol. 212, no. 3, p. 596, 2019. Publisher: NIH Public Access.

[197] Q. Chen, M. R. Larochelle, D. T. Weaver, A. P. Lietz, P. P. Mueller, S. Mercaldo, S. E. Wakeman, K. A. Freedberg, T. J. Raphel, A. B. Knudsen, P. V. Pandharipande, and J. Chhatwal, "Prevention of Prescription Opioid Misuse and Projected Overdose Deaths in the United States," *JAMA Network Open*, vol. 2, p. e187621, 2 2019.

[198] U. S. D. O. H. A. H. S. S. A. A. M. H. S. A. C. F. B. H. S. A. Quality, "National Survey on Drug Use and Health, 2014: Version 1," 2016. Version Number: v1 Type: dataset.

[199] S. Wright, "Evolution in Mendelian Populations," *Genetics*, vol. 16, pp. 97–159, Mar. 1931.

[200] R. A. Fisher, "XXI.—On the Dominance Ratio," *Proceedings of the Royal Society of Edinburgh*, vol. 42, pp. 321–341, 1923. Publisher: Royal Society of Edinburgh Scotland Foundation.