Inverted topologies in sequential fitness landscapes enable evolutionary control

Peng Chen[1,2+], Nikhil Krishnan[3,+], Anna Stacy[1,2], Davis T. Weaver[1,4], Rowan Barker-Clarke[1], Michael Hinczewski[2], Jeff Maltas[1,4,*], and Jacob G. Scott[1,2,4,*]

[1]Translational Hematology Oncology Research, Lerner Research Institute, Cleveland Clinic, US
[2]Department of Physics, Case Western Reserve University, US
[3]Cavendish Laboratory, Department of Physics, University of Cambridge, UK
[4]School of Medicine, Case Western Reserve University, US
[*]scottj10@ccf.org, jeff.maltas@gmail.com
[+]These authors contributed equally to this work.

ABSTRACT

Adaptive populations rarely evolve in a static environment. Therefore, understanding and ultimately controlling the evolution of a population requires consideration of fluctuating selective pressures. The fitness landscape metaphor has long been used as a tool for representing the selective pressures a given environment imposes on a population. Much work has already been done to understand the dynamics of evolution on a single fitness landscape. More recently, evolution on fluctuating or sequentially applied landscapes has come to the fore of evolutionary biology. As more empirical landscapes are described, metrics for describing salient features of paired landscapes will have uses for understanding likely evolutionary dynamics. Currently, Pearson correlation coefficient and collateral sensitivity likelihoods are used to quantify topological relatedness or dissimilarity of a pair of landscapes. Here, we introduce the edge flip fraction, a new metric for comparing landscapes, which quantifies changes in the directionality of evolution between pairs of fitness landscapes. We demonstrate that the edge flip fraction captures topological differences that traditional metrics may overlook. By applying this metric to both empirical and synthetic fitness landscapes, we show that it partially predicts the collateral sensitivity likelihoods and can inform the optimality of drug sequences. We show that optimal drug sequences that keep populations within lower fitness regions require shifts in evolutionary directions, which are quantified by the edge flip fraction. Edge flip fraction complements existing measures and may help researchers understand how populations evolve under changing environmental conditions.

Introduction

The ability to engineer and control evolution is a major outstanding question in biology with wide ranging applications for human health[1], agriculture[2, 3], industry[4], and ecosystems[5, 6]. However, successful prediction and control remains difficult[7, 8]. Fitness (adaptive) landscapes, a mathematical construct that maps an individuals genotype to a phenotypic measure of fitness, have been proven to be a valuable model for diverse evolving systems from bacteria to yeast[9–13]. Over the last few decades, rapid experimental advances have made it increasingly possible to enumerate enormous genotype-to-phenotype maps, thus driving continued interest in theoretical modeling of fitness landscapes[14–20]. These experimental landscapes have revealed that mutations often have a complex relationship with the genetic background in which they arise (G×G interactions), their evolutionary environment (G×E interactions), or both (G×G×E interactions)[21, 22].

In particular, the topography of these fitness landscapes has sparked recent interest due to its role in defining evolutionary trajectories[23, 24]. Critically, the topography defines how beneficial or deleterious emerging mutants are and therefore how likely they are to fix in the population. However, the topography of a fitness landscape is not static. Instead, the topography can fluctuate due to nonlinear interactions between genes (known as epistasis)[25–28], or environmental variation such as a varying antibiotic concentration[29], sequential drug application[30], or simply variable weather[31] and climate[6]. For example, collateral sensitivity — a phenomenon whereby a mutation that confers resistance to one environment also confers increased sensitivity to another unseen environment — would lead to fitness landscapes with negatively correlated topography [32–37]. In this context, a mutation that was once beneficial and fixed in one environment, is now deleterious and subsequently shed (or compensated for) in a following environment[38–40]. As a result, evolving populations are constantly challenged by fluctuating fitness landscape topographies that lead to dynamic fixation/extinction rates[41–49].

In this work we seek to quantify the importance of this phenomenon — fitness landscape edge flips – where environmental variation leads to a dynamic fitness landscape topography that alters which mutations are beneficial or deleterious (and as

a result “flips" the preferred direction of evolution on the fitness landscape). We hypothesized that edge flips are critical to evolutionary control as they provide a genotype from which an experimentalist or physician by exert a particular selection pressure, favoring the selection of a mutant of their interest. We define a new metric for paired fitness landscapes the “edge flip fraction" which quantifies the extent to which two fitness landscapes preferred evolutionary trajectories differ. We show that edge flip fraction is highly correlated with past metrics for paired fitness landscapes such as their Pearson correlation coefficient and collateral sensitivity likelihood. However, because edge flip fraction is sensitive to evolutionary transition rates themselves, rather than global structure like correlation coefficients, edge flip fraction proves to be a more valuable metric for identifying potential evolutionary control. To demonstrate this we leverage a previously published set of fitness landscapes to numerous β-lactam antibiotics [50] and reveal that the optimal sequence of antibiotics to minimize resistance enriches specifically for sequences of fitness landscapes with numerous edge flips. Finally, we show that because evolutionary control is leveraged at the genotypic level, it is the genotypic edge flip fraction (local edge flip fraction) that determines controllability.

Methods

A Markov Model of Evolution on Fitness Landscapes

We begin by considering Wright’s Fitness Landscape, or a mapping of genotype and fitness. As established in previous studies [18, 49], we imagine that each of N loci can be either wild-type or contain a point mutation. Any given genotype is then represented by a bit-string of length N where bit 0 represents the absence of a point mutation, and bit 1, presence. This gives us 2^N possible genotypes which can be represented as an N dimensional hypercube with vertices in $\{0,1\}^N$ and edges connecting mutational neighbors. The fitness of a genotype, driven by selection from the environment, can be presented by the following fitness function:

$$f:\ \{0,1\}^N \rightarrow \mathbb{R}^{\geq 0} \tag{1}$$

It is assumed that the entire population occupies one genotype at a time, or in other words, the population is isogenic. This assumption is valid when mutation rate is sufficiently low and selective advantage is sufficiently high relative to population size such that a mutation fixes on much shorter times scales than the waiting time for the next mutation. This regime is also known as the strong selection weak mutation (SSWM) regime, as dubbed by Gillespie [10]. The details of how the population moves from genotype to genotype depending on their fitness varies with the model of evolution employed. In this paper the population’s evolution will be described as a Markov chain governed by the random move model [49]. At each evolutionary time step, the population transitions to a fitter neighboring genotype (at Hamming distance of 1) with uniform probability, and the probability of these transitions is encoded in the transition matrix $P = \mathbb{P}[(i \rightarrow j)$. By leveraging the properties of Markov chains, it can be shown that the population reaches a local fitness peak in finite evolutionary steps [49].

Calculating the Edge Flip Fraction between Pairs of Fitness Landscapes

The resultant transition matrices described in the previous section define directed graphs with self-loops at the local fitness peaks. In this work we are interested in understanding changes in the direction of evolution between pairs of fitness landscapes. For the transition matrix of two landscapes **A** and **B**, the edge flip matrix **E** is given by:

$$\mathbf{E} = [(U^{\lceil \mathbf{A} \rceil} - L^{\lceil \mathbf{A} \rceil T}) - (U^{\lceil \mathbf{B} \rceil} - L^{\lceil \mathbf{B} \rceil T})]/2 \tag{2}$$

Here, $U^{\lceil \mathbf{A} \rceil}$ and $U^{\lceil \mathbf{B} \rceil}$ denote the upper triangular portions of the matrices **A** and **B**, respectively, after applying the ceiling function to each element. The ceiling function rounds each element up to the smallest integer greater than or equal to the given number, resulting in matrices consisting only of 0s and 1s. Similarly, $L^{\lceil \mathbf{A} \rceil T}$ and $L^{\lceil \mathbf{B} \rceil T}$ refer to the transposed lower triangular portions of the matrices **A** and **B**, respectively, with the ceiling function also applied to each element.

Then, the edge flip fraction, denoted by $\mathscr{E}$, quantifies the extent of directional changes in evolution between the two landscapes. It is calculated by summing over all elements $|\mathbf{e}_{ij}|$ within the edge flip matrix **E** and normalizing by the total number of edges in the N dimensioinal hypercube:

$$\mathscr{E} = \frac{\sum_{ij} |\mathbf{e}_{ij}|}{N \cdot 2^N/2} \tag{3}$$

If the edge flip fraction, $\mathscr{E} = 1$, it indicates that all the directions of evolution differ between the two landscapes, whereas $\mathscr{E} = 0$ signifies no changes in evolutionary direction. This formulation provides an analytical framework for quantifying the evolutionary direction changes between pairs of landscapes.

Box 1: Edge Flip Fraction

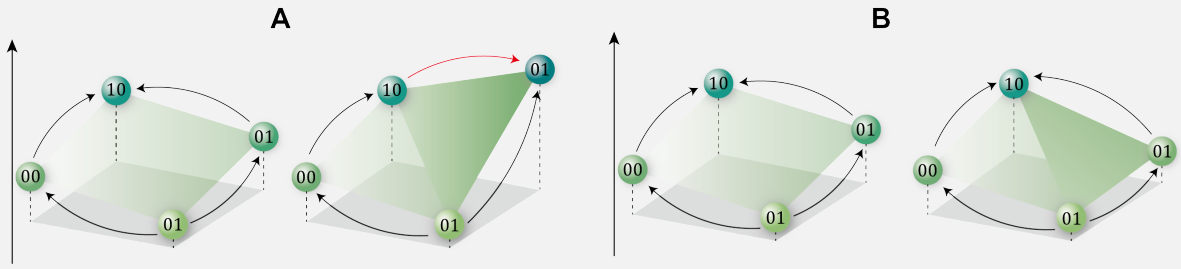

Fitness $f_A = [0.5, 0.4, 0.7, 0.6]$ $f_B = [0.5, 0.4, 0.7, 0.8]$ $f_A = [0.5, 0.4, 0.7, 0.6]$ $f_B = [0.5, 0.4, 0.7, 0.45]$

Transition Matrix

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ .5 & 0 & 0 & .5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad \mathbf{B} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ .5 & 0 & 0 & .5 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ .5 & 0 & 0 & .5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad \mathbf{B} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ .5 & 0 & 0 & .5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

Edge Flip Fraction

$$\mathbf{E} = \left[U^{\lceil \mathbf{A} \rceil} - L^{\lceil \mathbf{A} \rceil T} - \left(U^{\lceil \mathbf{B} \rceil} - L^{\lceil \mathbf{B} \rceil T}\right)\right]/2 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\mathscr{E} = \frac{\sum_{ij} |\mathbf{e}_{ij}|}{N \cdot 2^N/2} = \frac{1}{4}$$

$$\mathbf{E} = \left[U^{\lceil \mathbf{A} \rceil} - L^{\lceil \mathbf{A} \rceil T} - \left(U^{\lceil \mathbf{B} \rceil} - L^{\lceil \mathbf{B} \rceil T}\right)\right]/2 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\mathscr{E} = \frac{\sum_{ij} |\mathbf{e}_{ij}|}{N \cdot 2^N/2} = 0$$

Correlation Coefficient

$r = 0.849$ $r = 0.834$

Figure 1. **A.** A pair of $N = 2$ landscapes in which increasing the fitness of the double mutant '11' to 0.8, results in a Pearson correlation coefficient of $r = 0.849$ and leads to an edge flip, with an edge flip fraction of 0.25. As fitness of '11' exceeds its Hamming distance-1 neighbor '10', the direction of the edge between them changes. **B.** A pair of $N = 2$ landscapes with the same fitness values as in panel A, in which decreasing the fitness of the double mutant '11' to 0.45 yields a similar Pearson correlation coefficient of $r = 0.834$, but does not result in an edge flip, with an edge flip fraction of 0. As long as the fitness of '11' doesn't go below that of its Hamming distance-1 neighbor '01', the direction of the edge remains unchanged, so the changes in fitness are only reflected in the Pearson correlation coefficient.

Collateral Sensitivity Likelihoods (CSLs)

Collateral Sensitivity Likelihoods (CSLs), is an established metric for quantifying the probability that a population, after evolving to the highest local fitness peak under one drug, will exhibit increased sensitivity when subsequently exposed to a second drug [37]. Building on the Markov model described above and following established methods [37], we conducted simulations across pairs of fitness landscapes to measure CSLs.

Construction of Paired Synthetic Fitness Landscapes

We generated 50 synthetic fitness landscapes using a many-peaked "rough Mt. Fuji" model of size N=4 (yielding 16 total genotypes) [51, 52]. As in previous work[11], for each of these 50 landscapes (denoted as A), we created 50 paired landscapes (B) with Pearson correlation coefficients ranging from -1 to 1, resulting in 2500 total landscape pairs that span the full range from anticorrelated to correlated. We repeated the same procedure using the Kauffman NK model [53], producing another 50 landscapes and pairing each with 50 correlated variants. Finally, using an empirically grounded Tradeoff-Induced Landscape (TIL) model [16], which provides genotype-specific dose-response curves (fitness seascapes), we extracted a fitness landscape at the concentration corresponding to IC_{90} as landscape A and generated 50 correlated landscapes B with the same correlation methodology.

Markov Decision Process (MDP)

To systematically identify drug policies that minimize or maximize population fitness, we employed a Markov Decision Process (MDP) framework, as previously described [8, 30, 35]. In this framework, each genotype corresponds to a discrete state. At each timestep, the system transitions stochastically to a new state according to predefined probabilities, and one of several possible actions (e.g., drug choices) is selected. The chosen action determines both an immediate reward (e.g., a change in fitness) and the probabilities governing the subsequent state transition. By iterating over all possible states and actions, the MDP computes a policy that optimizes a given objective function. In this work, we applied the MDP to the fitness landscapes

from Mira et al. to calculate a "optimal" policy that minimizes population fitness over time and a "worst" policy that drives the population toward higher fitness.

Results

Antibiotic	Class	Abbreviation
Ampicillin	Penicillin	AMP
Amoxicillin	Penicillin	AM
Cefaclor	Cephalosporin Gen. 2	CEC
Cefotaxime	Cephalosporin Gen. 3	CTX
Ceftizoxime	Cephalosporin Gen. 3	ZOX
Cefuroxime	Cephalosporin Gen. 2	CXM
Ceftriaxone	Cephalosporin Gen. 3	CRO
Amoxicillin + Clav.	Penicillin + Inhib.	AMC
Ceftazidime	Cephalosporin Gen. 3	CAZ
Cefotetan	Cephalosporin Gen. 2	CTT
Ampicillin +Sulbactam	Penicillin + Inhib.	SAM
Cefprozil	Cephalosporin Gen. 2	CPR
Cefpodoxime	Cephalosporin Gen. 3	CPD
Piperacillin + Tazobactam	Penicillin + Inhib.	TZP
Cefepime	Cephalosporin Gen. 4	FEP

Table 1. Name, class, and abbreviation for each antibiotic included in the Mira et al. study [50].

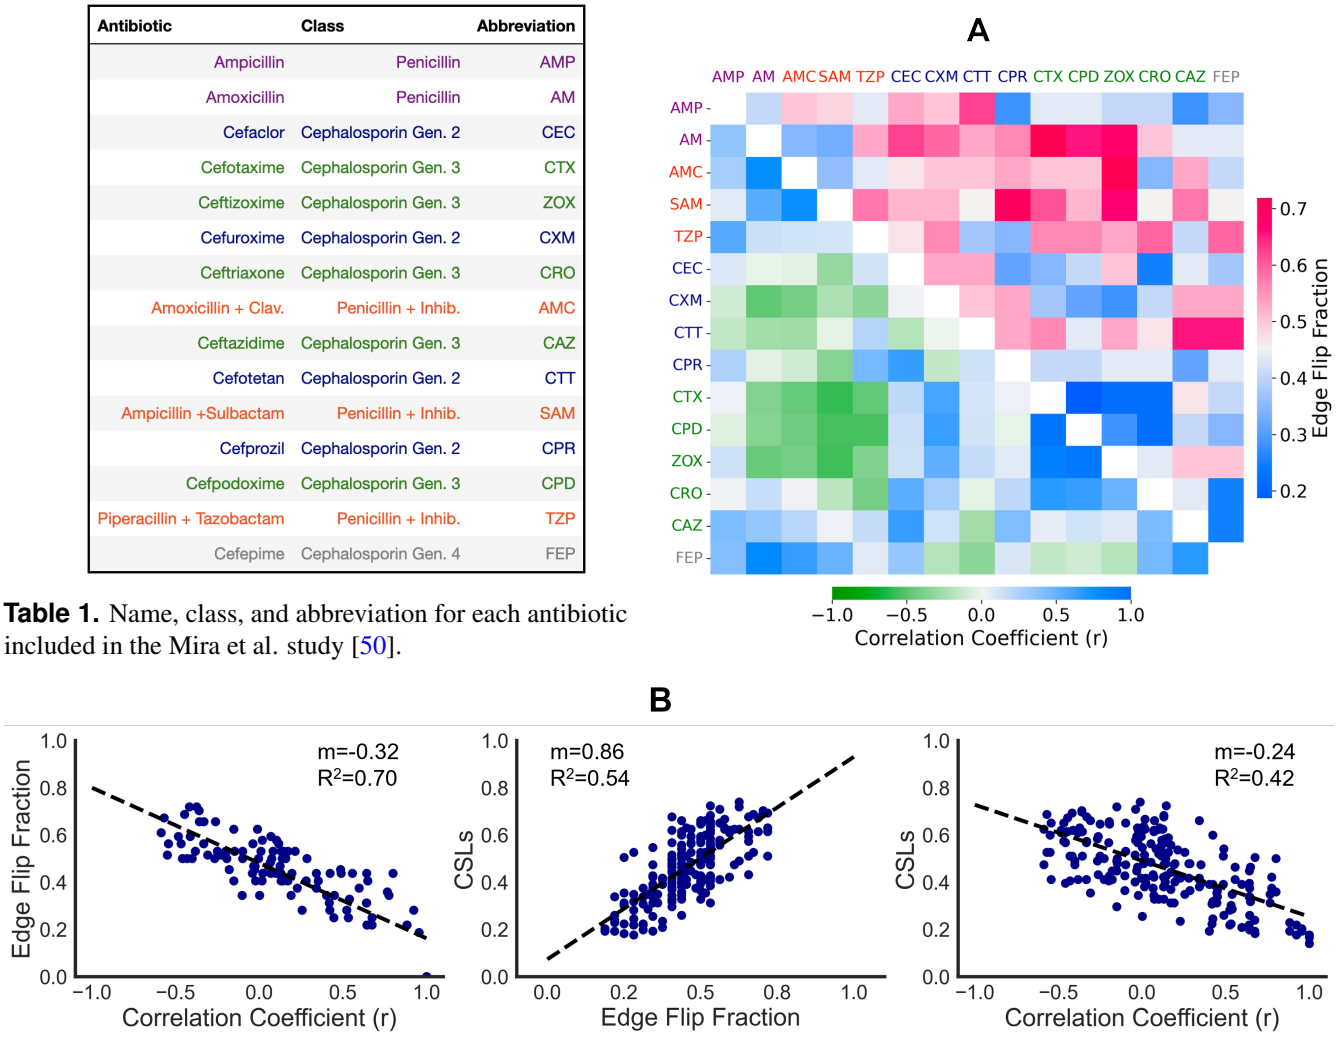

Figure 2. A. Heatmap showing the edge flip fraction (upper triangular) and the Pearson correlation coefficient (lower triangular) between pairs of empirical landscapes measure by Mira et al, with the drugs color-coded by their classes as detailed in **Table 1**. **B.** Scatterplots with linear fits illustrating the relationships between edge flip fraction ($\mathscr{E}$), Pearson correlation coefficient (r), and collateral sensitivity likelihood (CSLs) for the same set of empirical landscapes presented in heatmap A [50].

Explicit Relationship Between Edge Flip Fraction and Pearson Correlation Coefficient in Pairwise Landscapes

Consider two pairs of $N = 2$ fitness landscapes with genotypes '00', '01', '10', and '11' (**Figure 1A, B**). These examples illustrate how the edge flip fraction $\mathscr{E}$ reveals distinct and statistically significant information beyond what is provided by the Pearson correlation coefficient r. For the first pair of landscapes (**Figure 1A**), the Pearson correlation coefficient is $r = 0.849$, and an edge flip is present. Conversely, the second pair (**Figure 1B**) exhibits a similar $r = 0.834$ but no edge flip. This comparison suggests that the edge flip fraction can capture topological features of landscape relationships that may be overlooked by the Pearson correlation coefficient.

Nevertheless, there are cases where the edge flips do not fully account for significant fitness changes. For instance, in **Figure 1B**, if the fitness of genotype '11' decreases but remains higher than that of its Hamming distance-1 neighbor '01',

the edge flip fraction $\mathscr{E}$ remains zero. Meanwhile, the Pearson correlation coefficient can vary substantially within the range $r \in [0.73, 1]$. This underscores the complementary nature of these metrics, with $\mathscr{E}$ providing topological information and r reflecting proportional fitness changes.

To investigate this relationship in empirical fitness landscapes, we calculated the pairwise edge flip fraction and Pearson correlation coefficient across an ensemble of 15 four-locus empirical fitness landscapes reported by Mira et al. [50]. The 15 drugs were categorized and color-coded by their classes, as detailed in **Table 1** . The edge flip fractions observed in the data ranged from 0.19 (CTX-CPD) to 0.72 (CTX-AM) (**Figure 2A**). Drug pairs from the same class, such as CTX and CPD (both classified as Cephalosporin Gen. 2), generally exhibit fewer changes in evolutionary direction, as evidenced by lower edge flip fractions. In contrast, cross-class drug pairs, such as CTX (Cephalosporin Gen. 2) and AM (Penicillin), tended to have higher edge flip fractions, suggesting that the broad biochemical mechanisms defining drug categories substantially influence the topological structure of fitness landscapes. Moreover, a comparison across all drug pairs revealed a moderate anti-correlation between edge flip fractions and Pearson correlation coefficient (**Figure 2B**). Pairs of landscapes with consistent evolutionary direction tended to be more correlated. These results underscore the complementary roles of edge flip fraction and Pearson correlation coefficient in characterizing the relationships between paired fitness landscapes.

Edge Flip Fraction and Its Correlates in Empirical and Synthetic Fitness Landscapes

To go beyond descriptions of landscape topology, we explored how information derived from edge flip fraction could be utilized for rational treatment design. In particular, we examined the suitability of pairs of fitness landscapes for evolutionary control by leveraging collateral sensitivity likelihoods (CSLs) [49]. Within the context of the specific model employed here, collateral sensitivity likelihoods (CSLs) represent the likelihood that a population, having evolved to the highest local fitness peak under one drug, will transition to the lowest fitness peak when subsequently exposed to a second drug. Drug pairing with high CSLs are thus advantageous for therapeutic sequencing as they maintain sensitivity.

As the topological factors that facilitate such a relationship are more complex and model dependent than Pearson correlation coefficient, we expect direction change in evolution to contribute, though not completely, to the CSLs of a given pair. Indeed, a direct comparison of edge flip fraction and CSLs revealed a strong positive correlation in the empirical fitness landscapes of Mira et al. (**Figure 2B**).

To further investigate the edge flip fraction in both empirical and synthetic landscapes, we analyzed pairs of $N = 4$ fitness landscapes constructed using the rough Mount Fuji (RMF) model, Kauffman NK model and trade-off-induced fitness landscapes (TIL) model [11, 52–54]. Despite the differing characteristics of these three types of stimulated landscapes, we consistently observed an anti-correlation between edge flip fraction and Pearson correlation coefficient and a positive correlation between edge flip fraction and CSLs across all models (**Figure 3**). This relationship persisted in higher-dimensional $N = 8$ fitness landscapes, as shown in **Figure S1**.

Our results suggest that beyond just the edge flip fraction, the transition probabilities associated with edges and the magnitude of subsequent changes also contributes to CSLs. A notable challenge in interpreting the relationship between edge flip fraction and CSLs lies in their differing dependencies: CSLs are order-dependent, while edge flip fraction is not. This distinction underscores that the direction of evolution between two genotypes gives a potentially useful yet incomplete picture of the evolutionary dynamics between pairs of landscape. Nevertheless, edge flip fraction remains a computationally efficient metric for predicting CSLs and guiding evolutionary control strategies.

Linking Edge Flip Fraction to Drug Sequence Optimality

Understanding how drug sequences can be leveraged to optimize treatment outcomes is a critical area of research in both clinical infectious disease and oncology alike. In this work, we employed a Markov Decision Process (MDP), a mathematical framework we have previously used to identify optimal drug policy that minimizes resistance over the course of treatment [8, 30, 35]. In an MDP, the system transitions stochastically between discrete states (in our case, genotypes), and at each time step. an action (in our case, drug choice) is taken. This action determines both the immediate reward and influences the subsequent state of the system. The goal of the MDP is to compute a policy that optimizes a predefined objective function. Applying this framework to the fitness landscapes measured by Mira et al., we calculated an optimal policy that identifies the best drug to administer at each genotype to minimize population fitness over a specified time period.

To better understand the characteristics of drugs selected by the MDP's optimal policy, we examined the rank of the fitness of the genotype and the genotype net flow index (calculated as the ratio of less-fit neighbors to fitter neighbors minus one) under the selected drug. The genotype net flow index quantifies the evolutionary flow under drug selection: a value of -1 indicates all edges connected to the genotype lead to fitter neighbors (a local fitness minimum), while a value of 1 indicates all edges point to the genotype from less fit neighbors (a local fitness maximum). For the Mira et al. empirical landscapes, the MDP selects the drug associated with the lowest fitness among available options for 75% of genotypes (12 out of the 16), as shown in **Figure 4A**. In the remaining 25% of cases where the MDP does not select the lowest-fitness drug, 75% of these genotypes (3 out of 4) are local fitness minima (**Figure 4B**). By generating 1000 sets of procedurally generated fitness landscapes, each

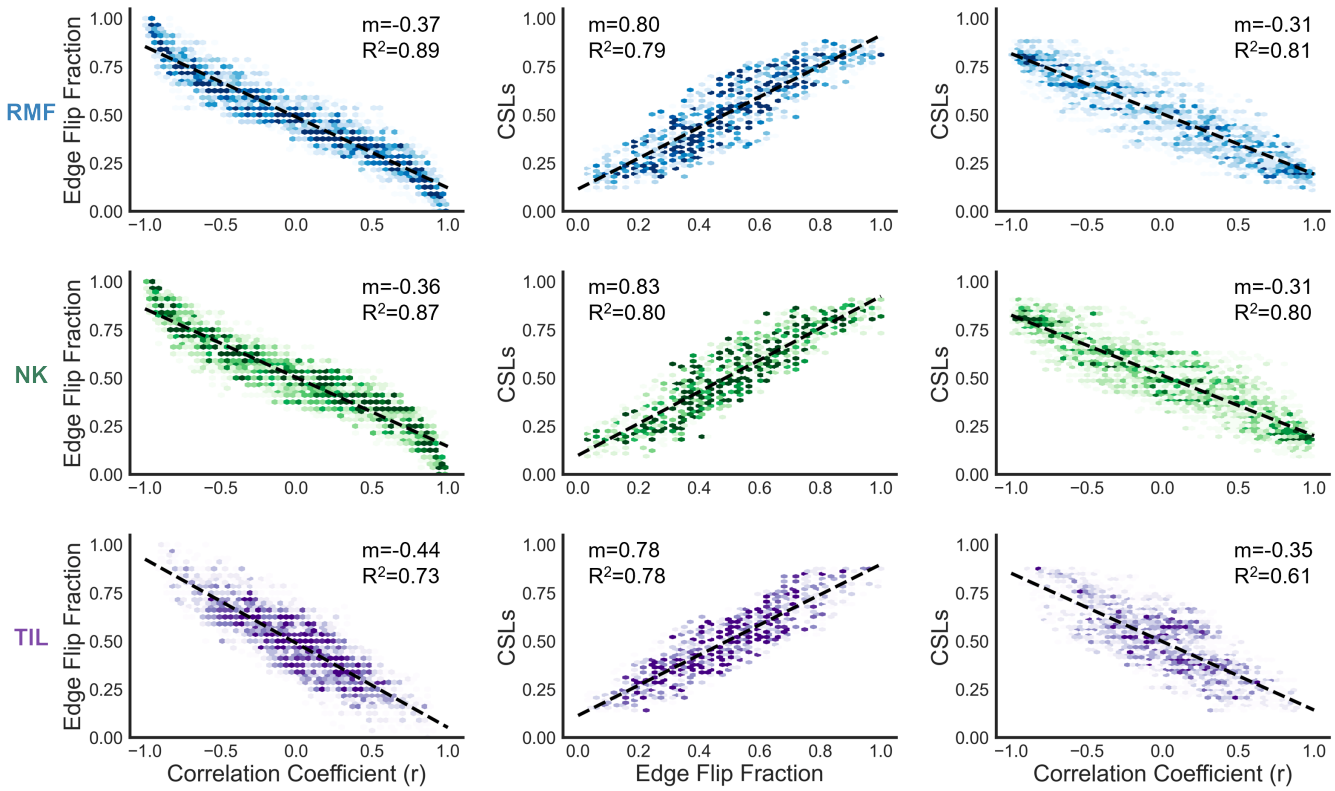

Figure 3. Edge flip fraction ($\mathscr{E}$), collateral sensitivity likelihood (CSLs), and Pearson correlation coefficient (r), all partially explain the variance within each other in synthetic fitness landscapes – including the rough Mount Fuji (RMF), Kauffman NK, and trade-off-induced fitness landscape (TIL) [11, 50, 52–54]. Hexbin plots show the density of data points, with darker regions indicating higher density, and dashed lines represent linear regression fits for the displayed relationships. As expected, $\mathscr{E}$ and CSLs and are positively correlated, while both $\mathscr{E}$ and r, CSLs and R are negatively correlated.

containing 15 fitness landscapes of $N = 4$, we further explored this behavior. Across these stimulated landscapes, an average of 71% of MDP-selected drugs were the lowest-fitness option, and 60% of non-lowest-fitness drug choices occurred at local minima (**Figure S2**). These findings suggest a consistent tendency for the MDP to prioritize drugs that minimize fitness or target genotypes near local minima.

To elucidate the relationship between the edge flip fraction and the optimality of drug sequences, we simulated 1000 drug sequences of length-20 by repeatedly applying the MDP optimal drug policy. For each drug sequence, we computed both global edge flips - occurring across the entire landscapes - and local edge flips specific to the current genotype when switching between drugs. As a contrasting baseline, we also stimulated a "worse-case" MDP policy designed to maximize fitness over time. We then compared global and local edge flips across four sequence types: optimal drug sequences, shuffled optimal sequences (using the same drugs as the optimal sequence but in randomized order), random drug sequences (drugs chosen randomly from the full set), and the worst-case sequences (**Figure 5**).

The results revealed that the worst drug sequences exhibited the lowest fraction of global and local edge flip, concomitant with the highest fitness levels. This outcome arises because the worst MDP policy tends to switch drugs less often, allowing the genotype evolve toward the highest fitness peak more rapidly. While global edge flip fraction did not show significant differences among the optimal, shuffled optimal and random drug sequence - since the MDP optimizes the drug choice for each genotype rather than altering the entire fitness landscape - there were significantly more local edge flips observed in the optimal drug sequences. This phenomenon occurs because the MDP tends to select drugs that drive the genotype into unstable local minima to maintain lower fitness levels. In subsequent evolutionary steps, the genotype is likely to evolve stochastically to a fitter neighbor. The optimal drugs applied thereafter tend to reverse evolutionary direction, resulting a local edge flip.

By inducing oscillations between genotypes near local minima, the optimal MDP policy effectively traps the pathogen in less favorable states. This strategy leverages the topological features of the fitness landscapes, exploiting the arrangement of local minima and the connectivity between genotypes to restrict adaptive pathways. The increased fraction of edge flips in the optimal drug sequences signifies a deliberate manipulation of evolutionary pathways to hinder the development of higher

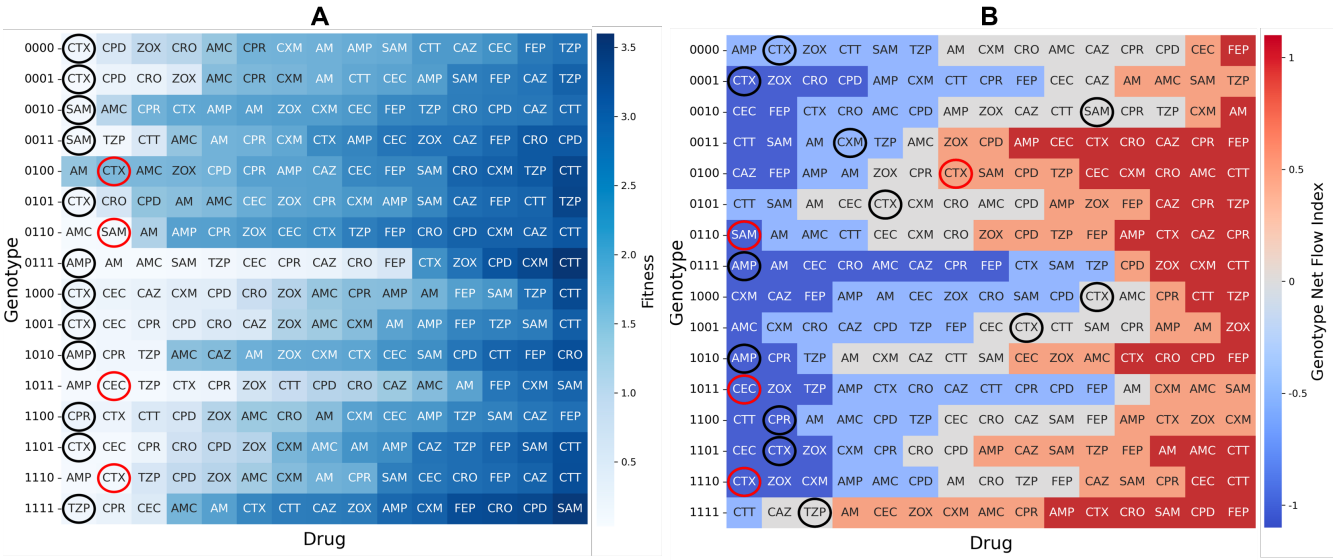

Figure 4. A. Heatmap of fitness values for 16 genotype under various drugs. The color gradient from light blue to dark blue reflects fitness levels, with darker shades representing higher fitness. Drugs selected by optimal MDP are circled: black circles denote chosen drugs with the lowest fitness, and red circles denote chosen drugs without the lowest fitness. **B.** Heatmap of the Genotype Net Flow Index for 16 genotypes across different drugs. A value of -1 (blue) represents a local fitness minima, where all edges connected to the genotype lead to fitter neighbors, while a value of 1 (red) corresponds to a local fitness maxima, where all edges point toward the genotype from less fit neighbors. As in heatmap A, drugs selected by optimal MDP are marked with black circles for the lowest-fitness choice and red circles otherwise.

fitness genotypes. These findings underscore the potential of adaptive drug sequencing strategies in controlling the evolution of drug resistance.

Discussion

The fitness landscape model of evolution is a very useful tool for understanding how selective pressures shape population genotypic frequencies over time. As a mathematical model, fitness landscapes provide a framework for how evolutionary histories impact the subsequent phenotype of a population, and the evolutionary forces involved. By representing a fitness landscape as a directed hypergraph, the relationship between two fitness landscapes applied in succession can be characterized by the change in these edge directions. While the Pearson correlation coefficient captures proportional changes in fitness values across genotypes, it does not account for alterations in evolutionary pathways - the directions in which populations are likely to evolve. Our novel metric, edge flip fraction addresses this gap by quantifying the extent to which these evolutionary direction differ between landscape, which holds practical implications for understanding and controlling the evolution of drug resistance.

Our analysis of empirical fitness landscape from Mira et al. further reinforces the utility of edge flip fraction. Drug pairs that can be used to bind a population initially at a high fitness peak to a lower fitness peak with the administration of a second drug, known as a collaterally sensitive drug pair, require changes in edge direction. However, the edge flip fraction alone is likely not sufficiently predictive of collateral sensitivity likelihood.

Our investigation into the relationship between edge flip fraction and drug sequence optimality using Markov Decision Process (MDPs) further emphasizes the practical relevance of our findings. MDPs allowed us to simulate optimal drug policies; the optimal drug sequences that minimize the fitness over time, as identifies by the MDPs, were characterized by a higher fraction of local edge flips compared to non-optimal sequences. This suggests that effective treatment strategies often involve inducing changes in evolutionary direction, thereby trapping the population in less favorable fitness states.

An intriguing avenue for future research involves extending the concept of edge flip fraction to fitness seascapes, where environmental conditions change continuously rather than abruptly between distinct drugs [29]. If we consider similar environmental conditions with slight differences - instead of two distinct condition - such as the same drug at different dose, changes in edge direction between vertices of a fitness landscape can convey more nuanced information. For instance, if the fitness of two genotypes switches rank order with increasing dose of a drug, and these genotypes have mutational distance (i.e. hamming distance) of one, a edge flip must have occurred [55]. In this way, changes in evolutionary direction between

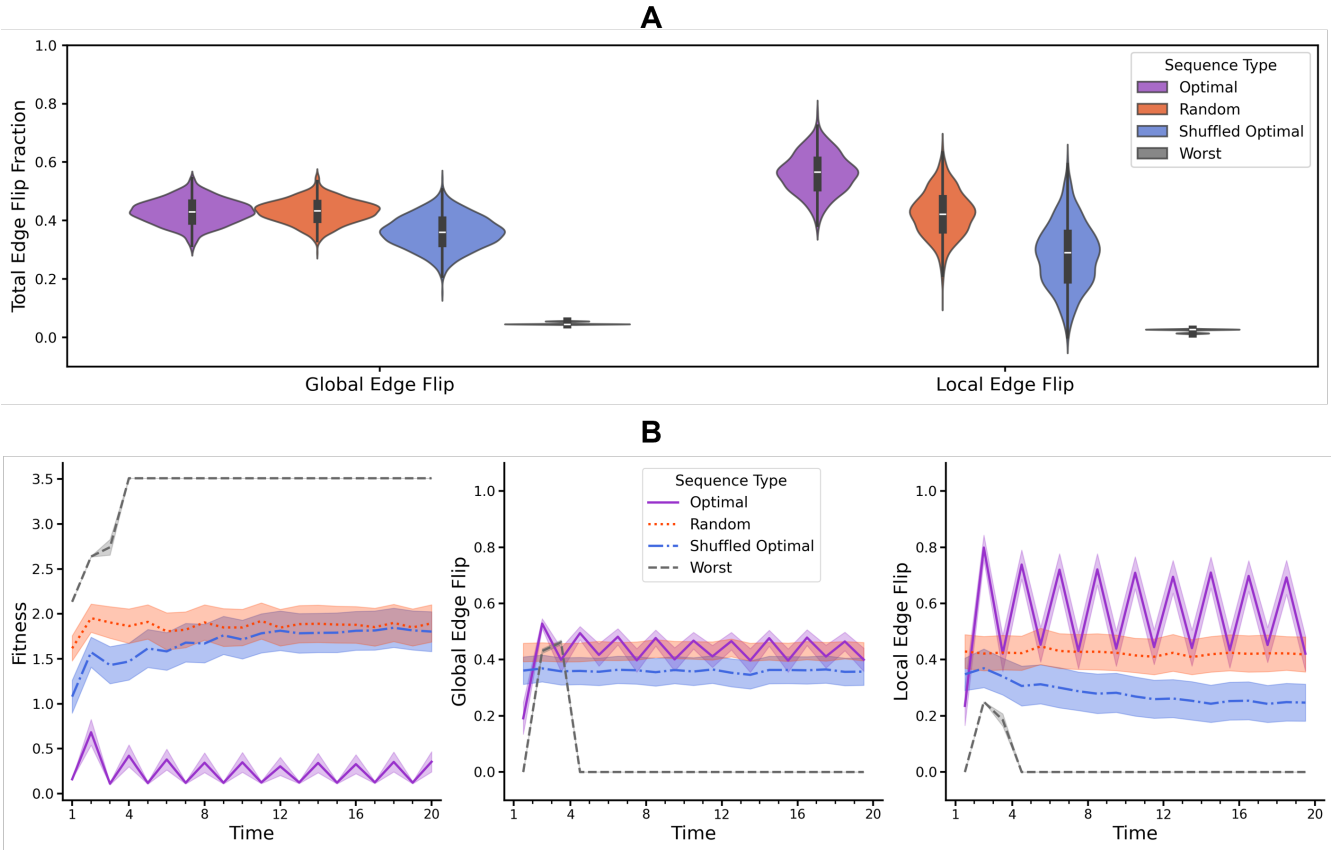

Figure 5. Comparison of edge flip dynamics and fitness across optimal, shuffled optimal, random and worst drug sequences. **A.** Violin plots showing the distribution of total edge flip fractions for global edge flip (across the entire landscape) and local edge flip (specific to the current genotype) across four drug sequence types: Optimal, Random, Shuffled Optimal, Random and Worst. The worst sequences show significantly lower global and local edge flip fraction compared to other sequences. While global edge flips do not differ significantly among optimal, shuffled optimal and random sequences, optimal drug sequences exhibit notably higher local edge flips compared to other sequence types. **B.** Time-series line plots tracking changes in fitness, global edge flips, and local edge flips over time. Optimal sequences (purple) maintain low fitness values with high local edge flip in cyclical dynamics. Shuffled optimal (blue) and random (orange) sequences show moderate fitness levels with moderate local edge flips. The worst sequence (gray) consistently exhibits the highest fitness and minimal edge flipping. Shaded areas indicate the standard error of the mean across replicates.

distinct conditions can relate to topological similarity, where as edge direction changes with changing of some continuous environmental variable like drug dose or temperature can be the result of a fitness cost.

The analysis presented here establishes how the fitness landscape analogy applied to theoretical and empirical landscapes alike can provide insights into evolution and can facilitate improved evolutionary control techniques.

References

1. Koel, B. F. *et al.* Substitutions near the receptor binding site determine major antigenic change during influenza virus evolution. *Science* **342**, 976–979 (2013).

2. Kreiner, J. M. *et al.* Rapid weed adaptation and range expansion in response to agriculture over the past two centuries. *Science* **378**, 1079–1085 (2022).

3. Hawkins, N. J., Bass, C., Dixon, A. & Neve, P. The evolutionary origins of pesticide resistance. *Biol. Rev.* **94**, 135–155 (2019).

4. Promdonkoy, P., Mhuantong, W., Champreda, V., Tanapongpipat, S. & Runguphan, W. Improvement in d-xylose utilization

and isobutanol production in s. cerevisiae by adaptive laboratory evolution and rational engineering. *J. Ind. Microbiol. Biotechnol.* **47**, 497–510 (2020).

5. Dixon, G. B. *et al.* Genomic determinants of coral heat tolerance across latitudes. *Science* **348**, 1460–1462 (2015).

6. MacDougall, A. S., McCann, K. S., Gellner, G. & Turkington, R. Diversity loss with persistent human disturbance increases vulnerability to ecosystem collapse. *Nature* **494**, 86–89 (2013).

7. Iram, S. *et al.* Controlling the speed and trajectory of evolution with counterdiabatic driving. *Nat. Phys.* **17**, 135–142 (2021).

8. Weaver, D. T., King, E. S., Maltas, J. & Scott, J. G. Reinforcement learning informs optimal treatment strategies to limit antibiotic resistance. *Proc. Natl. Acad. Sci.* **121**, e2303165121, DOI: 10.1073/pnas.2303165121 (2024).

9. Gillespie, J. H. A simple stochastic gene substitution model. *Theor. population biology* **23**, 202–215 (1983).

10. Gillespie, J. H. Molecular evolution over the mutational landscape. *Evolution* 1116–1129 (1984).

11. Maltas, J., McNally, D. M. & Wood, K. B. Evolution in alternating environments with tunable interlandscape correlations. *Evolution* **75**, 10–24 (2021).

12. Tan, L., Serene, S., Chao, H. X. & Gore, J. Hidden randomness between fitness landscapes limits reverse evolution. *Phys. review letters* **106**, 198102 (2011).

13. Kaznatcheev, A. Computational complexity as an ultimate constraint on evolution. *Genetics* **212**, 245–265 (2019).

14. Wagner, A. Genotype sampling for deep-learning assisted experimental mapping of a combinatorially complete fitness landscape. *Bioinformatics* **40**, btae317 (2024).

15. Dadonaite, B. *et al.* Full-spike deep mutational scanning helps predict the evolutionary success of sars-cov-2 clades. *biorxiv* (2023).

16. Das, S. G., Direito, S. O., Waclaw, B., Allen, R. J. & Krug, J. Predictable properties of fitness landscapes induced by adaptational tradeoffs. *Elife* **9**, e55155 (2020).

17. Papkou, A., Garcia-Pastor, L., Escudero, J. A. & Wagner, A. A rugged yet easily navigable fitness landscape. *Science* **382**, eadh3860 (2023).

18. Weinreich, D. M., Delaney, N. F., DePristo, M. A. & Hartl, D. L. Darwinian evolution can follow only very few mutational paths to fitter proteins. *science* **312**, 111–114 (2006).

19. Greenbury, S. F., Louis, A. A. & Ahnert, S. E. The structure of genotype-phenotype maps makes fitness landscapes navigable. *Nat. Ecol. & Evol.* **6**, 1742–1752 (2022).

20. De Visser, J. A. G. & Krug, J. Empirical fitness landscapes and the predictability of evolution. *Nat. Rev. Genet.* **15**, 480–490 (2014).

21. Johnson, M. S., Reddy, G. & Desai, M. M. Epistasis and evolution: recent advances and an outlook for prediction. *BMC biology* **21**, 120 (2023).

22. Blount, Z. D., Lenski, R. E. & Losos, J. B. Contingency and determinism in evolution: Replaying life's tape. *Science* **362**, eaam5979 (2018).

23. Bajić, D., Vila, J. C., Blount, Z. D. & Sánchez, A. On the deformability of an empirical fitness landscape by microbial evolution. *Proc. Natl. Acad. Sci.* **115**, 11286–11291 (2018).

24. Diaz-Colunga, J., Sanchez, A. & Ogbunugafor, C. B. Environmental modulation of global epistasis in a drug resistance fitness landscape. *Nat. Commun.* **14**, 8055 (2023).

25. Diaz-Colunga, J. *et al.* Global epistasis on fitness landscapes. *Philos. Transactions Royal Soc. B* **378**, 20220053 (2023).

26. Ardell, S. M., Martsul, A., Johnson, M. S. & Kryazhimskiy, S. Environment-independent distribution of mutational effects emerges from microscopic epistasis. *Science* **386**, 87–92 (2024).

27. Starr, T. N. & Thornton, J. W. Epistasis in protein evolution. *Protein science* **25**, 1204–1218 (2016).

28. Da Silva, J., Coetzer, M., Nedellec, R., Pastore, C. & Mosier, D. E. Fitness epistasis and constraints on adaptation in a human immunodeficiency virus type 1 protein region. *Genetics* **185**, 293–303 (2010).

29. Ogbunugafor, C. B., Wylie, C. S., Diakite, I., Weinreich, D. M. & Hartl, D. L. Adaptive landscape by environment interactions dictate evolutionary dynamics in models of drug resistance. *PLoS computational biology* **12**, e1004710 (2016).

30. Maltas, J. & Wood, K. B. Pervasive and diverse collateral sensitivity profiles inform optimal strategies to limit antibiotic resistance. *PLoS biology* **17**, e3000515 (2019).

31. McBride, R. C., Ogbunugafor, C. B. & Turner, P. E. Robustness promotes evolvability of thermotolerance in an rna virus. *BMC evolutionary biology* **8**, 1–14 (2008).

32. Imamovic, L. & Sommer, M. O. Use of collateral sensitivity networks to design drug cycling protocols that avoid resistance development. *Sci. translational medicine* **5**, 204ra132–204ra132 (2013).

33. Carolus, H. *et al.* Collateral sensitivity counteracts the evolution of antifungal drug resistance in candida auris. *Nat. Microbiol.* **9**, 2954–2969 (2024).

34. Maltas, J., Krasnick, B. & Wood, K. B. Using selection by nonantibiotic stressors to sensitize bacteria to antibiotics. *Mol. biology evolution* **37**, 1394–1406 (2020).

35. Maltas, J., Huynh, A. & Wood, K. B. Dynamic collateral sensitivity profiles highlight challenges and opportunities for optimizing antibiotic sequences. *bioRxiv* 2021–12 (2021).

36. Barbosa, C. *et al.* Alternative evolutionary paths to bacterial antibiotic resistance cause distinct collateral effects. *Mol. biology evolution* **34**, 2229–2244 (2017).

37. Nichol, D. *et al.* Antibiotic collateral sensitivity is contingent on the repeatability of evolution. *Nat. communications* **10**, 1–10 (2019).

38. Wang, S. & Dai, L. Evolving generalists in switching rugged landscapes. *PLoS computational biology* **15**, e1007320 (2019).

39. Dhawan, A. *et al.* Collateral sensitivity networks reveal evolutionary instability and novel treatment strategies in alk mutated non-small cell lung cancer. *Sci. reports* **7**, 1–9 (2017).

40. Scarborough, J. A. *et al.* Identifying states of collateral sensitivity during the evolution of therapeutic resistance in ewing's sarcoma. *Iscience* **23** (2020).

41. Acar, M., Mettetal, J. T. & Van Oudenaarden, A. Stochastic switching as a survival strategy in fluctuating environments. *Nat. genetics* **40**, 471–475 (2008).

42. Cook, R. D. & Hartl, D. L. Uncorrelated random environments and their effects on gene frequency. *Evolution* 265–274 (1974).

43. Mustonen, V. & Lässig, M. Molecular evolution under fitness fluctuations. *Phys. review letters* **100**, 108101 (2008).

44. Mustonen, V. & Lässig, M. From fitness landscapes to seascapes: non-equilibrium dynamics of selection and adaptation. *Trends genetics* **25**, 111–119 (2009).

45. Kim, S., Lieberman, T. D. & Kishony, R. Alternating antibiotic treatments constrain evolutionary paths to multidrug resistance. *Proc. Natl. Acad. Sci.* **111**, 14494–14499 (2014).

46. Yoon, N., Krishnan, N. & Scott, J. Theoretical modeling of collaterally sensitive drug cycles: shaping heterogeneity to allow adaptive therapy. *J. mathematical biology* **83**, 47 (2021).

47. West, J. *et al.* Towards multidrug adaptive therapy. *Cancer research* **80**, 1578–1589 (2020).

48. Ogbunugafor, C. B. & Eppstein, M. J. Competition along trajectories governs adaptation rates towards antimicrobial resistance. *Nat. ecology & evolution* **1**, 0007 (2016).

49. Nichol, D. *et al.* Steering evolution with sequential therapy to prevent the emergence of bacterial antibiotic resistance. *PLoS Comput. Biol* **11**, e1004493 (2015).

50. Mira, P. M., Meza, J. C., Nandipati, A. & Barlow, M. Adaptive landscapes of resistance genes change as antibiotic concentrations change. *Mol. biology evolution* **32**, 2707–2715 (2015).

51. Aita, T. & Husimi, Y. Adaptive Walks by the Fittest among Finite Random Mutants on a Mt. Fuji-type Fitness Landscape. *J. Theor. Biol.* **193**, 383–405, DOI: 10.1006/jtbi.1998.0709 (1998).

52. Neidhart, J., Szendro, I. G. & Krug, J. Adaptation in tunably rugged fitness landscapes: the rough mount fuji model. *Genetics* **198**, 699–721 (2014).

53. Kauffman, S. A. & Weinberger, E. D. The NK model of rugged fitness landscapes and its application to maturation of the immune response. *J. Theor. Biol.* **141**, 211–245, DOI: 10.1016/S0022-5193(89)80019-0 (1989).

54. Das, S. G., Direito, S. O., Waclaw, B., Allen, R. J. & Krug, J. Predictable properties of fitness landscapes induced by adaptational tradeoffs. *eLife* **9**, e55155, DOI: 10.7554/eLife.55155 (2020).

55. King, E. S., Tadele, D. S., Pierce, B., Hinczewski, M. & Scott, J. G. Diverse mutant selection windows shape spatial heterogeneity in evolving populations. *PLOS Comput. Biol.* **20**, e1011878, DOI: 10.1371/journal.pcbi.1011878 (2024).

Supplementary Material

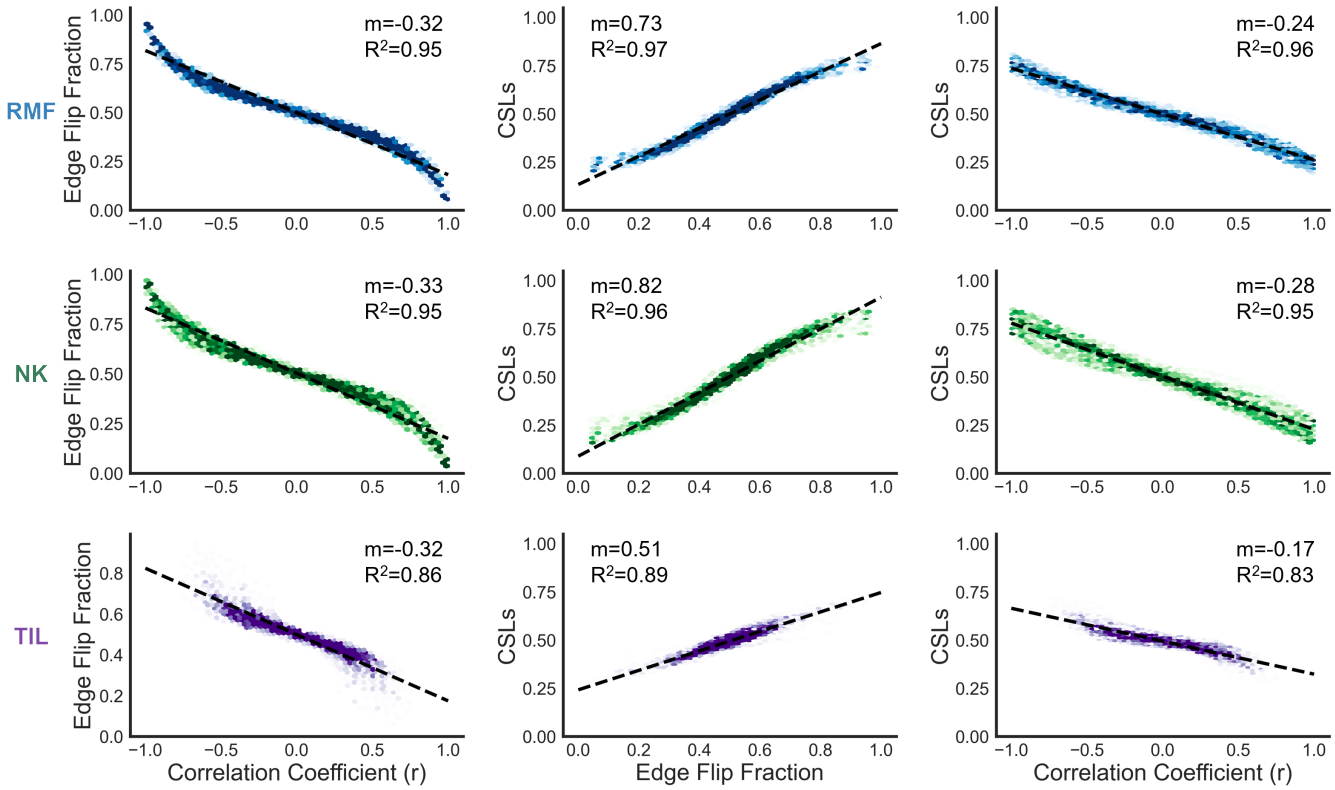

Figure S1. Comparison of synthetic $N = 8$ fitness landscapes generated using the rough Mount Fuji (RMF), Kauffman NK model, and trade-off-induced landscapes (TIL) [11, 50, 52–54]. The observed relationships - $\mathscr{E}$ and CSLs being positively correlated, while $\mathscr{E}$ and r, as well as CSLs and R, being negatively correlated - remain consistent in pairs of higher-dimensional fitness landscapes.

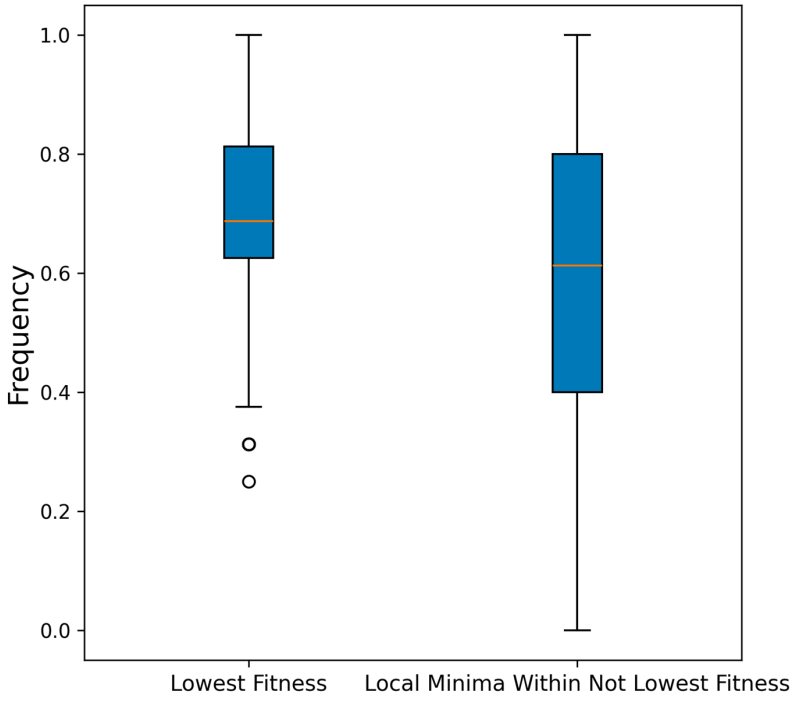

Figure S2. Boxplot depicting the frequency distribution of MDP-selected drugs being the lowest-fitness option (left) and non-lowest-fitness drug choices occurring at local minima (right) across 1000 sets of fitness landscapes. Each set contains 15 fitness landscapes with $N = 4$. On average, 71% of MDP-selected drugs were the lowest-fitness option, while 60% of non-lowest-fitness choices corresponded to local minima."

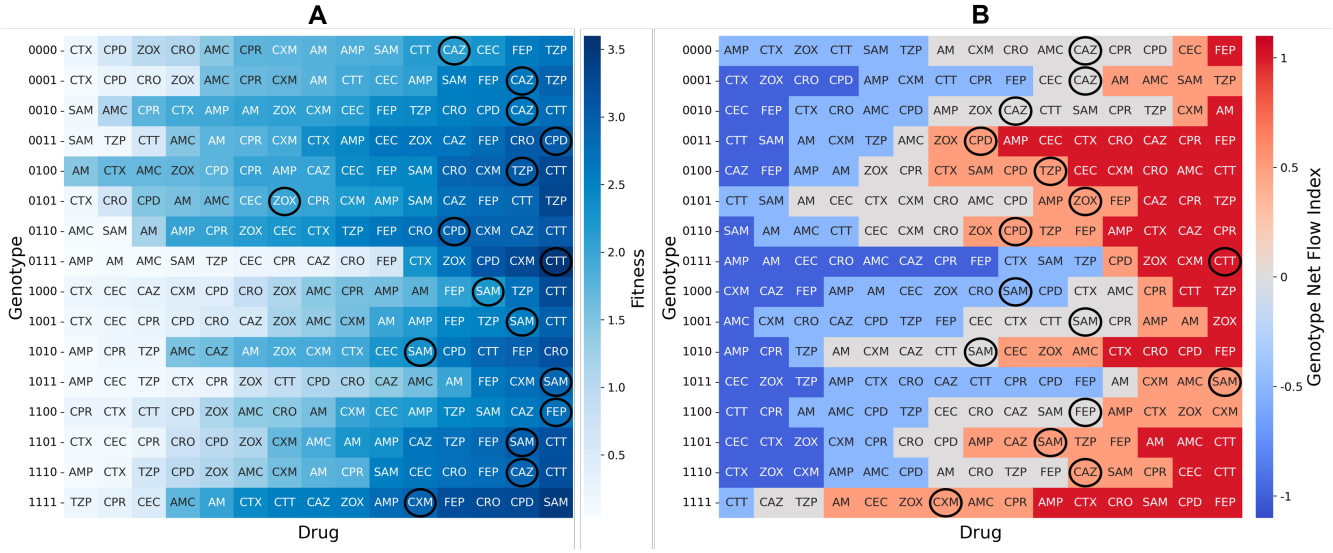

Figure S3. **A.** Heatmap of fitness values for 16 genotype under various drugs. The color gradient from light blue to dark blue reflects fitness levels, with darker shades representing higher fitness. Drugs chosen by the "worst-case" MDP policy for each genotype are highlighted with black circles. **B.** Heatmap of the Genotype Net Flow Index for 16 genotypes across different drugs. A value of -1 (blue) represents a local fitness minima, where all edges connected to the genotype lead to fitter neighbors, while a value of 1 (red) corresponds to a local fitness maxima, where all edges point toward the genotype from less fit neighbors. Drugs selected by "worst-case" MDP policy are marked with circles.