Simulation for Steering Evolution

Sebastian Rene Towers

Reuben College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Hilary 2026

Abstract

Whilst conceptually simple, evolution is a powerful force that pushes biological objects toward superior variants. Or it would, if the world were more simple. In reality, there is no objective notion of "superior": natural selection is immensely complex, representing an interplay between organisms, their environment, and the underlying genetic landscape, gently favouring organisms which can survive and reproduce in a particular environment, at a particular point in time. This complexity makes evolution difficult to study, but advances in both high-throughput experimental systems and deep learning have expanded our understanding of the space of genetic possibilities more than ever before. Given these insights, a natural question arises: can we not only understand evolution, but actively *steer* it?

This thesis is devoted to the question of how evolutionary systems can be manipulated to behave to our advantage. The investigation is carried out *in silico*, leveraging modern hardware and scientific computing tools to enable large-scale computational experiments. The work proceeds in two parts. The first focuses on *accelerating* evolution for the development of improved proteins; the second on *constraining* evolution, specifically by limiting the rate of viral antigenic escape through appropriate antibody choice.

In the first part, on directed evolution, we introduce the problem setting in **Optimisation strategies for directed evolution without sequencing**, demonstrating how alternative experimental strategies can yield superior outcomes in fewer generations. This is complemented by **Sequence-free landscape inference for directed evolution**, where we develop methods to infer properties of the underlying genetic landscape in order to guide strategic decision-making.

The second part consists of **ADIOS: Antibody Development via Opponent Shaping**, where ideas from multi-agent reinforcement learning are applied to the problem of viral evolution. In a mechanistic model of viral adaptation, we show that antibody choice can dramatically influence the rate of viral evolution, and that an antibody that performs well over long time horizons may differ substantially from one that performs best against the initial virus.

Together, these works show how evolution can be steered in the systems studied, providing a first step towards understanding how to control real-world evolutionary systems. They offer concrete computational methods and experimentally testable strategies, opening a path towards steering increasingly realistic models and, ultimately, real-world systems, including evolving viruses.

Simulation for Steering Evolution



Sebastian Rene Towers

Reuben College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Hilary 2026

Acknowledgements

To my supervisors, Harrison Steel and Jakob Foerster: I could not have asked for a better combination of brilliant, sparky supervisors. Jakob, a source of relentless dynamic energy and ambitious optimism, always directing me towards the most daring ideas. Harrison, an admirable, wise stoic with a cunning wit, doing your best to contain and refine my own chaotic nature. To you both, thank you for all your advice, your wisdom, and your good humour.

To my close collaborators, Ola, Jess, and Peer. All the major works in my PhD were in joint first authorship, a decision which ensured a successful and highly enjoyable four years. This is no doubt thanks to being able to work with such intelligent, resilient, and likeable people. We have shared a lot of highs and lows together, and I am immensely thankful to have such excellent companions on this journey.

To all the people in both the Steel Group and FLAIR: I must say I am sincerely amazed by the density of both intellect and charming personality in both groups. I can firmly say that you are all destined for great things. Thank you for the camaraderie, support, and putting up with my endless odd lunchtime conversations.

To the 2021 AIMS student cohort, I could not have asked for a finer group of people to do the CDT with. I look back on our time together with deep nostalgia, as I am sure we all do.

Thank you to the AIMS CDT for funding my PhD, and in particular Wendy Poole, the legendary administrator who ensures the success of the CDT. Everyone in the CDT is truly blessed by your presence.

To my parents, Renata and Paul, and my sister, Stazie, who have all given me unconditional love and support. Without you, I would never have been close to starting on this path, and I dearly love you all.

Finally, to Duygu, my love. These years with you have been the best of my life, and you have been a bottomless well of love, support, and comfort. Whenever I talk about you, which is extremely often, my adoration for you is obvious, and I intend for this paragraph to be no different. So, thank you for being my funny, powerful, smart, outrageous, sharp, loud, lifelong partner. I love you.

Abstract

Whilst conceptually simple, evolution is a powerful force that pushes biological objects toward superior variants. Or it would, if the world were more simple. In reality, there is no objective notion of "superior": natural selection is immensely complex, representing an interplay between organisms, their environment, and the underlying genetic landscape, gently favouring organisms which can survive and reproduce in a particular environment, at a particular point in time. This complexity makes evolution difficult to study, but advances in both high-throughput experimental systems and deep learning have expanded our understanding of the space of genetic possibilities more than ever before. Given these insights, a natural question arises: can we not only understand evolution, but actively *steer* it?

This thesis is devoted to the question of how evolutionary systems can be manipulated to behave to our advantage. The investigation is carried out *in silico*, leveraging modern hardware and scientific computing tools to enable large-scale computational experiments. The work proceeds in two parts. The first focuses on *accelerating* evolution for the development of improved proteins; the second on *constraining* evolution, specifically by limiting the rate of viral antigenic escape through appropriate antibody choice.

In the first part, on directed evolution, we introduce the problem setting in **Optimisation strategies for directed evolution without sequencing**, demonstrating how alternative experimental strategies can yield superior outcomes in fewer generations. This is complemented by **Sequence-free landscape inference for directed evolution**, where we develop methods to infer properties of the underlying genetic landscape in order to guide strategic decision-making.

The second part consists of **ADIOS: Antibody Development via Opponent Shaping**, where ideas from multi-agent reinforcement learning are applied to the problem of viral evolution. In a mechanistic model of viral adaptation, we show that antibody choice can dramatically influence the rate of viral evolution, and that an antibody that performs well over long time horizons may differ substantially from one that performs best against the initial virus.

Together, these works show how evolution can be steered in the systems studied, providing a first step towards understanding how to control real-world evolutionary systems. They offer concrete computational methods and experimentally testable strategies, opening a path towards steering increasingly realistic models and, ultimately, real-world systems, including evolving viruses.

Contents

1 Introduction **1**

2 Background **7**

2.1 Proteins, Folding and Binding 7
2.2 Antibodies 9
2.3 Directed Evolution 11
2.4 Fitness Landscape Theory 12
2.4.1 Ruggedness 13
2.4.2 Empirical and Model Landscapes 14
2.5 Opponent Shaping 16
2.6 JAX and Hardware Acceleration 17
2.7 Platforms for High-Throughput Single-Cell Observation 18

3 Sequence-Free Directed Evolution **21**

3.1 Motivation 21
3.2 Summary 22
3.3 Contributions 24
3.4 Paper 24
3.5 Limitations and Future Work 47
3.5.1 Limitations of the NK Model 47
3.5.2 Empirical Validation and Scale 48
3.5.3 Strategy Selection and Ruggedness Inference 48

4 Parameter Inference for Directed Evolution **49**

4.1 Motivation 49
4.2 Summary 50
4.3 Contributions 51
4.4 Paper 52
4.5 Limitations and Future work 93
4.5.1 Experimental Viability 93
4.5.2 NK landscapes 95
4.5.3 Ruggedness Metrics 95

5 Antibody Shaping **99**
5.1 Motivation 99
5.2 Summary 100
5.3 Contributions 102
5.4 Paper 102
5.5 Limitations and Future work 123

6 Conclusion **125**
6.1 Summary 125
6.2 Limitations and Future work 127
6.3 Future Direction 129
6.4 Final Remarks 130

References **131**

The only joy in the world is to begin.

Cesare Pavese

1 Introduction

Over the past several decades there has been exponential growth in the capabilities of experimental methods for creating, measuring, and analysing biological information [1–3]. We are now able to synthesise and functionally characterise millions of variants of a biological part or genetic circuit [4], opening the potential of linking high-throughput measurements of sequence-to-function relationships. Meanwhile, tools from AI have also been increasingly applied to biological questions [5–7], enabled by advances in fundamental AI methodologies as well as approaches to incorporate biophysical constraints into models to allow them to embed assumptions about fundamental properties of biology [8, 9].

Taking a step back, advances in ML have benefited substantially from the effective use of hardware [10]. Hardware accelerators such as GPUs can offer orders-of-magnitude speedups over CPU implementations for effectively parallelised calculations [11]. Toolkits for making ML models run effectively on accelerated hardware, such as JAX [12], are also being used in scientific computation, including molecular dynamics and cosmology [13, 14]. Evolutionary simulations are natural candidates for utilising such toolkits: evolution is effectively already a parallel search over the space of genotypes.

For an area which demonstrates how understanding evolutionary systems can translate to real-world benefits, consider directed evolution. It is a broad family of approaches which aim to emulate biological evolution to improve the functions of biomolecules or cells for human-defined goals [15]. The basic idea of directed evolution is to test a wide variety of proteins (or cells, or any kind of biological entity) for some application, pick the best, make many slightly altered copies of the best performer, test these, and repeat ad infinitum.

Originally developed as a tool for screening and enriching mutants-of-interest in bulk assays of different biological materials [16], directed evolution has subsequently been applied to a multitude of applications. These include improving specificity of genetic circuits [17], optimisation of cellular chassis to grow in different environments [18], and implementing biosynthesis pathways [19].

Directed evolution follows the same broad loop and inspiration as computational evolutionary algorithms, but takes place in a wet lab rather than in silico. The primary difference is that directed evolution needs to respect real-world biological constraints. For example, in a computer, effectively any variant may be tested, but in the wet lab, creating a specific variant may be challenging, and it is often easier to perform certain biologically enabled mutation operations from current variants than specify a variant directly. Going further, in directed evolution, it is often non-trivial to find which genotype resulted in a certain phenotype as genetic sequencing is somewhat expensive. Meanwhile, in a computer algorithm, finding the genotype is as simple as looking in the computer memory. It is precisely these kinds of differences which motivate a new set of methods for understanding directed evolution.

Going beyond directed evolution, a specific subset of (directed) evolutionary challenges is concerned with co-evolution of two (or more) molecules/cells in which the fitness (or more broadly, function) of one has an impact on the function or

fitness of the other [20, 21]. Importantly, this also includes pathogenic viruses [22, 23], which cause enormous social and economic damage [24–27].

To summarise, this thesis aims to better understand and guide evolutionary systems with accelerated computation and simulation. This is undertaken in two primary strands:

- **Directed Evolution as Optimisation:** New experimental approaches to directed evolution now enable us to implement selective processes *directly*, rather than requiring direct coupling of traits of interest to growth of a cell or molecule. This research direction investigates how to optimally employ this new technique, for example in the context of how one would implement a selective strategy if it were possible to select based on *any* criteria defined on a specific set of measurements.

- **Directed Co-evolution:** Motivated by AI tools that (for example) allow effective structural prediction from sequence data, we develop methods to understand and potentially design biological systems which are optimised in terms of their predicted (co)-evolutionary trajectories. As a model, we look at the evolutionary system consisting of a pathogenic virus and antibody, to understand how the evolutionary trajectory of a virus may depend on the selective pressure induced by an antibody.

This thesis addresses these strands via three papers, **Optimisation strategies for directed evolution without sequencing**, **Sequence-free landscape inference for directed evolution**, and **ADIOS: Antibody Development via Opponent Shaping**. These form chapters 3, 4, and 5, respectively.

Together, these papers address three questions:

1. How can we frame directed evolution as an optimisation problem?

2. How can we infer useful properties of a fitness landscape from population-level measurements?

3. Can we apply principles from opponent shaping, anticipating and influencing an opponent's adaptation, to viral evolution?

Chapter 3 develops a simulation-first perspective on sequencing-free directed evolution, framing directed evolution as a controllable optimisation process. Via large-scale GPU-based simulations we explore how non-greedy selection schedules, population splitting, and multi-trait objectives perform on a range of landscapes. With this, we demonstrate that simple adjustments, such as allowing low-fitness variants to persist or running several smaller populations in parallel, yield dramatic improvements on both mathematical and empirical landscapes. The resulting design rules provide concrete strategies for experimental testing using new high-throughput mutation–selection technologies.

Chapter 4 builds on this optimisation view by asking how we can diagnose the landscape before we start evolving. It introduces a sequence-free ruggedness inference method grounded in spectral (Fourier) landscape theory, using the dynamics of a randomly mutating population to estimate a scalar measure of spectral ruggedness. By mapping this measure onto mathematical landscapes, the paper provides a principled bridge between measurable population statistics and the simulated strategy library established in Chapter 3, allowing practitioners to choose mutation rates, selection stringency, and population structures that are tailored to their specific system.

Chapter 5 pivots from accelerating evolution to constraining it, presenting the ADIOS framework for antibody design under adversarial viral adaptation. By optimising antibodies against simulated long-horizon viral escape trajectories, the paper shows, in simulation, that "shaper" antibodies can remain potent even as the virus adapts and can shepherd viral evolution toward variants that are easier

for diverse antibodies to bind. Robustness checks with mixed selective pressures highlight that opponent shaping is a promising route toward therapies resilient to the evolutionary future, not just the present.

Taken together, these contributions chart a path from efficient search procedures, through data-efficient landscape characterisation, to proactive control of co-evolutionary dynamics. They illustrate how hardware-accelerated modelling can both expedite the discovery of beneficial molecules and harden our defences against rapidly evolving pathogens.

There are a couple of additional lines of work that merit a brief mention.

Firstly, **Computational Synthetic Biology Enabled through JAX: A Showcase** [28], by Olivia Gallup, Kirill Sechkar, and Sebastian Towers (under the supervision of Harrison Steel), published in *ACS Synthetic Biology* **13**(9), pp. 3046–3050 (2024), provides a concise demonstration of a few examples of how JAX can accelerate computational biology. The work highlights three primary use cases: Olivia Gallup writes about automatic differentiation and optimisation of gene circuits; Kirill Sechkar writes about stochastic cell model simulations; and Sebastian Towers writes about (the simulation of) directed evolution strategies.

Secondly, there is the work **Microscopic PhotoSelection (MiPS) of single cells in mother-machine microfluidic devices** [29] by Jessica James and Idris Kempf, with Sebastian Towers as a secondary author. The work is also available as an *eLife* Reviewed Preprint, version 1 (28 August 2026), DOI: `https://doi.org/10.7554/eLife.112391.1` [30]. The purpose of MiPS is to enable the simultaneous observation of up to approximately 10^6 cell lineages, tracking their properties over time, and to exert single-cell-level control through an optical targeting system. Coupled with an optogenetic circuit, this allows targeted activation of specific genetic pathways in individual cells, greatly expanding the range of possible perturbations and experimental designs.

The MiPS platform has broad potential applications. The rich dataset generated from hundreds of thousands of cells provides detailed insight into cellular behaviour and responses to environmental conditions. Such high-throughput data are particularly valuable for deep-learning-based analyses, where large datasets enable pattern discovery and hopefully yield mechanistic understanding. In addition, the optical targeting system allows each cell to be controlled independently based on its own history, enabling precise, adaptive interventions.

The most immediate application of these tools is directed evolution, in which cells are repeatedly reproduced and selected to evolve a trait of interest. This was the primary application explored in **Microscopic PhotoSelection (MiPS) of single cells in mother-machine microfluidic devices** [29], where UV illumination was used to suppress the reproduction of undesired cells. Established directed-evolution methods already allow researchers to vary mutagenesis conditions, population sizes and selection thresholds. MiPS extends the information available for selection by allowing decisions to depend on the observed histories and dynamic responses of hundreds of thousands of individual cells. MiPS enables selection rules of great complexity, supporting the evolution of intricate traits and the systematic exploration of the fitness landscape using information derived from the behaviour of all observed cells, both past and present.

These capabilities motivate Chapters 3 and 4: Chapter 3 investigates how to use control over selection and population allocation, while Chapter 4 investigates how population-level fitness measurements can inform the choice of strategy. MiPS provides a prospective platform for testing these methods experimentally.

The remainder of this thesis expands on each strand in detail, and I conclude the introduction with optimism that these tools will help convert the current wave of biological data and modelling power into therapies and technologies that are able to not only understand evolution, but actively steer it.

Patience is also a form of action.

Auguste Rodin

2 Background

Contents

2.1 Proteins, Folding and Binding 7
2.2 Antibodies 9
2.3 Directed Evolution 11
2.4 Fitness Landscape Theory 12
2.4.1 Ruggedness 13
2.4.2 Empirical and Model Landscapes 14
2.5 Opponent Shaping 16
2.6 JAX and Hardware Acceleration 17
2.7 Platforms for High-Throughput Single-Cell Observation 18

2.1 Proteins, Folding and Binding

Proteins are large molecules that are essential to many of the core functions of life. They act as enzymes catalysing biochemical reactions, as structural components providing cellular integrity, as transporters and signalling agents coordinating biological processes, and as regulators controlling gene expression and metabolic pathways.

The sequence in which these 20 standard[1] amino acids are chained together determines how the protein folds into its three-dimensional conformation. The

[1]There are a few more, less common, amino acids.

chemical properties of each amino acid, such as charge, polarity, and hydrophobicity, drive interactions that cause the chain to bend, twist, and fold into stable structural motifs like alpha helices and beta sheets. These secondary structures then assemble into larger tertiary and quaternary arrangements, giving rise to the protein's final shape and, consequently, its biological function.

Predicting a protein's behaviour from its amino-acid sequence is difficult because the process by which a linear chain folds into a stable three-dimensional structure involves navigating a vast and complex energy landscape [31], so even small changes in the sequence can lead to dramatic alterations in structure and function. Moreover, a protein's structure is not static: it fluctuates and adapts as it interacts with other molecules, meaning that even accurate predictions of a single folded state capture only part of its functional behaviour. Recent advances in machine learning, most notably AlphaFold [32], have substantially improved our ability to predict static structures from sequence, yet significant challenges remain in accounting for conformational flexibility, ligand interactions, and dynamic function [33].

While the three-dimensional structure of a protein is very important, what truly matters are the functional properties implied by that structure. Most important among these properties is how the protein binds to other molecules. Binding determines whether two molecules will form a stable complex, how long that complex will persist, and what biochemical consequences follow. Such complexes underpin an enormous range of biological phenomena, from enzymes catalysing reactions, to RNA inhibitors blocking gene expression, to the assembly of intricate biomolecular machines such as the bacterial flagellum.

Accurately modelling these interactions remains extremely challenging. The most detailed approaches are full atomistic molecular dynamics simulations [34, 35]. These simulations can capture binding and unbinding events as well as the accompanying conformational changes, but this level of detail comes at substantial computational cost, often necessitating months of high-performance computing resources. As a result, molecular dynamics is typically restricted to relatively short timescales or limited system sizes, although advances in methodology and

computational resources continue to expand what is feasible. Nevertheless, these simulations remain expensive to run, so many approaches instead focus on more approximate and targeted methods, such as molecular docking, empirical scoring functions, and physics-based free-energy calculations [36, 37]. These provide a less complete representation of protein–protein interactions, but estimating binding affinity alone is sufficient for many applications.

This field has also been bolstered by recent advances in deep learning. Extensions of AlphaFold to multimeric systems, such as AlphaFold-Multimer, can generate plausible joint conformations for multiple interacting proteins, with predicted complex structures that often correlate well with experimentally determined binding modes [38, 39]. However, despite these advances, binding remains an extremely complex problem, and developing faster and more accurate models will continue to be an important area of study.

2.2 Antibodies

Antibodies, or immunoglobulins, are large proteins produced by the immune system to recognise and bind to specific molecules known as antigens. Note "antigen" is not a class of molecule, but rather refers to the role of being targeted by an antibody. The typical Y-shaped antibodies have two identical binding sites located at the tips of the Y. These sites determine what the antibody can recognise and bind to, while the rest of the molecule serves as a structural scaffold and as a handle for the immune system to act upon [40].

The diversity of possible antibodies is enormous. The human body effectively performs a randomised search over a vast combinatorial space of possible binding sites, generated by rearranging and mutating genetic segments. Candidate antibodies are then tested. Those that bind strongly to foreign molecules while avoiding self-binding are retained and reproduced, whereas unsuitable ones are discarded. This biological process is conceptually similar to stochastic optimisation or evolutionary search, but operating at the molecular level [41].

Once an antibody binds its target, it can perform several functions. It might physically block the target's activity, as when an antibody prevents a virus from attaching to a host cell, or it may act as a signal, flagging the bound molecule for removal by other components of the immune system. It may also not provide a meaningful function; binding does not guarantee neutralisation. Regardless, antibodies may be thought of as precision recognition tools, linking specific binding to physiological response.

Having effective antibodies against a pathogen greatly reduces the threat it poses. Once the immune system has produced effective antibodies, subsequent exposure to the same antigen often, but not always, leads to rapid neutralisation before infection can take hold. This principle underlies immune memory and the success of vaccination, where exposure to a harmless form of a pathogen drives the immune system to develop appropriate antibodies and respond with them in future infections.

However, antibodies are not a magic silver bullet. Several criteria are required. They must be specific, binding primarily to the pathogen and not other "off-target" molecules. In the worst case, if they bind to human molecules, it can cause the immune system to attack the human body itself. They must also have high affinity, binding to the pathogen with sufficient strength. Their concentration needs to be high enough to effectively target the pathogen. And then, even with all this, they must actually help control the pathogen by either directly limiting some vital function (neutralisation) or by successfully marking the pathogen to be targeted by the immune system (opsonisation). This can be more or less difficult depending on the pathogen in question, and life has found some truly ingenious solutions. For example, the parasite *Trypanosoma brucei* can exploit the very antibodies bound to its surface: as it swims, they act as "molecular sails", with hydrodynamic drag sweeping antibody–surface-protein complexes towards its rear, where they are internalised [42].

Since antibodies are so essential to the human immune system, there is great interest in designing them *in vitro*, that is, outside a living organism. Artificially generated antibodies have a wide range of applications, from direct therapeutic

use in treating infections and cancers [43], to serving as stabilising agents in protein crystallography and structural biology [44]. Recently, advances in artificial intelligence have accelerated this field. Machine-learning models can now predict antibody binding efficacy, guide optimisation of candidate sequences, and even propose entirely new designs [45].

2.3 Directed Evolution

There are many reasons to design proteins ourselves, rather than simply relying on what nature has created. Proteins mediate the majority of fundamental processes in living systems, reflecting how they are an immensely rich and versatile class of biomolecules. Being able to create or modify them opens possibilities for new enzymes, therapeutics, and materials. One approach is rational, or *de novo*, design, where computation and biochemical reasoning are used to engineer proteins from scratch or adapt existing ones for new purposes. However, rational design remains constrained by our limited understanding of the complex relationships between sequence, structure, and function.

Directed evolution offers a complementary strategy. Rather than reasoning about design directly, it mimics natural evolution by generating large libraries of protein variants, experimentally testing them, and selecting the best performers for the next round. Through repeated cycles of mutation and selection, proteins can be "evolved" towards improved or entirely new functions without requiring detailed mechanistic insight [46]. This approach leverages high-throughput screening, automation, and selection systems that can efficiently handle thousands or even millions of variants in parallel.

Directed evolution has been applied to a wide range of challenges. It has been used to evolve antibodies with enhanced affinity and specificity [47], to optimise enzymes that convert waste cooking oil into biodiesel [48], and even to create enzymes capable of catalysing chemical reactions not found in nature, such as forming carbon–silicon bonds [49].

Rational design and directed evolution are not mutually exclusive. In recent years, computational methods have begun to merge with directed evolution, combining the strengths of both empirical and rational design. Machine learning models can now predict which mutations are most likely to improve performance, allowing only variants with the greatest chance of improvement to be tested [50]. In some cases, the evolutionary loop is even performed entirely *in silico*, using directed evolution to effectively explore the space of model outputs [51]. Together, these hybrid approaches represent a new generation of design tools that blend experimental selection with computational reasoning, accelerating discovery while reducing experimental cost.

2.4 Fitness Landscape Theory

Before we discuss fitness landscapes, we must first discuss "fitness". The term, originally from evolutionary biology, generally refers to an organism's ability to survive, thrive, and, most importantly, multiply. Already, this is a somewhat imprecise term, as an organism's ability to pass on its genes depends on its environment, competition from other organisms, and sheer dumb luck. There is also a question of time horizon. For example, asexual reproduction can offer an immediate advantage, while the genetic variation generated by sexual reproduction can support faster adaptation over subsequent generations. Experiments in rotifers have demonstrated this contrast between the immediate advantage of asexual offspring and the longer-term benefits of sex during adaptation to a new environment [52].

This term is then inherited by evolutionary algorithms and directed evolution, where fitness broadly refers to the trait of interest being selected. This is, ironically, a kind of reversal of the original meaning: in evolutionary biology fitness is the ability to reproduce, but in directed evolution fitness is instead the particular property the researcher happens to want to be reproduced. This thesis uses aspects of both definitions. For Chapters 3 and 4, fitness refers primarily to some trait of interest we deem desirable and is, more or less, assumed to be given a priori. In Chapter 5, fitness is more understood in the evolutionary biology sense, with simulated viral fitness being the capacity of the virus to pass on its genes in the current environment.

In both cases, fitness can be understood as a kind of mapping from the genotype to a fitness score, a **fitness landscape**. The idea of visualising evolution as taking place on a high-dimensional fitness landscape was first put forward by Wright [53]. Originally called an adaptive landscape, Wright proposed that evolution can be thought of as a process attempting to climb this surface, with higher points corresponding to greater reproductive success. These notions naturally extend to directed evolution, where fitness no longer refers to reproductive success but rather to how well a protein performs its intended function.

Formally, we can denote a fitness landscape as a mapping

$$f : \mathcal{S} \to \mathbb{R},$$

where $\mathcal{S}$ is the genetic sequence space. $\mathcal{S}$ will typically take the form of a discrete, high-dimensional space, such as $\{0, 1\}^N$ for N **loci** (locations along a gene), where each locus can be either unmutated (0) or mutated (1). In the case of proteins, since there are 20 standard amino acids, a more appropriate representation is $\{1, 2, \ldots, 20\}^N$. In general, the possible variants at a given locus are called **alleles**, so a protein can be said to have 20 alleles per position.

2.4.1 Ruggedness

Ruggedness, in the context of fitness landscapes, is the notion that certain fitness landscapes are more difficult to evolutionarily traverse, in part due to the presence of more local optima (genotypes with fitness higher than all their neighbours) despite not necessarily being universally optimal. Suppose we have some population evolving via natural selection, in some static environment. Since evolution is generally restricted to exploring nearby genotypes, which must be of higher (or, at least, not significantly lower) fitness for the population to transition there, if all nearby genotypes are less fit than the current one, then the population may remain clustered at the current genotype, even if it is far less fit than the global maxima.

However, this is a simplified view of ruggedness. In reality, there are many reasons why a landscape may be hard to traverse beyond the mere presence of local

optima. Thus, while there is an intuitive sense of what ruggedness means, there is no single, universally accepted definition. Several mathematical measures of ruggedness have been proposed, each capturing a slightly different aspect of the concept [54–56].

To give some examples, in Chapter 4 we use the average Hamming distance to the nearest local maximum as one measure of ruggedness, reflecting how numerous and spread out the peaks are. Reaching a nearby local maximum is not necessarily desirable: evolution can become trapped at a suboptimal peak instead of reaching the more desirable global maximum, as discussed by Poelwijk et al. [54]. A more direct measure of evolutionary accessibility is provided by Weinreich, Watson, and Chao [55], who count the number of direct, monotonic paths leading to the global maximum from the opposite end of the genetic landscape (its antipole). Others avoid considering paths or optima altogether, for example, Aita, Iwakura, and Husimi [56] quantify ruggedness via the **roughness-to-slope ratio**, which measures how close a landscape is to being perfectly linear. Despite their conceptual differences, these measures tend to be correlated in practice [57], suggesting that they capture consistent facets of what it means for a landscape to be rugged.

A closely related concept to ruggedness is epistasis. Epistasis refers to the fact that the effect of a genetic mutation often depends on the presence or absence of other mutations. This means that for a given gene, a deleterious mutation may become beneficial when combined with another, potentially deleterious, mutation. Epistasis can be viewed as ruggedness on a local scale, with rugged landscapes typically exhibiting a high degree of epistatic interaction [54, 55].

2.4.2 Empirical and Model Landscapes

In order to study fitness landscapes and their properties, it is useful to have concrete examples. Broadly speaking, these fall into two categories: **empirical landscapes** and **model landscapes**. Empirical landscapes are explicit mappings from sequence to function for real-world proteins, obtained using high-throughput experimental methods to capture the immense size of even relatively small sequence spaces. Model landscapes, by contrast, are mathematical constructions from which landscapes

can be sampled. Because they are defined mathematically, their properties can be analysed theoretically, for example, to determine the expected number of local optima. This comes at the cost of realism, but given the difficulty of constructing empirical landscapes, the trade-off is often necessary.

Empirical landscapes may either be sparse or **combinatorially complete**. Combinatorially complete landscapes represent all possible variations for some section of the protein, whereas sparse landscapes include only a fraction of these variants, typically clustered around a single important genotype. Because the number of possible variants grows exponentially with the number of loci varied, combinatorially complete landscapes are restricted to considering only a few loci, in practice up to 4, leading to $20^4 = 160{,}000$ variants. Despite this limitation, having a full landscape is important for determining certain key properties, such as if a genotype is truly a local optimum, and so we restrict our analysis to combinatorially complete landscapes.

The alternative is to use a mathematical model of fitness landscapes. Formally, these take the form of a parametrised probability distribution over possible landscapes. For example, a classic model landscape is the **House of Cards** landscape [58, 59], where every genotype has a random fitness value, statistically independent of the fitness values of all other genotypes. So long as a distribution with full support is chosen, such as the normal distribution, technically *any* landscape is possible under the House of Cards model. Despite this, certain properties become extremely likely in the limit [60], meaning this is still an informative model to study.

However, it is intuitively clear that genetically similar proteins are more likely to have at least somewhat similar fitness values, which the House of Cards model cannot represent. The **NK model** addresses this, allowing neighbouring proteins to have correlated fitness values. This is achieved by creating N groups[2]. Each group consists of $K + 1$ loci, with the i-th group guaranteed to contain locus i, and the other K loci being typically selected at random. Given the loci they contain, each group can be thought of as a smaller House of Cards sub-landscape, having

[2]These groups are often referred to as loci. For consistent terminology within this thesis, we instead opt to refer to them as groups.

a random, independent fitness value for each of the A^{K+1} possible values for the alleles at its $K+1$ loci. Then, the fitness value of the landscape as a whole is simply the sum of the N groups' fitness values. This allows correlation, as when a locus is changed, some of the groups will change their fitness values, but not necessarily all of them, and these terms that do not change lead to the correlation.

At one extreme, for $K = 0$, each group is disjoint from all others, containing only a single locus, meaning the fitness landscape as a whole is linear. At the other extreme, for $K = N - 1$, every locus is part of every group, in effect meaning there is just one group, and the model reduces to a House of Cards landscape. For intermediate values of K, the landscapes sampled somewhat interpolate between the two extremes. This means the NK model has tuneable ruggedness, something that sets it apart from many other models. There are also extensions, the most well-known being the NKp and NKq landscapes; however both of these only change the distribution of fitness values each group draws from, and so do not change the core structure of the NK model.

2.5 Opponent Shaping

In classical multi-agent reinforcement learning (MARL), agents often treat their counterparts as part of the environment, a regime known as **independent learning**. However, when two or more adaptive processes interact, each one effectively defines part of the other's fitness landscape. This leads to non-stationarity: the loss surface for one model shifts every time its opponent updates. As a result, traditional optimisation techniques often fail to converge, becoming trapped in suboptimal cycles.

Opponent shaping encompasses a class of methods that explicitly account for this coupling by anticipating an opponent's learning step and selecting actions, or parameter updates, that steer those updates toward favourable regions of the joint landscape. Rather than optimising for immediate reward under a static environment assumption, opponent-shaping agents optimise over a composite objective that includes the downstream parameter changes they can induce in others [61].

One canonical framework is **Learning with Opponent-Learning Awareness** (LOLA). In LOLA, an agent i optimises its policy by considering the anticipated update of agent j. Mathematically, this involves differentiating through the opponent's learning step, requiring the calculation of second-order derivatives to account for how agent i's current parameters θ_i affect the gradient of agent j. Follow-up work, such as **Stable Opponent Shaping** (SOS), addresses the inherent instabilities of these differentiable games, particularly the "look-ahead" dynamics that occur when multiple players simultaneously attempt to manipulate one another's gradients [62]. These methods demonstrate that reasoning about the learning dynamics of peers can foster sophisticated strategic behaviours, such as the emergence of reciprocity and the prevention of exploitative cycles in social dilemmas.

While these concepts originated in the context of neural network optimisation, they provide a powerful conceptual lens for laboratory co-evolution and synthetic biology. Many directed-evolution protocols operate in inherently adversarial or symbiotic regimes: for instance, antibodies evolving against a mutating antigen or phage-assisted continuous evolution (PACE). In these biological systems, the "opponent" (the virus, phage, or host cell) is a learning agent with its own adaptive trajectory. Framing these systems as multi-agent games emphasises that a researcher or a co-evolving population can be viewed as an agent, acting as an external "shaper", that deliberately perturbs the environment to guide the evolutionary path of the target population toward a desired fitness peak.

2.6 JAX and Hardware Acceleration

Modern machine learning relies heavily on hardware accelerators such as Graphics Processing Units (GPUs) and Tensor Processing Units (TPUs). These devices enable massively parallel computation, allowing for efficient execution of large-scale numerical workloads that would be prohibitively slow on conventional CPUs. To make use of such hardware, a variety of software libraries have been developed, including TensorFlow [63], PyTorch [64], and JAX [12].

Among these, JAX has gained particular traction for scientific and simulation-based applications. Designed to closely mirror the syntax and semantics of NumPy, JAX allows users to write high-level numerical code that can be easily automatically differentiated, vectorised, and just-in-time compiled, to be executed efficiently on accelerators through the XLA (Accelerated Linear Algebra) compiler.

This flexibility has made JAX an increasingly popular choice outside simply implementing deep learning architectures. For instance, by implementing environment dynamics directly in JAX, reinforcement learning pipelines can perform simulation, policy updates, and gradient computation entirely on the GPU, avoiding costly data transfers between CPU and device memory. This approach, exemplified by frameworks such as Gymnax [65], yields up to $\sim 10^3$-fold speedups for large-scale and multi-agent experiments.

These advantages extend beyond machine learning into other scientific applications, such as cosmological simulation [14] or molecular dynamics [13]. This also applies to biology, enabling faster stochastic simulators for cell dynamics, and the simulation/optimisation of genetic circuits [28]. For directed evolution, GPU-based simulation frameworks such as JAX open up new possibilities for parallel fitness evaluation and large-scale evolutionary search. Because JAX enables vectorised operations over entire populations, it becomes feasible to evaluate thousands of sequences or variants in parallel. This capacity for massively parallel in silico evolution allows simulations to approach the large scale of real-world biological systems, giving us a greater understanding of evolutionary dynamics.

2.7 Platforms for High-Throughput Single-Cell Observation

For decades, the standard for analysing single cells at scale has been **Fluorescence-Activated Cell Sorting** (FACS). This technique passes cells through a laser beam to measure their physical and chemical properties, sorting them into distinct populations based on fluorescence [66]. While FACS classifies populations efficiently,

it only provides a "snapshot" of a cell's state. It cannot track how a specific cell changes over time, nor can it map the lineage of a cell to see how traits are inherited.

To capture these dynamic histories, the field developed microfluidic devices like the **Mother Machine** [67]. In this system, cells are trapped in dead-end channels while nutrients flow over them. The "mother" cell remains at the end of the channel while its progeny are washed away, allowing for the continuous observation of a single lineage over hundreds of generations. This platform allowed researchers to investigate long-standing questions about cellular ageing. For instance, it demonstrated that *E. coli* maintains a remarkably constant growth rate over time, challenging previous assumptions that cells inevitably deteriorate with replicative age.

While observation is incredibly useful, being able to manipulate the cells enables a large array of possibilities by having interventions tailored to each cell's observed behaviour in a closed-loop control system. This motivated the integration of optogenetics, where light is used to activate specific genes in real time [68]. While early implementations illuminated the entire chip, more advanced setups have incorporated Digital Micromirror Devices (DMDs). These devices allow researchers to project precise light patterns onto the bacteria, enabling the independent computer-interfaced control of hundreds of individual cells within the same population [69].

As these platforms scaled, they began generating terabytes of data, outpacing human analysis. This necessitated the use of deep learning tools like DeLTA [70], which utilise U-Net neural networks to automatically segment and track cells. This combination of microfluidic containment, DMD-based optical control, and deep learning forms the technological basis for the **MiPS** platform introduced in **Microscopic PhotoSelection (MiPS) of single cells in mother-machine microfluidic devices** [29].

The MiPS platform builds upon these foundations to scale single-cell observation and control to approximately 10^6 concurrent lineages. Crucially, it expands the capabilities of the optical system beyond simple activation. By employing targeted UV illumination, MiPS can selectively suppress the growth of specific cells based

on their historical performance. This transforms the microscope from a passive tool into an active engine for directed evolution, enabling the systematic exploration of fitness landscapes and the evolution of complex traits driven by real-time data. The setup is shown in Figure 2.1.

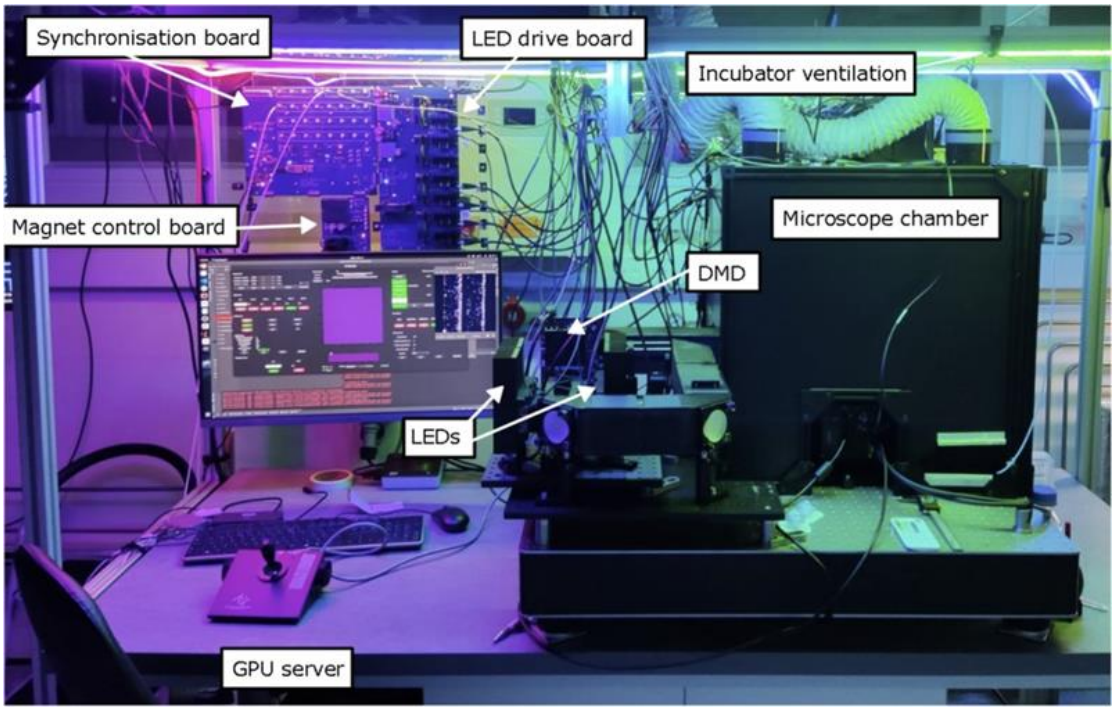


Figure 2.1: Photograph of the MiPS platform in the basement of the Thom Building, University of Oxford, used by the Engineered Biotechnology Research Group. The platform integrates high-throughput microfluidics, DMD-based optical targeting, deep-learning-based image segmentation, and stylish LED lighting; see James et al. [29] for details of the platform. Photograph: Idris Kempf.

Whatever it is you're seeking won't come in the form you're expecting.

Haruki Murakami *Kafka on the Shore*

3 Sequence-Free Directed Evolution

Contents

3.1 Motivation 21
3.2 Summary 22
3.3 Contributions 24
3.4 Paper 24
3.5 Limitations and Future Work 47
3.5.1 Limitations of the NK Model 47
3.5.2 Empirical Validation and Scale 48
3.5.3 Strategy Selection and Ruggedness Inference 48

This chapter presents *Optimisation strategies for directed evolution without sequencing*, published in *PLOS Computational Biology* (2024) [71].

3.1 Motivation

As discussed in Section 2.3, directed evolution represents one of the most powerful empirical approaches for protein engineering and the exploration of fitness landscapes. By iteratively mutating and selecting variants, it provides a general framework for optimisation that has achieved great success.

One of the major advantages of directed evolution as an optimisation method is that it does not rely on sequencing, the process of determining the amino acid

sequence that gives rise to a given protein. Whilst large improvements are being made, sequencing remains relatively expensive and time-consuming, so ideally an experimental setup will limit the amount of sequencing required. Similarly, by using physical mutation mechanisms, it is possible to randomly alter the sequence without knowing the exact change made, or the original sequence. Together, these properties allow for extremely high-throughput experimentation, albeit at the cost of precise control over which variants are explored.

These restrictions lead to a fundamentally new kind of optimisation problem. Unlike traditional optimisation, the experimenter does not observe the underlying parameters (i.e. the sequence) and cannot design explicit perturbations. The search therefore becomes a stochastic process over an unobserved, high-dimensional space, where control is limited to the strength and form of selection, the rate of mutation, and the management of population structure. Understanding how to use these to effectively optimise (i.e. evolve) a protein under these constraints is both practically important and theoretically non-trivial.

3.2 Summary

To investigate this problem, we constructed a computational model of directed evolution implemented in JAX, enabling large-scale simulation of mutation–selection dynamics on both synthetic (model-generated) and empirical (experimentally measured) fitness landscapes. The model represents genetic sequences as integer arrays, with fitness landscapes obtained either from an explicit empirical dataset or sampled from the NK model. This approach allows thousands of independent directed-evolution runs to be executed efficiently on GPU hardware, letting us find which strategies work best, and how this depends on the underlying landscape.

First, we examined the effects of non-greedy selection strategies. In conventional directed evolution, only the top-performing fraction of variants are propagated each round, which strongly exploits local improvements but risks getting stuck in local optima. To remedy this, we defined a family of **selection functions**, parametrised by **base chance**. While the highest-fitness variants always get

selected, all other variants still have a base chance probability of being selected. Simulations showed that this markedly improves performance on rugged landscapes, probably by promoting exploration through low-fitness "valleys", particularly at low mutation rates where many consecutive mutations are required to cross these valleys.

We also investigated the role of population structure. Given a fixed experimental budget (i.e. total number of fitness evaluations), we show that it is often more efficient to run multiple smaller, independent populations rather than a single large one. This **population splitting** approach increases trajectory diversity, effectively exploring multiple regions of the landscape in parallel. Across a range of conditions, population splitting improved both average performance and reproducibility, reducing variance in outcomes between replicate runs, and boosting the average performance.

However, these results are only on the model NK landscape, which is in many ways not representative of real landscapes. To remedy this, we validated these strategies on empirical fitness landscapes derived from the GB1 immunoglobulin-binding domain and TrpB tryptophan synthase. Incorporating both non-greedy selection and population splitting led to large gains, a 19-fold increase in the probability of reaching the global optimum on GB1 and a 7-fold increase on TrpB, relative to the standard greedy method. These results highlight how simple, experimentally implementable modifications to selection strategy can substantially accelerate discovery of high-performing variants, even in the absence of sequence information.

Overall, these findings demonstrate that directed evolution can be treated as a tunable optimisation process even without explicit sequence data. By appropriately balancing exploration and exploitation, and by managing population structure, it is possible to develop superior variants in fewer iterations. These principles provide a theoretical and computational foundation for future high-throughput sequencing-free directed evolution experiments, particularly as new single-cell and continuous-evolution technologies mature.

3.3 Contributions

This work represents a joint collaboration between Jessica James and Sebastian Towers, and so it is important to clearly state the contributions of each primary author. The first iteration of the directed evolution framework was made by J. James in MATLAB, allowing for initial experiments in how non-greedy selection methods can improve outcomes. This framework was then rewritten by S. Towers to utilise JAX, alongside optimising the NK model to allow on-the-fly computation. These changes allowed much larger populations to be tested on far larger landscapes, allowing for more detailed computational experiments. Using this computational framework, experimental design and evaluation represent a joint effort between J. James and S. Towers, with Harrison Steel providing valuable supervision, and Jakob Foerster giving additional insights throughout. More specifically, J. James produced Figures 3 (population splitting) and 5 (empirical landscapes), S. Towers produced Figure 4 (multi-factor selection), and both worked on Figure 2 (selection strategies).

3.4 Paper

PLOS COMPUTATIONAL BIOLOGY

RESEARCH ARTICLE

Optimisation strategies for directed evolution without sequencing

Jessica James[☯], **Sebastian Towers**[☯], **Jakob Foerster, Harrison Steel***

Department of Engineering Science, University of Oxford, Oxford, United Kingdom

☯ These authors contributed equally to this work.
* harrison.steel@eng.ox.ac.uk

Check for updates

OPEN ACCESS

Citation: James J, Towers S, Foerster J, Steel H (2024) Optimisation strategies for directed evolution without sequencing. PLoS Comput Biol 20(12): e1012695. https://doi.org/10.1371/journal.pcbi.1012695

Editor: Alexandre V. Morozov, Rutgers University, UNITED STATES OF AMERICA

Received: March 15, 2024

Accepted: December 4, 2024

Published: December 19, 2024

Peer Review History: PLOS recognizes the benefits of transparency in the peer review process; therefore, we enable the publication of all of the content of peer review and author responses alongside final, published articles. The editorial history of this article is available here: https://doi.org/10.1371/journal.pcbi.1012695

Copyright: © 2024 James et al. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Data and code used for running experiments and plotting is available on Software Heritage Archive at https://archive.softwareheritage.org/swh:1:dir:d61e1c81185705659c6ed545a77e7c84cd4aa6d5,

Abstract

Directed evolution can enable engineering of biological systems with minimal knowledge of their underlying sequence-to-function relationships. A typical directed evolution process consists of iterative rounds of mutagenesis and selection that are designed to steer changes in a biological system (e.g. a protein) towards some functional goal. Much work has been done, particularly leveraging advancements in machine learning, to optimise the process of directed evolution. Many of these methods, however, require DNA sequencing and synthesis, making them resource-intensive and incompatible with developments in targeted *in vivo* mutagenesis. Operating within the experimental constraints of established sorting-based directed evolution techniques (e.g. Fluorescence-Activated Cell Sorting, FACS), we explore approaches for optimisation of directed evolution that could in future be implemented without sequencing information. We then expand our methods to the context of emerging experimental techniques in directed evolution, which allow for single-cell selection based on fitness objectives defined from any combination of measurable traits. Finally, we explore these alternative strategies on the GB1 and TrpB empirical landscapes, demonstrating that they could lead to up to 19-fold and 7-fold increases respectively in the probability of attaining the global fitness peak.

Author summary

The standard approach to sorting-based selection in directed evolution is to take forward only the top-performing variants from each generation of a single population. There are, however, many possible approaches to exploring non-convex evolutionary fitness landscapes, and choosing this strategy as default may not always be the strongest approach. In this work, we begin to explore alternative selection strategies within a simulated directed evolution framework. We propose "selection functions", which allow us to tune the balance of exploration and exploitation of a fitness landscape, and we demonstrate that splitting a population into sub-populations can improve both the likelihood and magnitude of a successful outcome. We also propose strategies to leverage emerging selection methods that can implement single-cell selection based on any combination of measurable traits. We finally assess the space of alternative directed evolution strategies on the empirical

as well as GitHub at https://github.com/nesou2/direvo_sim.git. Raw data available at https://doi.org/10.5281/zenodo.12799997.

Funding: H.S. is supported in part by the Engineering and Physical Sciences Research Council (EPSRC, https://www.ukri.org/councils/epsrc/) projects EP/W000326/1 and EP/X017982/1. The funders did not play any role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

fitness landscapes of the GB1 immunoglobulin protein and of TrpB tryptophan synthase. Our resulting proposal is that researchers should consider moving away from the standard approach, which we find to be generally sub-optimal, and implement population splitting to improve experiments. With improved knowledge from fitness landscape inference, directed evolution strategies could be further tailored using the tools proposed here.

Introduction

Engineered biological systems hold immense potential for application across industries including medicine, manufacturing, and agriculture [1–3]. In recent decades, protein engineering in particular has demonstrated the potential of natural biological elements to be adapted for new functionalities. Advancements in computational methods are bringing *de novo* protein engineering closer to reality [4–8], however such approaches remain limited by our developing understanding of protein sequence-to-function relationships. To circumvent the need for such detailed prior knowledge, a range of techniques termed "directed evolution" have been developed [9]. Directed evolution techniques have delivered products across a range of applications, from cancer and autoimmune disorder drugs [10] to enzymes for converting cooking oil into bio diesel [11].

In a process that mimics nature, directed evolution consists of iteratively introducing random variation by mutagenesis, followed by selection biased toward desirable user-defined traits. Directed evolution methods are highly varied, and exist for both *in vivo* and *in vitro* settings. In this work, we have modelled *in vivo* directed evolution, in which fitness is evaluated inside a host organism such as a bacterial cell. Selection of living cells can be achieved broadly via two families of approaches; first, those that couple a trait-of-interest to growth of a host organism [12, 13], or second, those that couple a trait to a measurable output (e.g. expression of fluorescent proteins). In the latter, cells are screened to identify and isolate those with desirable output expression levels for future mutagenesis and propagation [14]. Each approach has different strengths and weaknesses. Growth-coupled selection utilises (comparatively) straightforward growth-based assays but requires engineering of trait-to-growth coupling, which is both challenging and can lead to "cheating" behaviour [15]. Meanwhile, screening-based methods only require a trait be *measurable* (i.e. they do not require coupling to growth), but in turn necessitate more complex experimental approaches to implement the selection for the measured winners (e.g. FACS [14]). In contrast to FACS, which takes only a single time-point measurement from each cell, emerging selection techniques leverage microfluidics to observe cells over long time periods prior to sorting [16, 17]. This produces large amounts of additional temporal information on which to sort cells, and allows selection to be performed based on dynamic phenotypes. Our work focuses on these emerging microfluidics-based methods due to their increased level of control, and consequently not all strategies we propose are suitable for growth-coupled directed evolution. This is a timely challenge, as new methods for single-cell selection pose novel theoretical questions—which our work aims to explore—regarding how their capabilities can be optimally implemented and exploited.

Mutagenesis methods are another essential component of directed evolution that accelerate the speed of evolution by introducing mutations at a greater than natural rate. *In vivo* mutagenesis methods introduce mutations to DNA inside living cells, which can then be cycled directly back into the selection process, resulting in a continuous and therefore less labour-intensive form of directed evolution. For example, methods including treatment with UV or chemical mutagens can lead to genome-wide increases in mutation rates, however, such

untargeted mutations may not always impact the trait of interest, and may cause unwanted effects elsewhere. An emerging family of more sophisticated targeted *in vivo* mutagenesis methods allows this limitation to be overcome in cases where specific sequence(s) are hypothesised to hold potential for improving a desired trait, by targeting mutagesis primarily within that region. Such methods, including EvolvR [18], MutaT7 [19], and OrthoRep [20], have already begun to be adopted by the wider community, and are continually undergoing further refinement for increased performance [21].

As a biological entity undergoes directed evolution, the process can be imagined as navigation across high-dimensional "fitness landscapes" [22]. Fitness landscapes map each genetic sequence to a measure of fitness (with "fitness" being performance for a desired function), and the goal of directed evolution is to find the highest peaks on that landscape. Fitness landscapes are known to exhibit variable degrees of ruggedness, which can create local optima that constrain paths of evolution [23–25]. Standard practice in directed evolution is to take forward and mutate only the top fraction of variants during each iteration [26, 28–31]. This "greedy" approach is prone to getting trapped in local optima, particularly on rugged landscapes [32]. With the aid of computational methods, however, it is possible to navigate protein fitness landscapes in a more active way. One of the earliest examples of such a method is ProSAR, which uses a statistical algorithm to identify specific residues that are correlated with high fitness. Each new generation of variants is designed to combine residues that were predicted to contribute most to fitness [33]. Methods that predict the fitness effects of mutations in this way are now able to accommodate machine learning [34–36] and Bayesian optimisation approaches [37, 38]. Such methods have been built upon by not only utilising the fitnesses of the sequences in isolation, but also time-series mutation data acquired *during* a directed evolution experiment [39]. The vast majority of these methods, however, are based upon the requirement for sequencing information from each generation. This means they are somewhat resource- and labour-intensive, and are not suited to maximising the benefits from *in vivo* mutagenesis methods for directed evolution [21].

In instances where sequencing information is unavailable, the field has established the default approach of selecting only the top variants with each generation [26, 28–31]. This approach may restrict population diversity, leading to a higher propensity to get trapped in local optima. Examples of previous work to address this include alternating between "on" and "off" states of selection [40], as well as modifying selection stringency [41] to increase population diversity. Here, we approach the challenge from several new angles. First, probability of selection is applied as a parameterised function of fitness that can be used to tune the balance exploration and exploitation on a fitness landscape. Second, we investigate the benefits that can be gained by splitting a population into sub-populations and allowing their trajectories to diverge. Finally, we explore the novel capabilities of the aforementioned emerging selection methods [16, 17], which enable effective optimisation of multiple properties in parallel. We assess the performance of our optimisation approaches by simulating directed evolution on the GB1 and TrpB empirical landscapes [23, 25].

Results

In order to test selection strategies, a computational model was implemented to mimic the process of directed evolution. Genes in the model are represented by one-dimensional arrays, which iterate through rounds of mutation and selection (Fig 1, Methods: Model). In the selection process, the fitness of each gene is calculated using an empirical landscape [23, 25] or an NK model [42, 43]; see Methods. Empirical landscapes are combinatorially complete fitness measurements for all variants of a protein (or protein region). NK models are computationally

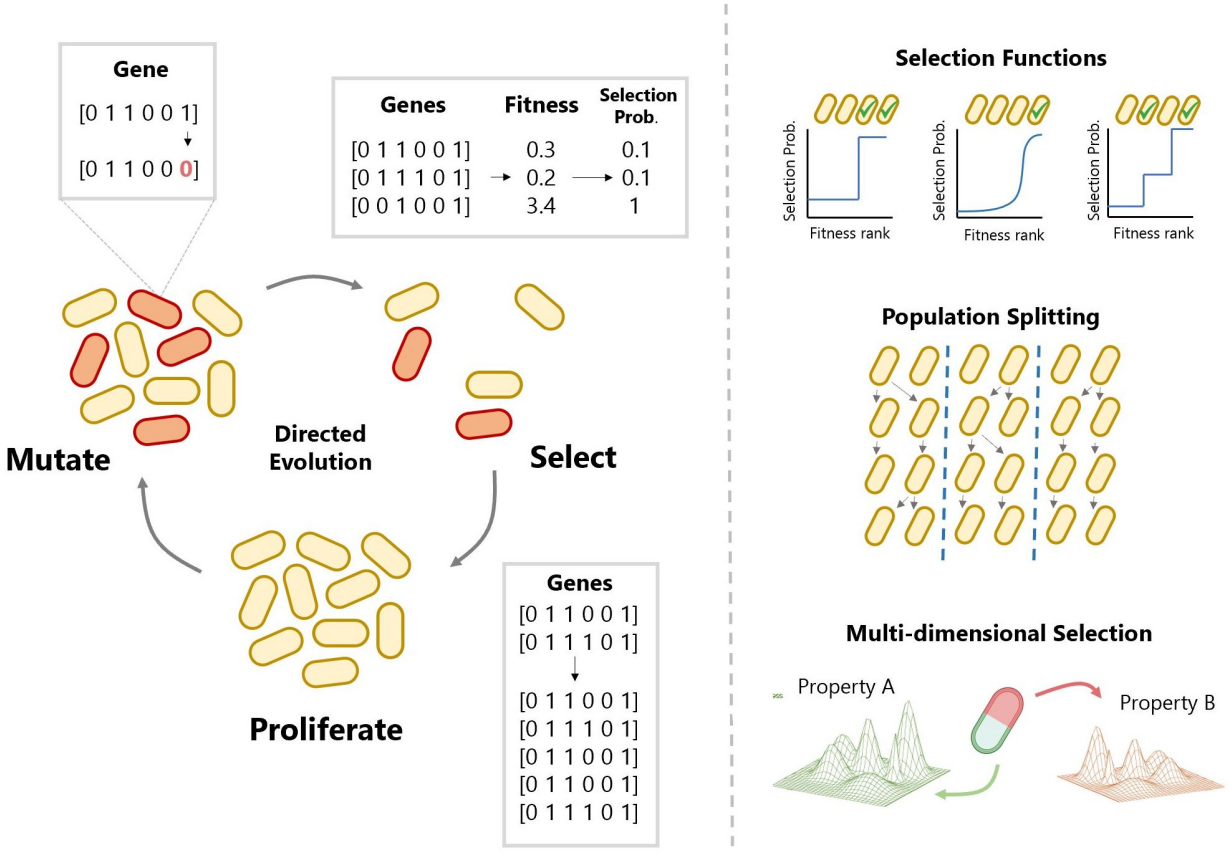


Fig 1. Schematic of the directed evolution simulation cycle. The model of directed evolution performs iterative round of mutation, selection and proliferation. Genes are represented by one-dimensional arrays. The fitness of each gene can be generated by feeding the array into a fitness landscape model. The probability of selection is determined by feeding the resulting fitness into a selection function. Proliferation is carried out by sampling with even probability up to a fixed population size. Mutation is carried out by introducing random changes to the arrays. For a more detailed description of the computational pipeline, see Methods: Model. Strategies explored using the model include selection functions, population splitting and selection across multiple properties.

https://doi.org/10.1371/journal.pcbi.1012695.g001

generated fitness landscapes that are made necessary by the limited availability of empirical landscapes. In both NK and empirical landscape implementations, the gene sequence is taken as input and the corresponding fitness value is given as output. As outlined above, the approaches explored with the model are possible in contexts where one actively selects "winning" variants to enrich (via FACS or microfluidic sorting), as opposed to growth-coupled directed evolution, which does not offer this type of control.

Selection functions for tuneable exploration vs exploitation

Selection functions are introduced as a means to tune the balance of exploration and exploitation on a fitness landscape. The selection functions proposed here are defined by two parameters: "fitness threshold" and "base chance" (Fig 2A). The fitness threshold is the fitness percentile above which variants have a 100% chance of selection, otherwise the chance of selection is equal to the base chance (Methods: Selection Functions). In this work, selection functions are normalised to select a constant fraction of variants. In a continuous directed evolution experiment, this ensures that proliferation time remains approximately constant

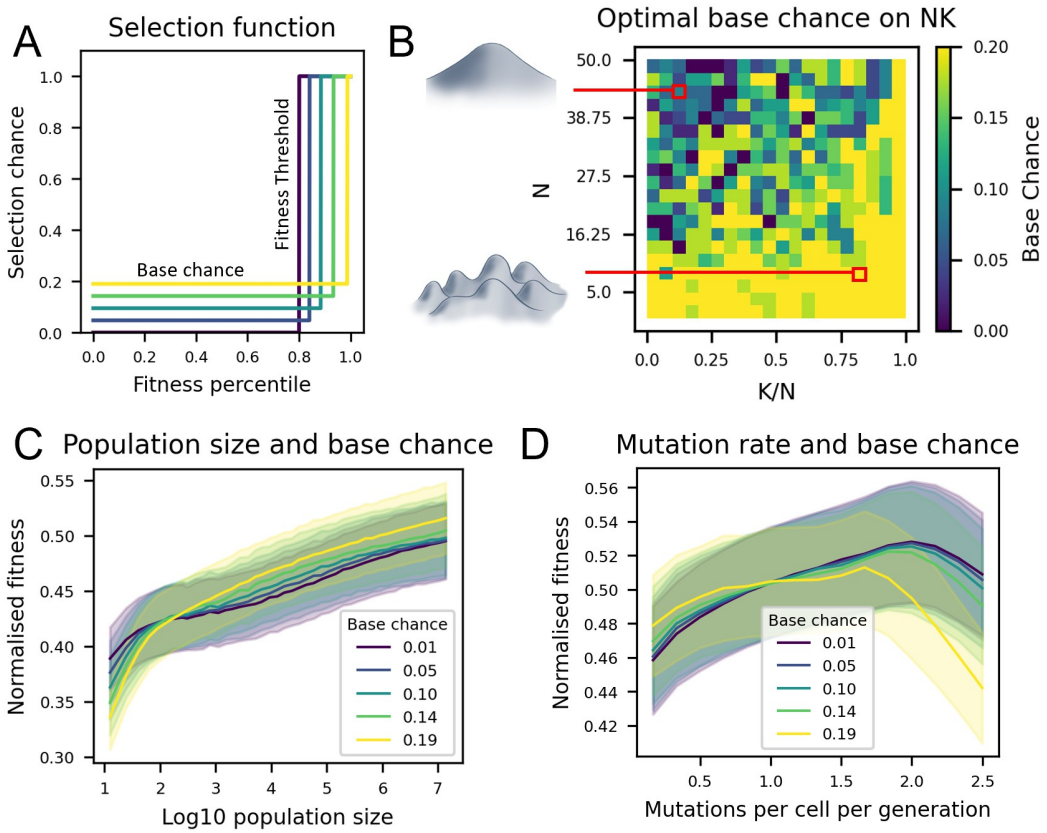


Fig 2. Investigating the performance of selection functions in directed evolution. A: Selection functions define probability of selection as a function of fitness (Methods: Selection Functions). The selection function used here is determined by two parameters: fitness threshold and base chance. Selection functions were normalised to select 20% of the population. B: Optimal base chance values on varying NK landscapes. Dependence of trajectory end-point fitness on C: population size, and D: mutation rate. "Normalised fitness" is maximum fitness across the population at the final time point (averaged over 100 runs) as a fraction of the global maximum on the landscape. Shaded areas represent standard deviation over repeat runs. Experiments ran for 300 generations. $0.01 \leq$ base chance ≤ 0.19, N = 25, K = 5, mutations per cell = 0.1, population size = 1000.

https://doi.org/10.1371/journal.pcbi.1012695.g002

between generations (hence, performance metrics can be considered as improvement in a trait *per unit time*). Fixing the selected proportion of cells also reduces the parameter space of selection functions to one dimension, as every base chance has only one fitness threshold value corresponding to a fixed proportion of the population being selected. We hypothesise that the base chance parameter will improve directed evolution by allowing a population to escape local optima on the fitness landscape. By selecting some cells unconditionally, they are allowed to accumulate more mutations, potentially allowing them reach higher performing variants via deleterious phenotypes.

The NK model describes a class of fitness landscape with tuneable ruggedness [42, 43] (Methods: NK Landscapes). N describes the number of variable sites, and K can be thought of as a metric of ruggedness ranging from 0 to $N - 1$, with high K meaning high ruggedness (i.e. more local optima). Fig 2B shows the optimal base chance over varying NK landscapes, as estimated via simulation. In particular, Fig 2B shows the optimal base chance increasing with respect to K (ruggedness), and decreasing with respect to N (dimensionality). Given that base chance is hypothesised to help escape local optima, one explanation for this result is that more

local optima are found in rugged landscapes, and they are more difficult to escape in low-dimensional landscapes, which offer fewer paths between any two points. For smooth, high-dimensional landscapes the opposite is true, therefore the function that most favours exploitation over exploration (i.e. base chance = 0) is found to be optimal. From this point onwards, the landscape N = 25, K = 5 was used as our standard landscape, as it displayed neither overly rugged nor overly smooth characteristics (S1 Fig). The robustness of base chance to noise on this standard landscape is displayed in S2 Fig.

Next, the interaction between base chance and population size was explored. Note that the NK landscape is non-linear, therefore a 1% increase in raw fitness may truly represent a larger underlying improvement, particularly at the high-fitness end of the distribution. Fig 2C shows that base chance can improve performance across a range of population sizes, particularly in large populations. Given that large populations can explore more of the landscape, they are likely to encounter more local optima. The ability for a population to escape a local optimum is dependent on its ability to reach a fitter state via at least one deleterious mutation. Evolution via deleterious mutations is shown to occur more readily in large populations [44, 45] and this dynamic may be promoted by base chance, hence the improvement in performance. In small populations, the cost of including detrimental variants is greater relative to the potential gain, therefore base chance is less beneficial.

The performance of the selection functions with varying mutation rates was also investigated (Fig 2D). When mutation rate is low, higher base chance values perform best and vice versa for high mutation rate. One explanation for this is in the balance of exploration and exploitation. Both base chance and mutation rate aid in escaping local optima by increasing the likelihood of a cell undergoing multiple mutations. Although this benefits exploration, *further* increases in mutation rate can come at the cost of not effectively exploiting a position on the landscape. For this reason, base chance performance drops off more quickly in a high mutation rate regime. Given that most directed evolution experiments operate in a low mutation rate regime to avoid detrimental side effects [46], base chance could act as a useful tool for promoting landscape exploration. The benefits of such an approach are that implementing a base chance has no direct impact on top-performing variants, whereas increasing mutation rate impacts all variants.

Population splitting for improved exploration

Until now, this work has assumed a single population undergoing directed evolution. However, in practice, one could run multiple, smaller, copies of the same directed evolution experiment by subdividing a population, and take the best outcome across all of them as the final result. Here, this method is referred to as population splitting. An example of such a situation is displayed in Fig 3A, where a population of size 500 is split into five equal sub-populations of size 100. In this example experiment population splitting performs better, and Fig 3D and 3E demonstrate the consistency of this result across parameter regimes. This may be because in a single, mixed population, mutations will eventually drift to fixation or extinction, therefore the population as a whole remains largely on the same trajectory. If one splits the population, sub-populations are able to drift on separate trajectories without cross-competition, effectively mimicking the process of speciation and increasing landscape exploration [47, 48].

Principal components analysis (PCA) was used to verify that sub-populations diverge on separate trajectories. The final timepoint gene sequences from the simulation in Fig 3A were collected. PCA was performed on the combined dataset to reduce the 25-dimensional genetic space to just 2 dimensions for visualisation (Fig 3B). The result shows that each sub-population forms a cluster, and the overall variation of the split population is significantly more than the

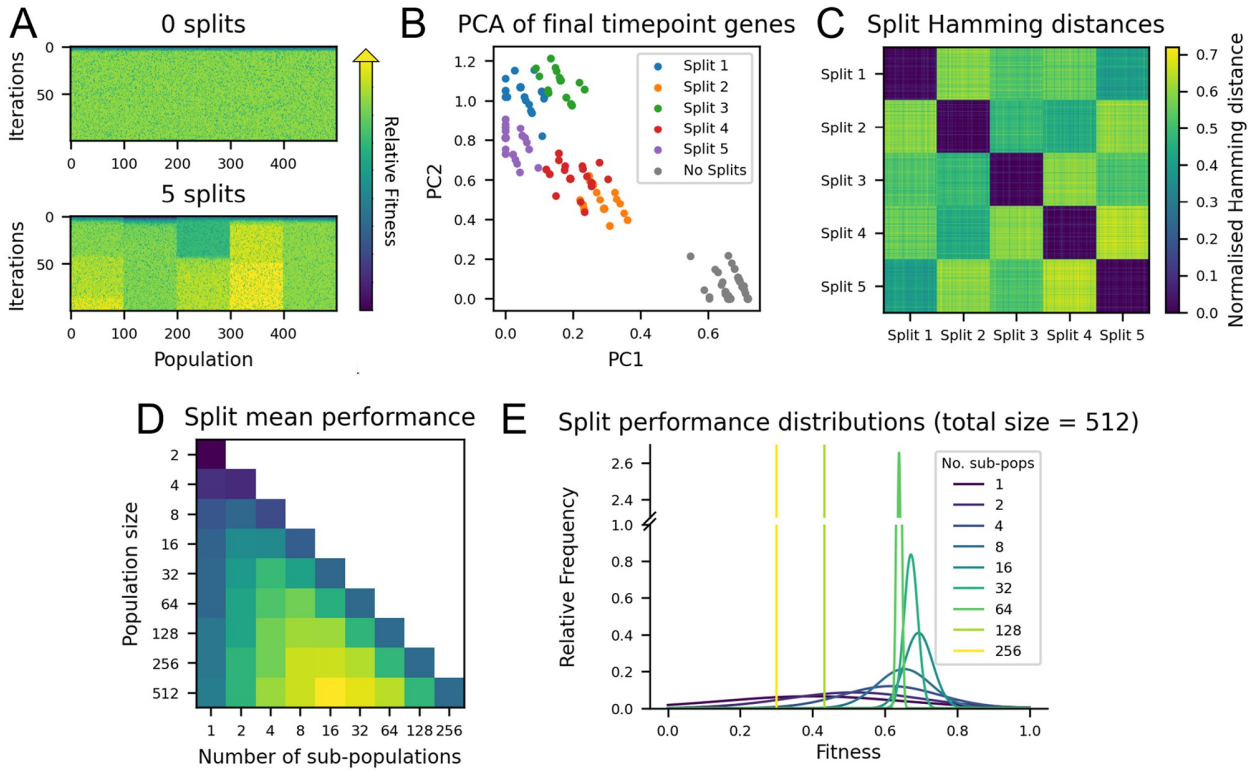


Fig 3. Investigating the effects of population splitting in directed evolution. A: Example of a directed evolution run split into five sub-populations vs a single large population. Colourbar indicates fitness value. B: Principal components analysis of the final time point sequences of a split vs non-split population. C: Hamming distances between the final timepoint sequences of the split population. D: Mean performance of directed evolution with varying total population and sub-population size. E: Mean and standard deviation of performance with varying number of sub-populations (fit to a normal distribution). Experiments ran for 100 generations. Mutations per cell = 0.1, base chance = 0, N = 25, K = 5.

https://doi.org/10.1371/journal.pcbi.1012695.g003

non-split population. This is further verified in Fig 3C, which displays that the average normalised Hamming distance between sub-populations (0.44) was far greater than that within sub-populations (0.011, similar to the average Hamming distance of 0.009 measured within the large single population).

Population splitting can clearly confer a benefit to performance, however there is a trade-off between splitting the population to maximise exploration, and keeping sub-populations large enough to effectively search around their local position on the landscape. Fig 3D summarises the performance of population splitting for different total population sizes, demonstrating that if the (total) population is too small, splitting is instead detrimental to performance. An equivalent plot on landscapes of different ruggedness is displayed in S3 Fig. Fig 3E shows the distribution of performance outcomes, over 1000 runs, for splitting a large population into increasingly smaller populations. This demonstrates an additional advantage of population splitting, which is that the variance of the final outcome decreases as the number of sub-populations increases.

Multi-dimensional selection with simulated novel selection methods

Previous sections have operated within the constraints of well-established directed evolution selection methods (e.g. FACS). Emerging methods for selection, however, may offer increased capabilities; notably using microfluidics to observe cells for long time periods prior to selection

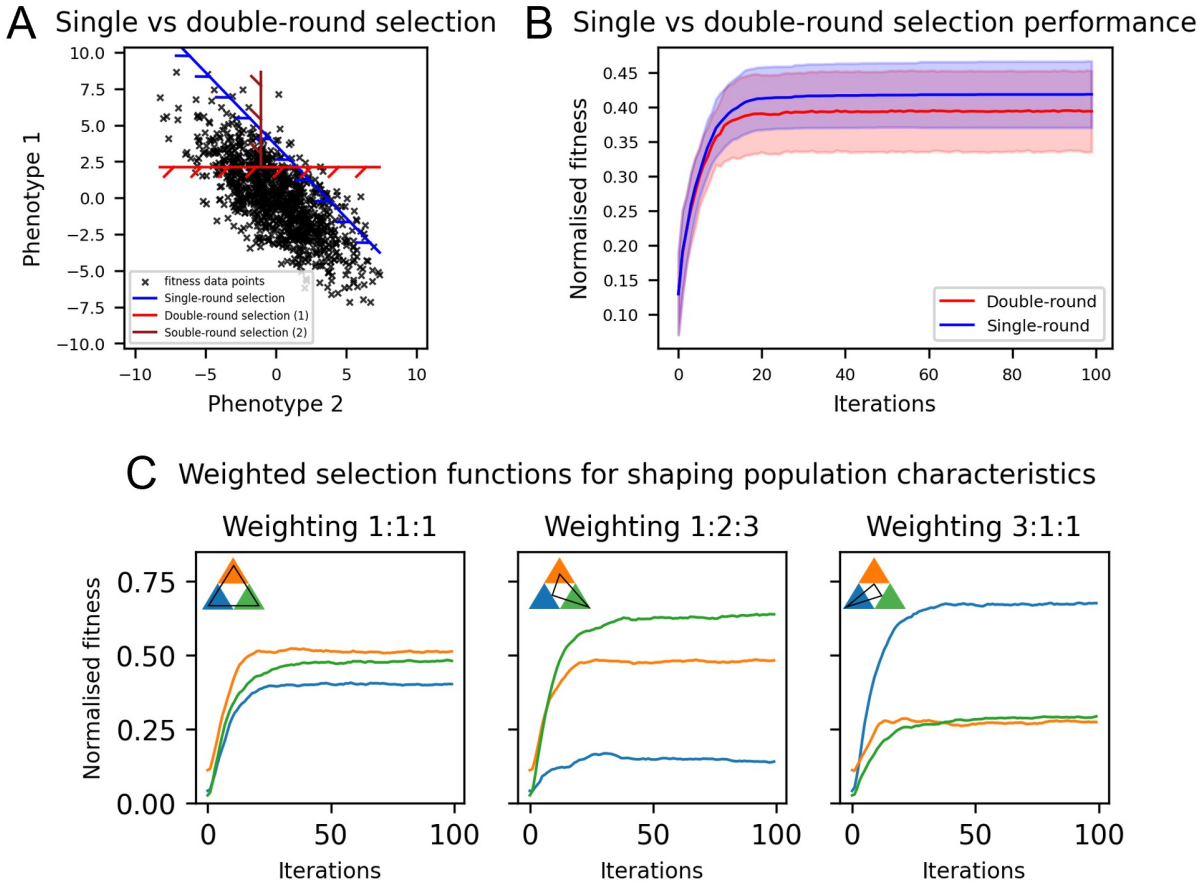


Fig 4. Optimisation of multi-dimensional selection. A: Demonstration of selection patterns with double-round (i.e. flow cytometry selection) and single-round selection. Randomly-generated fitness points normally-distributed around the mean 0 of phenotypes 1 and 2. B: Directed evolution performance of double-round vs single-round selection. Shaded region is standard deviation over repeated runs. C: Performance of weighted selection functions to perform directed evolution over three properties, with weightings of 1:1:1, 1:2:3 and 3:1:1. Experiments ran for 100 generations. N = 25, K = 5, mutations per cell = 0.1, population size = 1000, fitness threshold top 5%, as per [26].

https://doi.org/10.1371/journal.pcbi.1012695.g004

[16, 17]. Not only would long-term observation increase the reliability of readings and allow selection based on complex time-dependent traits, it would also allow for multiple properties (or responses to stimuli) to be measured in a single round of selection. Such multi-dimensional selection is highly applicable in the directed evolution of biosensors, in which one seeks to optimise both specificity and sensitivity [26].

The current standard approach to multi-dimensional selection (e.g. in FACS) is to perform sequential rounds of selection, one for each property [26, 27]. The sub-optimality of this approach is previously highlighted in [53]. This introduces a systematic error with respect to the selection objective, as shown in Fig 4A. The blue dashed line divides the true best cells from the population (as could be achieved by single-round selection), whereas the red dashed lines represent the cut-offs of a double-round selection setup. Cells that are poor in one property but excel in another are not selected by double-round selection. As a result, the overall performance of single-round selection is higher (Fig 4B).

Not only would the described emerging methods improve upon FACS in the simplest case, they would also offer the additional ability to *tune* the prioritisation of different properties. When selecting on a single property, the fitness value (F) used to determine selection is simply

the value of that property. When selecting on multiple properties, however, the overall fitness value used to determine selection is some combination. Given that most directed evolution experiments have a limited amount of time and/or physical resources, it is crucial to consider how much one prioritises each property in this combination. This prioritisation can be implemented by applying a weight (w_i) to the value of each property. So, $F = w_1f_1 + w_2f_2 + w_3f_3$, where f_i is the value of property i. In the simplest case, we allow all weightings to be equal (Fig 4C). By changing the weightings of the properties we observe proportional gains in the fitness of each property.

Translation to empirical fitness landscapes

The preceding results are based on simulations of NK landscapes, which although widely used, may not capture all important properties of real fitness landscapes. In order to test these strategies further, we therefore applied them two different empirical landscapes. Firstly the GB1 fitness landscape [23], which measures the binding strength of 149,631 variants of the GB1 immunoglobulin protein to IgG-Fc. Second, the TrpB empirical landscape [25] which measures the activity of 159,129 variants of typtophan synthase. In both GB1 and TrpB, a model was generated to impute the fitness of missing variants, bringing both combinatorially complete landscapes up to 160,000 variants (corresponding to 20 amino acid possibilities at each of the 4 sites).

We used these landscapes to assess the performance of strategies that employ base chance and/or population splitting. Fig 5A displays the performance of varying combinations of base chance and splitting. We observed a clear optimum in population splitting in the range of 50 sub-populations, which remains the optimum at most base chance values. The trends with respect to base chance are much weaker, and also dependent on the number of sub-populations. The impact of base chance is stronger when there is no population splitting.

Fig 5B displays the distribution of outcomes from GB1 directed evolution using four different strategies highlighted in Fig 5A (neither splitting nor base chance, each strategy in isolation, and both strategies combined). Of these four methods, the standard approach to directed evolution (employing neither splitting nor base chance) performs the worst, with over 60% of directed evolution runs on GB1 getting trapped at a local optimum (5.8x wildtype fitness). By introducing base chance that fraction is reduced to less than 30%, and by introducing splitting it is reduced to almost 0%. As for the fraction of runs which reach the *global optimum* (9.9x wildtype fitness), this reached 19% with the combined splitting and base chance approach compared to 1% using the standard approach, a 19-fold improvement.

Fig 5C and 5D display the equivalent measurements for the TrpB landscape. Compared to a strategy that employed neither base chance nor splitting, the optimal combined strategy generated a 7-fold increase in the probability of reaching the global optimum. Interestingly, greater benefits were achieved by base chance alone than by splitting alone, in contrast to GB1. The increased role of base chance in TrpB directed evolution could be on account of the fact that it is estimated to have 5x more local optima [25]. This may suggest that the role of base chance is more important in escaping local optima, whereas splitting strategies are more beneficial in terms of global exploration.

Despite the fact that GB1 and TrpB come from biologically very different contexts, they both responded best to a strategy that employs a combination of base chance and population splitting. Conversely, both exhibited the weakest performance with a strategy that employed neither approach (i.e. a typical "greedy" approach to directed evolution). Considering these observations, we hypothesise that the standard approach of no splitting and no base chance may *in general* significantly under-perform in real-world experiments. Deviating from the

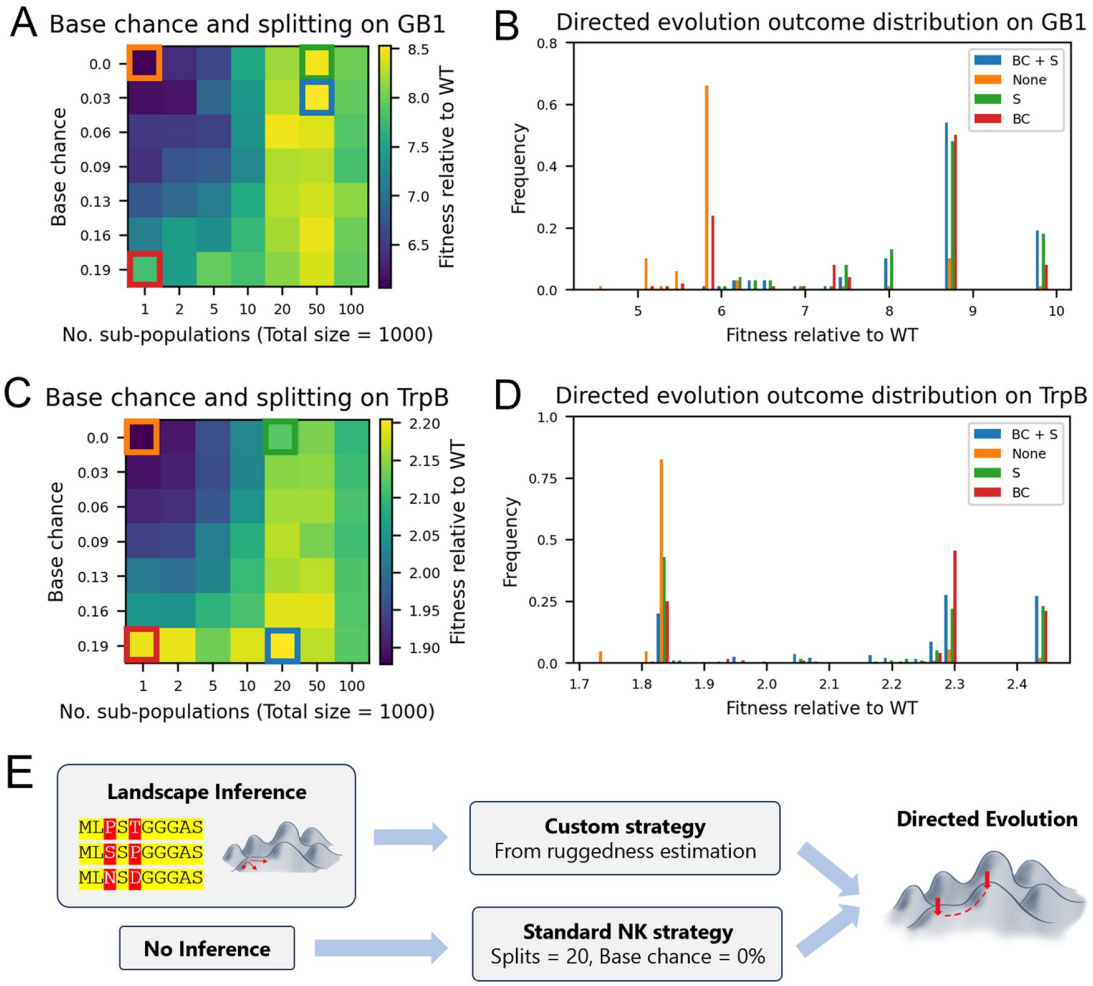


Fig 5. Application of base chance and population splitting to the empirical GB1 and TrpB landscapes. Mean performance (A) and outcome distribution (B) of varying base chance values and sub-population numbers in GB1 directed evolution. (C), (D): Same data respectively for TrpB. BC: base chance, S: splitting. Fitness score relative to WT (GB1: VDGV, TrpB: VFVS). Outcomes are from 1000 simulations with mutation rate = 0.01, total population size = 1000. E: Proposal directed evolution pipeline. Standard NK strategy comes from parameter sweep in S4 Fig.

https://doi.org/10.1371/journal.pcbi.1012695.g005

standard approach to directed evolution, in particular by population splitting and/or adding a non-zero base chance, may therefore offer benefits even if it is not the absolute optimal strategy.

In order to determine what an improved default approach might look like, a parameter sweep over many different NK landscapes, population sizes and mutation rates was conducted to identify the most robust strategy (S4 Fig). A range of landscapes from completely smooth to rugged (K/N = 0.25) was selected. In the absence of a more informed prior, we weighted each landscape equally to calculate a final mean score. This embeds an assumption about the properties of real landscapes that might be considered in such experiments, which might be refined in future work as more experimental datasets become available. Following this simulation, our results suggest that in the absence of any prior knowledge, employing splitting but no base chance (up to 20 splits, 0% base chance) may improve likelihood of a successful outcome. On average, this strategy was in the 65th percentile of all strategies, compared to the 42nd

percentile for the standard no splitting approach. This "optimal" splitting level is comparable to that observed for GB1 and TrpB. However, base chance was found to be generally disadvantageous even at low splitting (contrary to observations on GB1 and TrpB) which may be due to the chosen distribution of landscapes simulated (which includes completely smooth landscapes). The fact that base chance appears beneficial on the empirical, non-smooth landscapes of Fig 5 also supports this suggestion. Nonetheless, attempting to implement 0% base chance accounts for the possibility of a highly additive landscape, and in practise, a certain degree may arise naturally due to inaccuracies in the selection process. In summary, in the absence of further knowledge one could implement up to 20 population splits, and no base chance, which may improve the likelihood of a successful outcome.

In the future, selection strategies could be further refined with the aid of novel landscape inference methods, which can predict landscape parameters from sequence alignment [54] or experimental measurements of fitness in real-time as they are collected [39, 55]. Further work will be required on translating landscape inference to optimal strategy, which may either consider how they relate to NK landscapes (or other parameterisations of landscape ruggedness), or by conducting experiments to compare landscape inference outputs and directed evolution strategy performance. To summarise our recommendation, we propose two possibilities for the future (Fig 5E). With no prior data about the landscape available, the default NK strategy outlined above could be implemented. In order to refine the process further, landscape inference tools could be used to generate information on the ruggedness of their landscape and select an appropriate strategy.

Discussion

This study demonstrates that even in the absence of sequencing information, there are approaches that can be used to improve directed evolution outcomes. Such approaches were demonstrated both on simulated NK landscapes, and on protein empirical fitness landscapes. On the GB1 landscape, we showed alternative directed evolution strategies can lead to up to a 19-fold increase in the probability of reaching the global optimum, without any requirement for sequencing data.

There are several ways in which the methods outlined in this paper could be implemented experimentally. In its simplest form, a selection function of just "fitness threshold" and "base chance" parameters (Fig 2A) could be implemented by pipetting the appropriate portion of pre-selection cells into a post-selection population. A more sophisticated approach could incorporate the selection function into the Fluoresence Activated Cell Sorting (FACS) methodologies, or microfluidic-based methods for cell selection and sorting [16, 17]. Population splitting is the most intuitive to implement, and in many cases is already performed, in the form of biological replicates. Multi-dimensional selection in the form we describe can only be implemented with sorting technology that allows for multiple properties to measured prior to selection (i.e. with microfluidics) [16, 17].

The main limitation of this work is that optimal strategies are dependent on the shape of the fitness landscape, which is generally unknown. With the continued development of fitness landscape inference tools, however, it is becoming possible to make estimations about landscape properties. Of these there are two forms of fitness landscape inference we wish to highlight: Those that employ pre-existing sequence data, and those that are based on data as it is collected *during* a directed evolution experiment. The former utilises multiple sequence alignment (MSA) information, and identifies interacting residues from patterns of co-evolution (the frequency of interacting residues being an indicator of ruggedness). This approach is employed by EVmutation [51] and DeepSequence [52] to predict the effects of novel

mutations. It is also employed in SLIP [54], which uses MSA information to generate synthetic landscapes. These methods are not currently accurate enough in isolation to predict sequence-to-function fitness peaks reliably, but if used in combination with directed evolution (i.e. to inform ruggedness-informed selections of base chance and splitting, as well as promising starting points), they could produce a superior approach overall. The latter form of fitness landscape inference methods are those that take place during a directed evolution experiment and hence require no such bank of sequencing data to produce an MSA. Examples of such tools utilise both intermittent sequencing [39], and phenotypic measurements alone [55] to build a picture of the landscape as evolution takes place. The phenotypic approach uses a trained neural network to make ruggedness predictions from the behaviour of a randomly mutating population. Such an approach would be particularly powerful in combination with the methods outlined in our study, as the entire pipeline could be conducted with no sequencing required (neither prior MSA data nor sequencing from during the experiment). Work remains to be done, however, on bridging the gap between the output of each landscape inference method and selection of the most promising directed evolution strategy.

Another limitation of this study is the reliance on NK models. We chose to use such parameterisable fitness landscapes for testing because empirical landscapes, such as that of GB1 and TrpB, are scarce. NK models do not perfectly represent the statistical properties of natural fitness landscapes, therefore in order to improve the reliability of simulations such as these, future work could tailor the NK model to the specific application of protein fitness landscapes. For instance, by allowing K to be drawn from a distribution for each amino acid position, as opposed to a constant value, or by integrating the information offered by PAM substitution matrices into fitness estimations. NK model variants also exist that emphasise the role of neutral drift, another factor that could be integrated into an alternative NK model [49, 50].

With *de novo* protein design approaches still in their infancy, directed evolution remains an important component of the protein engineering toolbox. Here we demonstrate that the standard approach to directed evolution, which is to select the top variants from a single population, can be sub-optimal both in real and synthetic fitness landscape contexts. We propose two strategies to overcome the limitations of the standard approach; using base chance and population splitting to increase landscape exploration. By combining these techniques we observe that up to a 19-fold increase in the probability of reaching the global optimum is made possible. As our knowledge of how to make predictions about directed evolution strategy develops, most notably with the integration of landscape inference tools, it may become possible to unlock further such improvements in directed evolution without the need for sequencing.

Methods

Model

To test our selection algorithms we implemented *in silico* simulations of directed evolution, which can be applied to either synthetic or empirical landscapes. We model each genetic variant within a population using a one-dimensional array of N integer values (Fig 1), which is initialised at a random starting location (NK) or wildtype sequence (GB1, TrpB), creating a population of size P. To calculate the corresponding fitness value of a gene, this array is either used as input to the NK model (Methods: NK Landscapes), or used as coordinates to look up fitness in an empirical data set (Methods: Empirical Landscapes). Our simulation algorithm proceeds with three cyclic steps; selection, proliferation, and mutation.

Individual cell fitness values are input to a selection function (Fig 1), which translates relative fitness into probability of selection. Then, each cell is selected (or not) based only on its respective probability of selection. Cells that are selected are proliferated to bring the

population back up to its original size. To perform proliferation, a *new* population of size P is created by randomly sampling (with replacement) from the previously selected cells. This introduces a degree of stochasticity mirroring experimental error and biological variation, as described in other models of evolutionary processes [56].

Once a population of size P has been selected and proliferated, mutations are introduced by making random changes to genes in the population. For every cell's genetic code (an array), each residue (i.e. nucleotide or amino acid) has a random chance, p_I, of changing into another random *different* value (i.e. each other possibility is equally likely). Given the gene length N, the expected number of changes across the entire gene (the "mutations per cell") is given by $\mu = Np_I$. This quantity is useful to work with, as a fixed μ will give comparable results as N varies.

The final result is a new population of size P, and the process of selection, proliferation, and mutation may be repeated again. In this work, it is assumed that this runs for a fixed number of iterations, and the final result (by which we compare methods) is the maximum fitness across the final population. "Normalised fitness" divides this measure by the global maximum of the landscape.

Selection functions

In this work, we introduce the concept of a "selection function". This function takes a cell's fitness percentile, defined as the percentage of other cells in the population it is fitter than, and outputs its chance of being selected to be in the next population. The selection function used in this work is defined by two parameters: a "fitness threshold" and a "base chance"(Fig 2A). Above the "fitness threshold" the selection function outputs 1, otherwise it outputs the "base chance" (Eq 1). In this work, selection functions are normalised to select a constant fraction of cells (20% throughout this paper). This is to ensure that in a real continuous directed evolution experiment, the time required for proliferating cells between iterations remains effectively constant such that a "fair" comparison is made between strategies in terms of performance per-unit-time. By fixing the selected proportion (given by the integral of the selection function), our parameter search space is reduced to one dimension, as every base chance has only one possible corresponding fitness threshold value (Eq 2).

$$selection\ chance = \begin{cases} base\ chance & \text{if } fitness\ percentile < fitness\ \ threshold \\ 1 & \text{if } fitness\ \ percentile \geq fitness\ threshold \end{cases} \quad (1)$$

$$fitness\ threshold = \frac{1 - selected\ fraction}{1 - base\ chance} \quad (2)$$

NK landscapes

The NK model is a widely used approach for generating synthetic fitness landscapes with tuneable ruggedness [42, 43]. In the NK model, a gene is represented by an array of N sites, each of which has A possible values. In this work we will generally set $A = 2$ such that each entry has two possibilities (i.e. binary 1 or 0). Every entry corresponds to a "locus", which interacts with K other loci in the gene. The fitness contribution of each locus is dependent on the state of that locus and the state of the K other loci it interacts with. The fitness F of a gene G is the sum of the fitnesses of each locus. When $K = 0$, all loci are independent and the model is linear (and hence has a single peak). The other extreme, $K = N - 1$, in which every locus interacts with

every other locus, is maximally unstructured; the fitness landscape consisting of only random noise.

The NK model defines a parameterised *distribution* over functions $F : A^N \to \mathbb{R}$. To define how to generate samples from this distribution, first let L: $N \times K \to N$ be a "locus function", where $L(a, i)$ is the ith site that locus a interacts with. We require each locus to interact with precisely K other, uniformly random, sites, independent of all other loci. We then also have $N \cdot A^{K+1}$ independent, identically distributed standard normal ($\mathcal{N}(0, 1)$) random variables, which we denote as $X^a_{i_0,i_1,i_2\ldots,i_K}$, where $a \in \{1\ldots N\}$ and $i_{0,1,\ldots,K} \in \{1\ldots A\}$. These represent the possible fitness contributions of each loci, with a being the index of the loci in question, and $i_{0,1,\ldots,K}$ being the values of that loci and the K other loci it interacts with, plus itself (i_0).

Then, F is given as the following, where $g \in A^N$, and $g[i]$ denotes the ith site in g:

$$F(g) = \sum_{a=1}^{N} X^a_{g[a],g[L(a,1)],g[L(a,2)],\ldots,g[L(a,K)]} \tag{3}$$

Computationally, the algorithm used explicitly generates and stores L; in particular as a $N \times N$ binary matrix. However, it does not store the $N \cdot A^{K+1}$ random variables explicitly as doing so would require large amounts of memory. Instead, every time the value of $X^a_{i_0,i_1,i_2\ldots,i_K}$ is required, a, i_0, i_1, . . ., i_K are used as inputs to a pre-defined deterministic pseudorandom generation algorithm.

Empirical landscapes

Empirical fitness landscapes are real data sets produced by measuring the fitness of all possible sequential variants of a protein (or region of a protein). The empirical landscapes used in this paper are GB1 [23] and TrpB [25]. For our algorithm, the landscapes (each four amino acid positions) are stored as four-dimensional arrays, where each dimension corresponds to a variable residue. By assigning every amino acid a number from 1 to 20, each sequence of amino acids can therefore be mapped to a set of coordinates that points to the corresponding fitness value in the array. All directed evolution simulations were started from the wild-type sequence "VDGV" for GB1, and "VFVS" for TrpB.

Supporting information

S1 Fig. Competition plots showing base chance performance over varying NK landscapes. Data shows the proportion of runs that are won by either a 0 base chance, or 0.1 base chance strategy over different landscapes. Where the outcome is tied, the value is 0. Value of N increases vertically, and the value of K/N increases up horizontally. Mutations per cell = 0.1, population size = 1000.
(TIF)

S2 Fig. Competition plots showing the influence of noise on base chance performance over 100 iterations. To introduce noise, values sampled from $\mathcal{N}(0, noise^2)$ are added to fitness values each generation. N = 25, K = 5, mutations per cell = 0.1, population size = 1000.
(TIF)

S3 Fig. Comparing the performance of population splitting across landscapes of increasing ruggedness. Mutations per cell = 0.1.
(TIF)

S4 Fig. Parameter sweep over population sizes (5000,1000,100), mutation rates (0.05/N, 0.1/N, 0.15/N, 0.2/N, 0.25/N), N values (50, 25) and K values (0.25N, 0.2N, 0.15N, 0.1N,

0.05N, 0). Results were processed by ranking each of the 30 strategies (combinations of base chance and splitting) against one another for each set of parameters. A: Average rank overall. B: Histogram of ranks, top performing strategy vs standard strategy (no base chance, no splits).
(TIF)

Author Contributions

Conceptualization: Jessica James, Sebastian Towers, Jakob Foerster, Harrison Steel.

Data curation: Jessica James, Sebastian Towers.

Formal analysis: Jessica James, Sebastian Towers.

Funding acquisition: Harrison Steel.

Investigation: Jessica James, Sebastian Towers.

Methodology: Jessica James, Sebastian Towers.

Project administration: Jessica James, Sebastian Towers, Jakob Foerster, Harrison Steel.

Resources: Jakob Foerster, Harrison Steel.

Software: Jessica James, Sebastian Towers.

Supervision: Jakob Foerster, Harrison Steel.

Validation: Jessica James, Sebastian Towers.

Visualization: Jessica James, Sebastian Towers.

Writing – original draft: Jessica James, Sebastian Towers.

Writing – review & editing: Jessica James, Sebastian Towers, Jakob Foerster, Harrison Steel.

References

1. Yan X, Liu X, Zhao C, Chen G. Applications of synthetic biology in medical and pharmaceutical fields. Signal Transduct Target Ther. 2023; 8: 199. https://doi.org/10.1038/s41392-023-01440-5 PMID: 37169742
2. Scown CD, Keasling JD. Sustainable manufacturing with synthetic biology. Nat Biotechnol. 2022; 40: 304–307. https://doi.org/10.1038/s41587-022-01248-8 PMID: 35190687
3. Sargent D, Conaty WC, Tissue DT, Sharwood RE. Synthetic biology and opportunities within agricultural crops. Journal of Sustainable Agriculture and Environment. 2022; 1(2): 89–107. https://doi.org/10.1002/sae2.12014
4. Dauparas J, Anishchenko I, Bennett N, Bai H, Ragotte RJ, Milles LF et al. Robust deep learning–based protein sequence design using ProteinMPNN. Sci. 2022; 378(6615): 49–56. https://doi.org/10.1126/science.add2187 PMID: 36108050
5. Ferruz N, Schmidt S, Höcker B. ProtGPT2 is a deep unsupervised language model for protein design. Nat. Commun. 2022; 13(1): 4348. https://doi.org/10.1038/s41467-022-32007-7 PMID: 35896542
6. Singer JM, Novotney S, Strickland D, Haddox HK, Leiby N, Rocklin GJ et al. Large-scale design and refinement of stable proteins using sequence-only models. PLoS ONE. 2022; 17(3):e0265020. https://doi.org/10.1371/journal.pone.0265020 PMID: 35286324
7. Anishchenko I, Pellock SJ, Chidyausiku TM, Ramelot TA, Ovchinnikov S, Hao J et al. De novo protein design by deep network hallucination. Nat. 2021; 600(7889): 547–552. https://doi.org/10.1038/s41586-021-04184-w PMID: 34853475
8. Wicky BIM, Milles LF, Courbet A, Ragotte RJ, Dauparas J, Kinfu E et al. Hallucinating symmetric protein assemblies. Sci. 2022; 378(6615):56–61. https://doi.org/10.1126/science.add1964 PMID: 36108048
9. Arnold FH. Design by Directed Evolution. Acc. Chem. Res. 1998; 31(3): 125–131. https://doi.org/10.1021/ar960017f

10. Nixon AE, Sexton DJ, Ladner RC. Drugs derived from phage display. mAbs. 2014; 6(1): 73–85. https://doi.org/10.4161/mabs.27240 PMID: 24262785
11. Heater BS, Chan WS, Lee MM, Chanet MK. Directed evolution of a genetically encoded immobilized lipase for the efficient production of biodiesel from waste cooking oil. Biotechnol Biofuels. 2019; 12 (1):165. https://doi.org/10.1186/s13068-019-1509-5 PMID: 31297153
12. Neuenschwander M, Butz M, Heintz C, Kast P, Hilvert D. A simple selection strategy for evolving highly efficient enzymes Nat Biotechnol. 2007; 25(10): 1145–1147. https://doi.org/10.1038/nbt1341 PMID: 17873865
13. Wu Y, Jameel A, Xing X, Zhang C. Advanced strategies and tools to facilitate and streamline microbial adaptive laboratory evolution Trends Biotechnol. 2022; 41(1): 38–59. https://doi.org/10.1016/j.tibtech.2021.04.002 PMID: 33958227
14. Yang G, Withers SG. Ultrahigh-throughput FACS-based screening for directed enzyme evolution Chembiochem. 2009; 10(17): 2704–2715. https://doi.org/10.1002/cbic.200900384 PMID: 19780076
15. Trivedi VD, Mohan K, Chappel TC, Mays ZJS, Nair NU. Cheating the cheater: Suppressing false positive enrichment during biosensor-guided biocatalyst engineering ACS Synth Biol. 2022; 11(1): 420–429. https://doi.org/10.1021/acssynbio.1c00506 PMID: 34914365
16. Luro S, Potvin-Trottier L, Okumus B, Paulsson J. Isolating live cells after high-throughput, long-term, time-lapse microscopy Nat Methods. 2020; 17(1). https://doi.org/10.1038/s41592-019-0620-7 PMID: 31768062
17. Sheets MB, Tague N, Dunlop MJ. An Optogenetic Toolkit for Light-Inducible Antibiotic Resistance Nat Commun. 2023; 14(1). https://doi.org/10.1038/s41467-023-36670-2 PMID: 36823420
18. Halperin SO, Tou CJ, Wong EB, Modavi C, Schaffer DV and Dueber JE. CRISPR-guided DNA polymerases enable diversification of all nucleotides in a tunable window. Nat. 2023; 560: 248–252. https://doi.org/10.1038/s41586-018-0384-8
19. Moore CL, Papa LJ, and Shoulders MD. A Processive Protein Chimera Introduces Mutations Across Defined DNA Regions In Vivo. J. Am. Chem. Soc. 2018; 140: 11560–11564. https://doi.org/10.1021/jacs.8b04001 PMID: 29991261
20. Rix G, Watkins-Dulaney EJ, Almhjell PJ, Boville CE, Arnold FH and Liu CC. Scalable continuous evolution for the generation of diverse enzyme variants encompassing promiscuous activities. Nat. Commun. 2020; 11. https://doi.org/10.1038/s41467-020-19539-6 PMID: 33159067
21. Molina RS, Rix G, Mengiste AA, Álvarez B, Seo D, Chen H et al. In vivo hypermutation and continuous evolution. Nat Rev Methods Primers. 2022; 2(1): 1–22. https://doi.org/10.1038/s43586-022-00130-w PMID: 37073402
22. Wright S. The Roles of Mutation, Inbreeding, crossbreeding and Selection in Evolution. Proceedings of the XI International Congress of Genetics. 1932; 8: 209–222.
23. Wu NC, Dai L, Olson CA, LLoyd-Smith JO, Sun R. Adaptation in protein fitness landscapes is facilitated by indirect paths. eLife. 2016; 5:e16965. https://doi.org/10.7554/eLife.16965 PMID: 27391790
24. Papkou A, Garcia-Pastor L, Escudero JA, Wagner A. A rugged yet easily navigable fitness landscape. Sci. 2023; 382(6673):eadh3860. https://doi.org/10.1126/science.adh3860 PMID: 37995212
25. Johnston KE, Almhjell PJ, Watkins-Dulaney EJ, Liu G, Porter NJ, Yang J et al. A combinatorially complete epistatic fitness landscape in an enzyme active site. Proc. Natl. Acad. Sci. U.S.A. 2024; 121(32): e2400439121. https://doi.org/10.1073/pnas.2400439121 PMID: 39074291
26. Leopoldo FM, Currin A, Dixon N. Directed evolution of the PcaV allosteric transcription factor to generate a biosensor for aromatic aldehydes. J. Biol. Eng. 2019; 13(1):91. https://doi.org/10.1186/s13036-019-0214-z
27. Chen S, Yang Z, Zhong Z, Yu S, Zhou J, LI J et al. Ultrahigh-throughput screening-assisted in vivo directed evolution for enzyme engineering. Biotechnology for Biofuels and Bioproducts. 2024; 19(9). https://doi.org/10.1186/s13068-024-02457-w PMID: 38254175
28. LaCroix RA, Palsson BO, Feist AM. A Model for Designing Adaptive Laboratory Evolution Experiments. AEM. 2017; 83(8):e03115–16. https://doi.org/10.1128/AEM.03115-16
29. Tan Y, Zhang Y, Han Y, Hui H, Chen H, Ma F et al. Directed evolution of an α1,3-fucosyltransferase using a single-cell ultrahigh-throughput screening method Sci. Adv. 2019; 5(10). https://doi.org/10.1126/sciadv.aaw8451 PMID: 31633018
30. Tian D, Wang C, Liu Y, Zhang Y, Caliari A, Lu H et al. Cell Sorting-Directed Selection of Bacterial Cells in Bigger Sizes Analyzed by Imaging Flow Cytometry during Experimental Evolution Int. J. Mol. Sci. 2023; 24(4): 3243. https://doi.org/10.3390/ijms24043243 PMID: 36834655
31. Tu R, Martinez R, Prodanovic R, Klein M, Schwaneberg U. A Flow Cytometry–Based Screening System for Directed Evolution of Proteases SLAS Discov. 2011; 16(3): 285–294. https://doi.org/10.1177/1087057110396361 PMID: 21335599

32. Romero PA, Arnold FH. Exploring protein fitness landscapes by directed evolution. Nature Reviews Molecular Cell Biology. 2009; 10(12): 866–876. https://doi.org/10.1038/nrm2805 PMID: 19935669

33. Fox R, Roy A, Govindarajan S, Minshull J, Gustafsson C, Jones JT et al. Optimizing the search algorithm for protein engineering by directed evolution. PEDS. 2003; 16(8): 589–597. https://doi.org/10.1093/protein/gzg077 PMID: 12968076

34. Wu Z, Kan SBJ, Lewis RD, Wittmann BJ, Arnold FH. Machine learning-assisted directed protein evolution with combinatorial libraries. PNAS. 2019; 116(18): 8852–8858. https://doi.org/10.1073/pnas.1901979116 PMID: 30979809

35. Wittmann BJ, Yue Y, Arnold FH. Informed training set design enables efficient machine learning-assisted directed protein evolution. Cell. Syst. 2021; 12(11):1026–1045.e7. https://doi.org/10.1016/j.cels.2021.07.008 PMID: 34416172

36. Yang KK, Wu Z, Arnold FH. Machine-learning-guided directed evolution for protein engineering. Nat. Methods. 2019; 16(8): 687–694. https://doi.org/10.1038/s41592-019-0496-6 PMID: 31308553

37. Frisby TS, Langmead CJ. Bayesian optimization with evolutionary and structure-based regularization for directed protein evolution. Algorithms Mol Biol. 2021; 16(1): 13. https://doi.org/10.1186/s13015-021-00195-4 PMID: 34210336

38. Hu R, Fu L, Chen Y, Chen J, Qiao Y, Si T. Protein engineering via Bayesian optimization-guided evolutionary algorithm and robotic experiments. Brief. Bioinform. 2023; 24(1). https://doi.org/10.1093/bib/bbac570 PMID: 36562723

39. D'Costa S, Hinds EC, Freschlin CR, Song H, Romero PA. Inferring protein fitness landscapes from laboratory evolution experiments. PLoS. Comput. Bio. 2023; 19(3):e1010956. https://doi.org/10.1371/journal.pcbi.1010956 PMID: 36857380

40. Carpenter AC, Feist AM, Harrison FSM, Paulsen IT, Williams TC. Have you tried turning it off and on again? Oscillating selection to enhance fitness-landscape traversal in adaptive laboratory evolution experiments. Metab Eng Commun. 2023; 17: e00227 https://doi.org/10.1016/j.mec.2023.e00227 PMID: 37538933

40. Alpay B A, Desai M M. Effects of selection stringency on the outcomes of directed evolution. PLoS ONE.

42. Kauffman S, Levin S. Towards a general theory of adaptive walks on rugged landscapes. J. Theor. Biol. 1987; 128(1): 11–45 https://doi.org/10.1016/S0022-5193(87)80029-2 PMID: 3431131

43. Kauffman S, Weinberger ED. The NK model of rugged fitness landscapes and its application to maturation of the immune response. J. Theor. Biol. 1989; 141(2): 211–245 https://doi.org/10.1016/S0022-5193(89)80019-0 PMID: 2632988

44. Iwasa Y, Franziska M, Nowak MA. Stochastic tunnels in evolutionary dynamics. Genetics. 2004; 166 (3): 1571–1579 https://doi.org/10.1534/genetics.166.3.1571 PMID: 15082570

45. Ochs IE, Desai MM. The competition between simple and complex evolutionary trajectories in asexual populations BMC Evol. Biol. 2015; 15(1) https://doi.org/10.1186/s12862-015-0334-0 PMID: 25881244

46. Vedal LS, Isalan M, Heap JT, Ledesma-Amaro R. A primer to directed evolution: current methodologies and future directions. RSC. 2023; 4.

47. Gavrilets S. Perspective: Models of Speciation: What Have We Learned in 40 Years? SSE. 2003; 57 (10): 2197–2215. PMID: 14628909

48. Chen Z, Kang L. Multi-Population Evolutionary Algorithm for Solving Constrained Optimization Problems Artificial Intelligence Applications and Innovations (IFIP Conference). 2005.

49. Barnett L. Ruggedness and neutrality—the NKp family of fitness landscapes ALIFE: Proceedings of the sixth international conference on Artificial life. 1998. pp18–27.

50. Newman MEJ, Engelhardt R. Effects of neutral selection on the evolution of molecular species arXiv: adap-org/9712005v1. 1997. Available from: https://arxiv.org/abs/adap-org/9712005.

51. Hopf TA, Ingraham JB, Poelwijk FJ, Schärfe CP, Springer M, Sander C et al. Mutation effects predicted from sequence co-variation. Nat Biotechnol. 2017; 35(2): 128–135. https://doi.org/10.1038/nbt.3769 PMID: 28092658

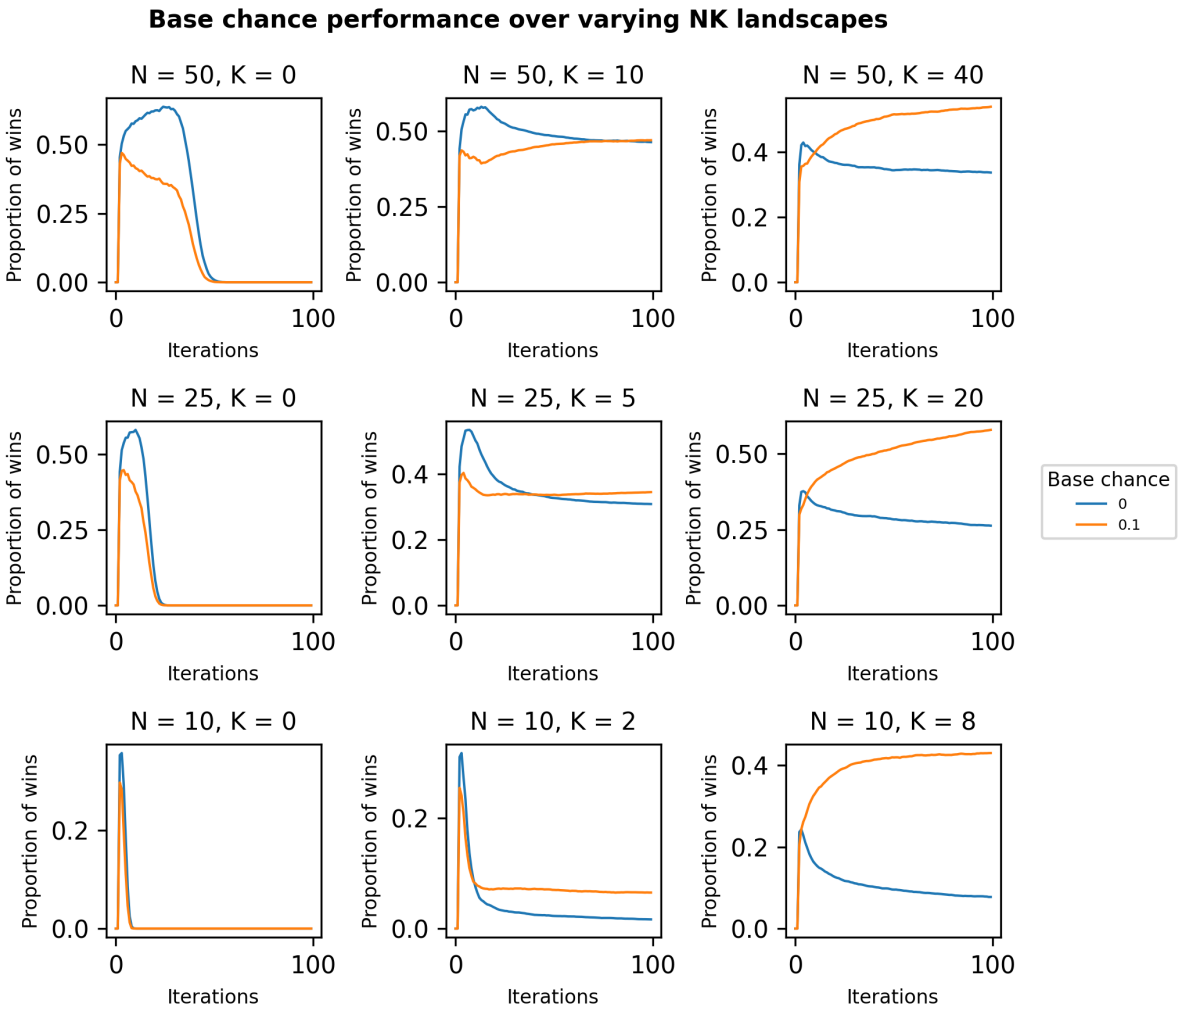
52. Riesselman AJ, Ingraham JB and Marks DS. Deep generative models of genetic variation capture the effects of mutations. Nat Methods. 2018; 15: 816–822. https://doi.org/10.1038/s41592-018-0138-4 PMID: 30250057

53. Rahi SJ, Gligorovski V. Directed evolution of dynamic, multi-state, and computational proteins. Cell Press. 2024; 123(3).

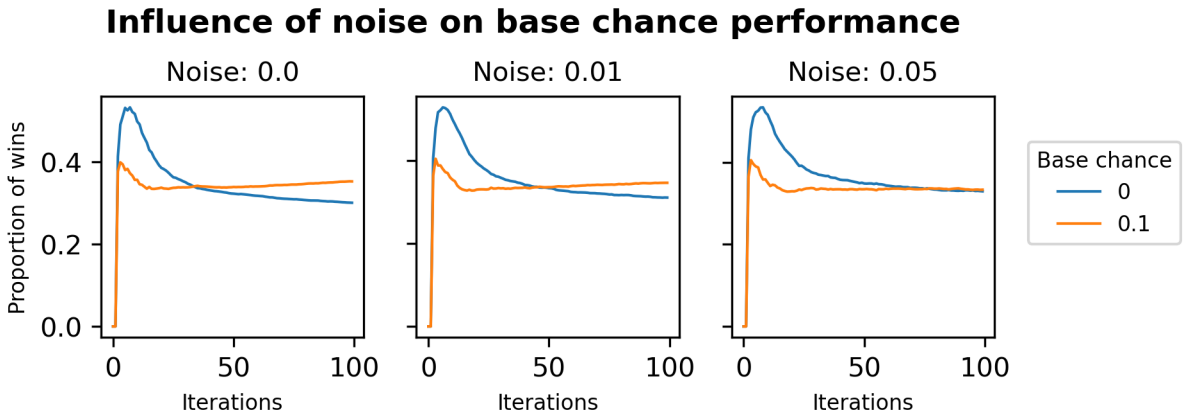
54. Thomas N, Agarwala A, Belanger D, Song YS, Colwell LJ. Tuned Fitness Landscapes for Benchmarking Model-Guided Protein Design. BioRxiv [Preprint]. 2022. Available from: https://www.biorxiv.org/content/10.1101/2022.10.28.514293v1.

55. Towers S, James J, Steel H, Kempf I. Learning-Based Estimation of Fitness Landscape Ruggedness for Directed Evolution. BioRxiv. [Preprint]. 2024. Available from: https://www.biorxiv.org/content/10.1101/2024.02.28.582468v1.

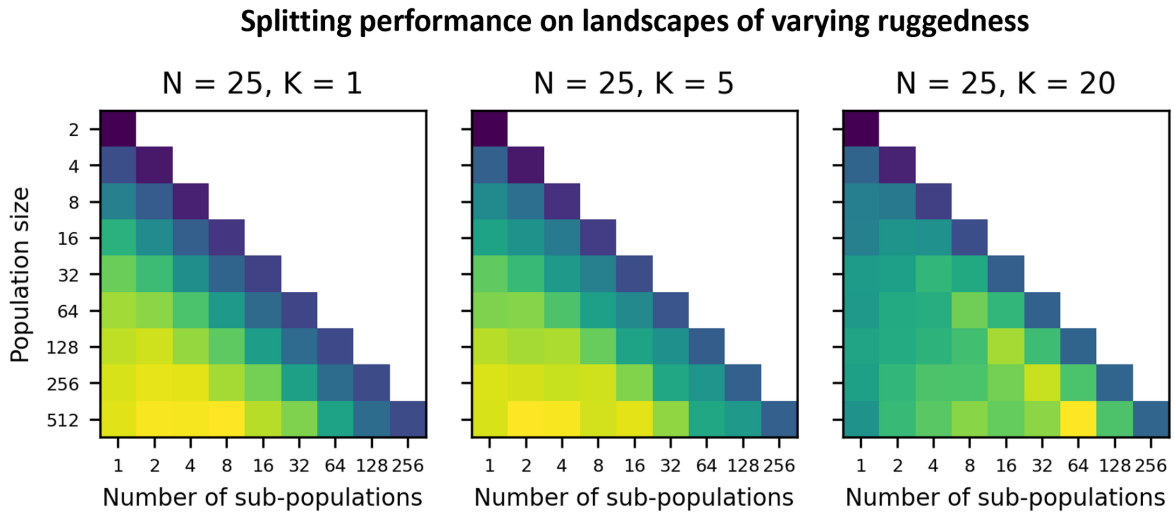
56. Desai MM, Fisher DS. Beneficial Mutation-Selection Balance and the Effect of Linkage on Positive Selection. Genetics. 2007; 176(1): 1759–1798. https://doi.org/10.1534/genetics.106.067678 PMID: 17483432



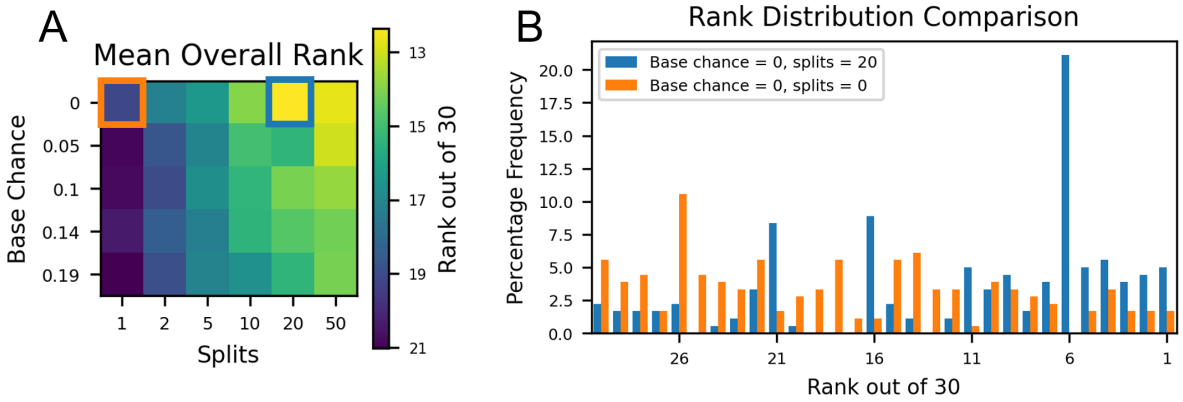
S1 Fig. Competition plots showing base chance performance over varying NK landscapes. Data shows the proportion of runs that are won by either a 0 base chance, or 0.1 base chance strategy over different landscapes. Where the outcome is tied, the value is 0. Value of N increases vertically, and the value of K/N increases up horizontally. Mutations per cell = 0.1, population size = 1000.



S2 Fig. Competition plots showing the influence of noise on base chance performance over 100 iterations. To introduce noise, values sampled from $\mathcal{N}(0, \text{noise}^2)$ are added to fitness values each generation. $N = 25$, $K = 5$, mutations per cell $= 0.1$, population size $= 1000$.



S3 Fig. Comparing the performance of population splitting across landscapes of increasing ruggedness. Mutations per cell = 0.1.



S4 Fig. Parameter sweep over population sizes $(5000, 1000, 100)$**, mutation rates** $(0.05/N, 0.1/N, 0.15/N, 0.2/N, 0.25/N)$**,** N **values** $(50, 25)$ **and** K **values** $(0.25N, 0.2N, 0.15N, 0.1N, 0.05N, 0)$**.** Results were processed by ranking each of the 30 strategies (combinations of base chance and splitting) against one another for each set of parameters. A: Average rank overall. B: Histogram of ranks, top performing strategy vs standard strategy (no base chance, no splits).

3.5 Limitations and Future Work

This work is subject to several key limitations regarding the underlying fitness landscapes and the generalisability of the selection strategies, which are outlined in this section.

3.5.1 Limitations of the NK Model

A primary limitation concerns the nature of the NK model itself. While the NK model is a widely accepted framework providing tunable ruggedness, it is highly simplified and fails to capture several statistical and structural properties inherent to real protein landscapes.

Specifically, the NK model produces statistically homogeneous landscapes; for a given landscape, the properties tend to remain consistent regardless of the position within the sequence space. Furthermore, the distribution of fitness values does not reflect biological reality. The model typically produces a symmetric normal distribution rather than the right-skewed distribution observed in nature, where the majority of variants exhibit low fitness and a small "tail" extends toward high-fitness peaks.

This introduces challenges when evaluating search strategies. Because the difference between the top 1% and 0.1% fitness values is marginal, and the gap to the top 0.01% smaller still, distinguishing between effective and suboptimal strategies becomes difficult. To mitigate this, we employ normalised fitness to map observed values to a uniform scale, highlighting performance deltas more clearly. However, in an ideal scenario, the underlying mathematical model would natively exhibit a realistic fitness distribution, obviating the need for such transformations.

Ultimately, the NK model, relying on only two parameters (N and K), is likely too simple to represent the rich complexity of biological fitness landscapes. While the impact of these limitations is difficult to quantify without a broader class of synthetic benchmarks, the absence of major qualitative features suggests that the development of more biologically representative synthetic landscapes remains a vital frontier for this research.

3.5.2 Empirical Validation and Scale

These theoretical constraints lead to the necessity of empirical validation. In this work, we utilise combinatorially complete empirical landscapes sourced from experimental protein data. While *de novo* wet-lab validation would be ideal, the cost of generating sufficient data to achieve statistical significance necessitated the use of these *in silico* empirical benchmarks.

Nevertheless, these landscapes harbour a fundamental trade-off: to be combinatorially complete, the sequence space must be restricted. In practice, this limits the scope to approximately four amino acid sites. This creates an unfortunate situation: if a landscape is small enough to be fully mapped, reaching the global optimum is often trivial. Real-world directed evolution experiments typically involve much larger sequence spaces, even if only a fraction of sites are mutated. Thus, while these empirical landscapes are a significant improvement over purely mathematical models, they do not fully mirror the scale or search dynamics of experimental directed evolution.

3.5.3 Strategy Selection and Ruggedness Inference

Finally, although we demonstrate that non-greedy selection strategies can outperform greedy ones, we provide no direct method for determining *a priori* which strategy is optimal for a specific protein of interest. While we characterise performance across various regimes of the NK model, there is no standardised protocol to map a physical protein to a specific (N, K) coordinate.

Consequently, while this work establishes the existence of performance gains, practical guidance for experimentalists remains limited to parameters that perform "broadly well". This limitation motivates the following chapter, which introduces an experimentally tractable approach for inferring landscape ruggedness, enabling practitioners to select more effective strategies for their specific experimental contexts.

The only thing that makes life possible is the permanent, intolerable uncertainty: not knowing what comes next.

Ursula K. Le Guin *The Dispossessed*

4

Parameter Inference for Directed Evolution

Contents

4.1 Motivation 49
4.2 Summary 50
4.3 Contributions 51
4.4 Paper 52
4.5 Limitations and Future work 93
4.5.1 Experimental Viability 93
4.5.2 NK landscapes 95
4.5.3 Ruggedness Metrics 95

This chapter presents *Sequence-free landscape inference for directed evolution*, published in *PLOS Computational Biology* (2026) [72].

4.1 Motivation

The previous chapter demonstrated that directed evolution presents us with a rich optimisation problem. However, a key limitation of that work was the absence of a practical method to infer the properties of a fitness landscape prior to running an experiment. Because different landscapes favour different search strategies, this leaves experimentalists with little guidance on how to choose an appropriate approach in advance.

More fundamentally, there is limited consensus on how fitness landscapes should be characterised. The most commonly discussed attribute is ruggedness, which informally describes how difficult a landscape is to traverse due to the presence of multiple local optima. Despite its intuitive appeal, ruggedness lacks a single unique definition: several mathematically distinct measures exist, each capturing different aspects of the same idea. Worse still, many of these measures require detailed knowledge of the entire landscape, such as the location of the global optimum, in order to compute them. Needless to say, if the global optimum were already known, directed evolution would be unnecessary.

What is needed, therefore, is a ruggedness measure that is both experimentally accessible and predictive of evolutionary behaviour. Such a measure would allow researchers to characterise new landscapes with minimal prior information and select strategies tailored to their underlying structure, improving the efficiency and interpretability of directed evolution experiments.

4.2 Summary

To address these challenges, this chapter, based on our paper **Sequence-free landscape inference for directed evolution**, develops a method to infer fitness landscape ruggedness directly from experimental data, without sequencing. We argue that the most informative and practically useful measure of ruggedness arises from **spectral landscape theory**, which analyses landscapes using tools from Fourier analysis. Much like a sound can be decomposed into high- and low-frequency components, a fitness landscape can be represented by its Fourier spectrum, where greater contributions from high-frequency terms indicate more rough topographies.

Although conceptually more involved, this approach offers two major advantages. First, we show that the Fourier spectrum can be inferred empirically by tracking the average fitness of a randomly mutating population over time. This relationship is not merely approximate: we show that, under ideal conditions, suitably averaged population-level fitness measurements determine the landscape's power spectrum. However, in reality, the number of samples needed to precisely estimate the

full spectrum is prohibitively large, so we instead focus on estimating a single parameter that summarises the spectrum. Regardless, this provides a rigorous and experimentally accessible route to estimate ruggedness without sequencing, relying solely on population-level fitness measurements.

Second, the inferred spectral measure aligns closely with existing notions of ruggedness, including those defined in the NK model, which allows us to map empirical landscapes to their closest NK equivalents. By leveraging this mapping, high-performing directed-evolution strategies identified *in silico* can inform the choice of strategies for wet-lab experiments. Using both simulated and empirical fitness landscapes, we demonstrate that this method reliably identifies near-optimal strategies from our strategy space, yielding substantial performance gains compared to naïve selection protocols.

We further characterise how the accuracy of the inferred spectrum depends on experimental design parameters, such as sample size, and benchmark our ruggedness measure against alternative metrics across both NK and empirical landscapes. Finally, significant effort is devoted to presenting the details of the Fourier approach in an intuitive manner, aiming to make this framework accessible to experimentalists. Overall, this work demonstrates that spectral ruggedness provides both a powerful theoretical tool and a practical method for guiding directed evolution.

4.3 Contributions

This work also represents a joint first authorship between Jessica James and Sebastian Towers, with Idris Kempf leading supervision of the project. The core methodology and the mathematical background were developed by S. Towers, with I. Kempf assisting in writing up the mathematics to be as clear and approachable as possible. Figure 2, explaining the methodology, was made by S. Towers, along with all the code to determine the ruggedness metrics from a landscape. J. James and S. Towers ran the experiments for Figure 3, seeing how accurately the Fourier spectra may be determined on NK landscapes, alongside how different ruggedness measures compare on the NK landscape. J. James ran the experiments for Figure

4, which is similar to Figure 3, but applied to empirical landscapes. J. James also ran the experiments for Figure 5, which demonstrate how one can determine the Fourier spectra of a landscape, look up the best-performing selection parameters from those tested, and then run directed evolution on said landscape. S. Towers ran the experiments for Figure 6, which examines mutation bias and landscape perturbations. Finally, I. Kempf ran the experiments for Figure 7, alongside writing all the work regarding statistical parameter estimation. Harrison Steel provided general advice and guidance throughout the project.

4.4 Paper

PLOS Computational Biology

RESEARCH ARTICLE

Sequence-free landscape inference for directed evolution

Sebastian Towers☯, **Jessica James**☯, **Harrison Steel***, **Idris Kempf***

Department of Engineering Science, University of Oxford, Oxford, United Kingdom

☯ These authors contributed equally to this work.
* idris.kempf@eng.ox.ac.uk (IK); harrison.steel@eng.ox.ac.uk (HS)

Abstract

Directed evolution is a method for engineering biological systems or components, such as proteins, wherein desired traits are optimised through iterative rounds of mutagenesis and selection of fit variants. The process of protein directed evolution can be envisaged as navigation over high-dimensional optimisation landscapes with numerous local maxima. The performance of any strategy in navigating such a landscape is dependent on the ruggedness of that landscape. However, this information is generally unavailable at the outset of an experiment. Here we propose **SLIDE**, **S**equence-free **L**andscape **I**nference for **D**irected **E**volution, which consists of two parts. First, SLIDE provides an estimation of landscape ruggedness from a mutating population using only population-level phenotypic data and an estimate of the mutation rate. Such ruggedness information in itself is valuable in protein design, for instance in predicting evolutionary stability. Second, SLIDE offers a framework for using the estimated ruggedness metric to identify high-performing selection strategies for directed evolution. Using theoretical *NK* landscapes and four empirical protein fitness landscapes, we demonstrate consistent *in silico* improvement upon the performance of fixed-parameter strategies, using a pipeline that could also be combined with emerging AI-based methods for driving directed evolution.

Author summary

Directed evolution (DE) navigates high-dimensional protein fitness landscapes to find improved variants, but the success of any strategy depends on the landscape's ruggedness, which is typically unknown and often requires genetic sequencing to derive. We present Sequence-free Landscape Inference for Directed Evolution (SLIDE), a two-step framework that (i) estimates landscape ruggedness from population-level phenotypic decay curves together with a mutation rate estimate, and (ii) uses that ruggedness estimate to choose DE control parameters that better balance exploration and exploitation. In contrast to existing ruggedness metrics, SLIDE relies on spectral analysis to link the exponential

OPEN ACCESS

Citation: Towers S, James J, Steel H, Kempf I (2026) Sequence-free landscape inference for directed evolution. PLoS Comput Biol 22(9): e1014713. https://doi.org/10.1371/journal.pcbi.1014713

Editor: Zachary Ardern, Wellcome Sanger Institute, UNITED KINGDOM OF GREAT BRITAIN AND NORTHERN IRELAND

Received: October 10, 2025

Accepted: August 12, 2026

Published: September 3, 2026

Copyright: © 2026 Towers et al. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data availability statement: Data and code used for running experiments and plotting are available on GitHub at https://github.com/Steel-Lab-Oxford/SLIDE. Raw data are available at https://zenodo.org/records/21282413.

Funding: This work was supported by the Engineering and Physical Sciences Research

Council (EPSRC, EP/Y014073/1 to HS and IK). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

decay rate of mean fitness under random, unbiased mutations to the dominant frequency content of the landscape. This decay rate is related to the normalised Dirichlet energy of the underlying genotype graph, and the framework naturally extends to biased mutation spectra. On homogeneous landscapes, such as *NK* landscapes, ruggedness can be accurately estimated from a single starting genotype. Heterogeneous empirical landscapes instead require sampling multiple starting points, naturally distinguishing local from global ruggedness estimates. We combine SLIDE with DE strategies developed in our previous work, but the method is strategy-agnostic and can be paired with any DE workflow, including emerging AI-driven workflows. *In silico* tests show that SLIDE reliably recovers landscape ruggedness and enables more effective selection of directed evolution strategies than fixed-parameter approaches across landscapes of varying ruggedness. Although heterogeneous landscapes require fitness decay measurements from multiple starting points, these can be generated in parallel using high-throughput mutagenesis methods, such as error-prone PCR or chemical mutagenesis, thereby avoiding the need for extensive sequencing or combinatorially complete genotype–phenotype maps while remaining scalable to large screens. Beyond DE, the approach can rapidly characterise landscapes arising in genetic circuits, enzyme libraries, and ecological models, providing a way to quantify genetic stability and guide experimental or computational search.

Introduction

Fitness landscapes map DNA or amino acid sequences to a measure of fitness such as enzyme activity, cell growth, binding affinity, or stability [1]. Fitness landscapes are characterised by multiple structural properties, such as ruggedness, epistasis, neutrality, and navigability, which have significant consequences for evolution and stability [2,3]. In this work, we focus on inferring the ruggedness of a fitness landscape, which unlocks the ability to assess the robustness of a protein to mutations and to optimise directed evolution (DE) protocols.

Various metrics have been proposed to define and quantify the ruggedness of fitness landscapes [4–8]. Although these metrics have been shown to correlate across several empirical landscapes [8], most of them require combinatorially complete genotype–phenotype maps to compute a ruggedness measure. Several of these metrics come from spectral landscape theory [9], which describes fitness landscapes using tools from Fourier analysis. In the same way that a signal can be decomposed into a sum of waves, a fitness landscape can be decomposed into a sum of landscapes of different frequencies (ruggedness). Where low-frequency terms dominate, the landscape can appear smooth, whereas where high-frequency terms dominate, the landscape can appear rugged.

In this paper, we develop Sequence-free Landscape Inference for Directed Evolution (SLIDE) to efficiently estimate the ruggedness of a landscape using phenotypic information alongside an estimated mutation rate (Fig 1), and we validate this

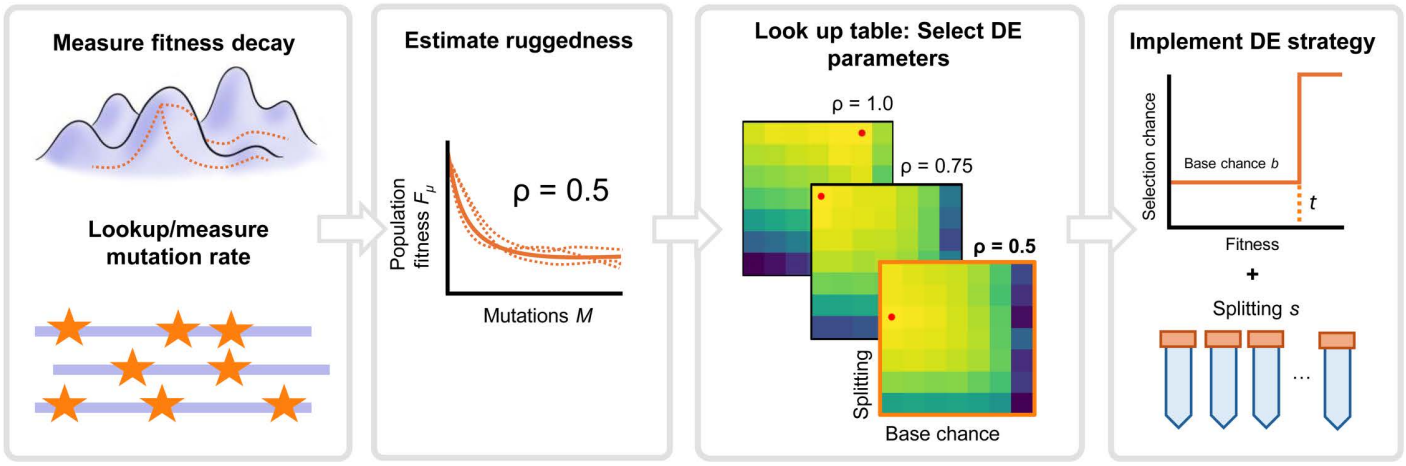


Fig 1. Schematic outline of a future DE pipeline. The initial phase of SLIDE involves random mutation away from one or several starting points and measurement of the resulting fitness decay curve, which, alongside a known mutation rate, can be used to estimate the fitted decay rate ρ^{fit} as a measure of ruggedness. A look-up table matches analytical ρ values of *NK* landscapes to optimal DE strategy parameters. This look-up table can be used to select appropriate parameters for the measured ρ^{fit}. Parameters explored in this paper include base chance, which assigns a fitness-independent probability of selection, and population splitting [10].

https://doi.org/10.1371/journal.pcbi.1014713.g001

pipeline *in silico*. Based on spectral landscape theory and under the assumption of random, unbiased mutations, we define a novel ruggedness metric as a normalised Dirichlet energy of the genotype graph, which can be related to the dominant Fourier terms of the landscape. We show that this metric can be estimated from the average fitness decay of a population that is randomly mutated from a starting point, unlike existing ruggedness metrics, which require combinatorially complete genotype–phenotype maps. The ruggedness metric characterises the rate at which the mean population fitness relaxes towards the landscape average; on smooth landscapes, this average fitness decays more slowly than on rugged landscapes. This phenomenon can be attributed to fitness correlations between variants on smooth landscapes, as random mutations are less likely to result in large fitness differences. Most importantly, the fitness decay curve can be acquired without large-scale sequencing data. Instead, its estimation requires only population-level fitness readouts and an estimate of the mutation rate.

In SLIDE, we combine ruggedness inference with DE. In DE, biological components such as proteins are engineered and improved through iterative rounds of selection and mutagenesis [11]. It can be seen as navigating a high-dimensional fitness landscape, aiming to guide the fitness of variants to a high-performing location on the landscape. Early DE optimisation techniques such as ProSAR used statistical models to correlate sequence with measured activity and iteratively design variants that maximise predicted fitness [12]. Since then, advances in machine learning have led to the development of more sophisticated algorithms that actively design the sequence of variants that form the next generation [13–15]. Although effective, these methods are limited by their need for high-throughput sequencing data. Without sequencing data, most DE experiments select a fixed top proportion of variants in each iteration [16–18], i.e., truncation selection (comparable to hill-climbing algorithms), which we refer to as the “baseline strategy”. This approach results in minimal landscape exploration, and is therefore prone to becoming trapped in local optima [10].

To address this limitation, several methods have been proposed for balancing landscape exploration and exploitation. Such methods include applying non-constant selection pressure [19], increasing the proportion of the population selected [20], splitting the population into sub-populations, and including a base chance of selection that is irrespective of fitness [10]. However, effective application of these methods depends on knowledge of landscape properties such as ruggedness [10]. Inferring the ruggedness beforehand is essential for selecting strategy parameters to ensure the balance of

exploration and exploitation. Here, we select these parameters based on a ruggedness estimate obtained using SLIDE, in which ruggedness is estimated and the strategy parameters are selected before the DE experiment. By doing so, SLIDE increases the likelihood of reaching higher-fitness variants within a fixed number of generations, even when the decay rate is not estimated accurately.

SLIDE is strategy-agnostic and can be integrated into DE workflows both with and without sequencing by using the inferred ruggedness to tune strategy parameters that control the balance between exploration and exploitation. Although we are not aware of prior DE workflows that explicitly use a measured ruggedness scalar for this purpose, many established methods already expose suitable control knobs. For example, machine-learning-assisted DE with combinatorial libraries uses measured sequence–function data to guide library design [13], while Bayesian-optimisation methods formalise exploration–exploitation trade-offs through acquisition-function choice and aggressiveness [21]. Recombination-based workflows provide another natural target, since the optimal selection pressure can depend on whether variation is supplied primarily by mutation or recombination [22]. In this paper, SLIDE is interfaced with strategies developed in our previous work [10], controlling base chance and population-splitting parameters. We evaluate the resulting pipeline *in silico* on theoretical landscapes from the *NK* model [23] and on four empirical landscapes [2,24–26]. SLIDE automatically selects landscape-dependent strategy parameters, improving over fixed-parameter strategies such as best-performer selection on rugged landscapes and overly exploratory selection on smoother landscapes.

Results

Ruggedness inference from fitness decay rate

Ruggedness of a fitness landscape can be defined in terms of various mathematically distinct metrics (see [4–8]), which aim to capture how fitness changes when the landscape is traversed. Fig 2A–2B presents examples of smooth and

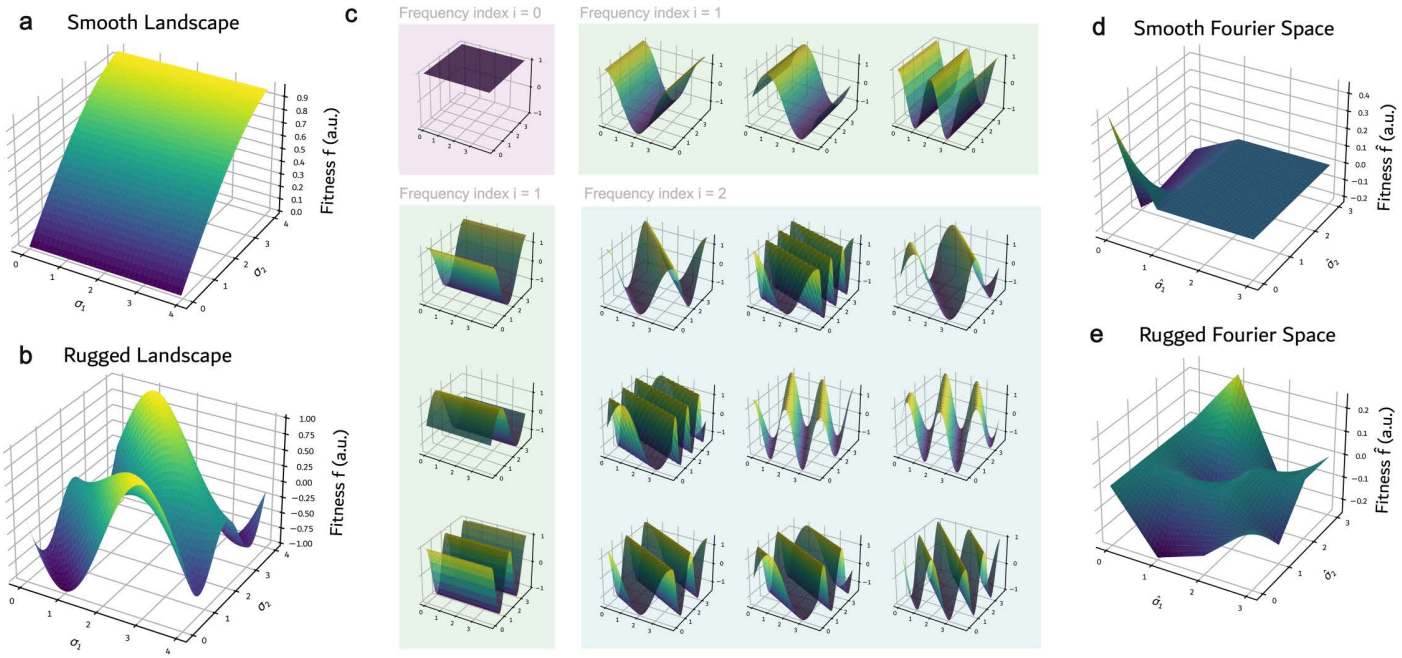


Fig 2. Demonstration of landscape transformations. Examples of **A:** smooth and **B:** rugged landscapes. **C:** Example Fourier basis. **D–E:** Smooth and rugged landscapes transformed into Fourier space with the example basis.

https://doi.org/10.1371/journal.pcbi.1014713.g002

rugged landscapes for a protein of length $N=2$ with 16 genotypes, where σ_i is used to denote the allele at sequence position $i = \{1, 2\}$, where each σ_i takes one of $A=4$ possible amino-acid values. Intuitively, smooth landscapes should have high fitness variants clustered together, leading to high fitness correlations and relatively few local extrema. In contrast, rugged landscapes should exhibit large fitness differences between adjacent nodes, resulting in more local extrema and lower fitness correlations.

We consider a population $\boldsymbol{\nu}$ that traverses a landscape via unbiased, single-point mutations and develop a method to measure ruggedness based on spectral landscape theory [9], which decomposes landscapes into low-frequency (smooth) and high-frequency (rugged) components using a graph Fourier transform (GFT). Analogous to a standard Fourier transform, the GFT decomposes the original landscape $\boldsymbol{f}$ to express it as a mixture of basis landscapes of increasing frequency, such as those shown in Fig 2C. Here, $\boldsymbol{f}$ can be interpreted as mapping from genotype space $\mathcal{G}$ to fitness, $\boldsymbol{f}: \mathcal{G} \to \mathbb{R}$, or as a vector $\boldsymbol{f} \in \mathbb{R}^{|\mathcal{G}|}$, where $|\mathcal{G}|$ denotes the cardinality of $\mathcal{G}$. For a protein of length N and A choices of amino acids, there are $|\mathcal{G}| = A^N$ basis landscapes grouped into $N+1$ frequencies $\lambda_0, ..., \lambda_N$. The weights of each basis landscape in $\boldsymbol{f}$ can be arranged to form a coefficient vector $\hat{\boldsymbol{f}}$ in the Fourier domain, as shown in Fig 2D–2E. In the Fourier domain, genotype indices are replaced by Fourier-mode indices, and the corresponding values are the Fourier coefficients of the landscape.

Let F_μ denote the expected population fitness after an average of μ accumulated mutations, where the expectation is taken over mutation trajectories. We further define the expected squared population-average fitness $G_\mu \stackrel{\text{def}}{=} \mathbb{E}_\nu[F_\mu^2]$, where the expectation is additionally taken over all single-genotype populations. Based on the GFT, we propose to extract the scalar ruggedness metrics ρ_1 and ρ_2, associated with F_μ and G_μ, respectively, as the weighted frequency averages

$$\rho_1 = \frac{1}{d}\frac{\sum_{i=0}^{N} a_i\lambda_i}{\sum_{i=0}^{N} a_i - a_0} = \frac{1}{d}\frac{\boldsymbol{f}^\mathsf{T}\Delta\boldsymbol{\nu}}{\boldsymbol{f}^\mathsf{T}\boldsymbol{\nu} - \bar{f}}, \tag{1a}$$

$$\rho_2 = \frac{1}{d}\frac{\sum_{i=0}^{N} b_i\lambda_i}{\sum_{i=0}^{N} b_i - b_0} = \frac{1}{d}\frac{\boldsymbol{f}^\mathsf{T}\Delta\boldsymbol{f}}{\boldsymbol{f}^\mathsf{T}\boldsymbol{f} - A^N\bar{f}^2}. \tag{1b}$$

where $d = N(A-1)$ is determined by the landscape size and alphabet, a_i denotes the population-weighted spectral coefficient at frequency index i, b_i the population-independent power spectral coefficient at frequency index i, $\bar{f}$ the landscape average fitness, and λ_i a frequency-dependent quantity determined by the genotype-graph Laplacian $\Delta = \Delta^\mathsf{T}$ associated with the genotype graph induced by the unbiased mutation process (see Methods). The left-hand sides of Eqs (1a) and (1b) correspond to the centroids of the population-weighted spectrum and the population-independent power spectrum, respectively. The right-hand sides provide equivalent representations in vector form. In particular, Eq (1b) shows that the power-spectral centroid is equivalent to a normalised Dirichlet energy on the genotype graph, quantifying the average squared fitness variation between mutational neighbours. Consequently, the frequency-domain interpretation of our metrics is naturally linked to the unbiased mutation model, although we show in the Methods how the framework extends to biased mutation operators beyond the GFT. A key advantage of the metrics in Eq (1) is that they can be estimated directly from phenotypic measurements obtained from randomly mutagenised populations. In Methods, we show that for small μ, the expected population fitness F_μ, the average squared fitness G_μ, and the metrics from Eq (1) are related by:

$$F_\mu = C_1 e^{-\mu\rho_1} + \bar{f} + \mathcal{O}\left(\mu^2\right), \tag{2a}$$

$$G_\mu = C_2 e^{-2\mu\rho_2} + \bar{f}^2 + \mathcal{O}\left(\mu^2\right), \tag{2b}$$

where C_1 is a constant that depends on the initial population and can be negative if the initial population fitness is lower than $\bar{f}$, and C_2 is a population-independent constant with $C_2 \geq 0$ if G_μ is obtained from *all* single-genotype populations. The ruggedness metrics ρ_1 and ρ_2 therefore govern the rate at which the mean population fitness and its square relax towards the landscape average $\bar{f}$ and $\bar{f}^2$. In both cases, the decay rates can be interpreted as the ruggedness of the landscape by relating them to the different Fourier components from Fig 2: for large ρ_1 or ρ_2, higher-frequency terms dominate the corresponding decay, and for small ρ_1 or ρ_2, lower-frequency terms dominate. This observation is the phenomenon that our work exploits to estimate the landscape ruggedness through measurements of the mean fitness of a population as it accumulates mutations. Suppose that each population member is mutagenised with θ single-point mutations per generation on average, and the mean fitness is recorded as $F_{\mu_1}, F_{\mu_2}, \ldots, F_{\mu_M}$ with $\mu_i = \theta i$ over M generations. These values can be used to estimate the parameters C_1, ρ_1, and $\bar{f}$ in Eq (2a). Equivalently, by squaring the mean-fitness trajectory for each starting population and then averaging the squared trajectories, they can be used to estimate C_2, ρ_2, and $\bar{f}^2$ in Eq (2b). The analytical quantities ρ_1 and ρ_2 can be computed from Eq (1), whereas the fitted quantities are obtained by approximating finite observations with a single exponential decay. In the following, we denote decay rates fitted from finite, noisy observations by ρ_1^{fit} for F_μ and ρ_2^{fit} for G_μ. We further distinguish between local estimates ρ_2^{loc}, obtained from one or a small number of starting genotypes, and global estimates ρ_2^{glob}, obtained after averaging over all, or a large representative set of, starting genotypes.

In homogeneous landscapes, such as *NK* landscapes, local and global notions of ruggedness broadly coincide. In particular, the epistatic interaction parameter K defines the closed-form ruggedness value $\rho_{NK} \stackrel{\text{def}}{=} (K+1)/N$, and the analytical decay rates satisfy $\rho_1 \approx \rho_2 \approx \rho_{NK}$ (see Eq (S2) in Section D of S1 Appendix). Fig 3A illustrates the behaviour of these two quantities on *NK* landscapes. For F_μ, the exponential approximation is accurate when the starting population lies sufficiently far from the landscape average fitness, whereas populations starting close to the average exhibit little measurable decay and consequently less reliable fitted decay rates (top panel of Fig 3A). In general, ρ_1 represents a local notion of ruggedness associated with a particular starting population. By contrast, G_μ averages the squared fitness over starting populations, making the decay substantially more robust to the choice of starting genotype. In the bottom panel of Fig 3A, ρ_2^{fit} is computed from 25 starting populations and matches ρ_{NK}. This behaviour is confirmed in Fig 3B, which shows decay rates computed on a grid of *NK* landscapes for $N = 10, \ldots, 50$ and shows that ρ_2^{fit} and ρ_1^{fit} computed from high-fitness starting populations both correlate strongly with ρ_{NK}. Beyond *NK* landscapes, these metrics broadly coincide if the landscape is dominated by a single spectral frequency. In heterogeneous landscapes where more than one frequency dominates, however, local and global notions of ruggedness need not agree, and fitted estimates may differ from the corresponding analytical values.

Due to the noise associated with observing a randomly mutating population, the reliability of ρ_1^{fit} and ρ_2^{fit} estimates also depends on the parameters of the observed population, including the population size, the mutation rate θ, and for G_μ, the number of averaged starting points. Fig 3C–3E analyses the sensitivity of ρ_2^{fit} with respect to these parameters, where red vertical lines denote the fixed parameter values used when other parameters are varied, and dashed lines denote the analytical ρ_2. Fig 3C shows ρ_2^{fit} estimates over an increasing number of starting points. If fewer than 10 starting points are used, ρ_2^{fit} estimates can exhibit a large variance.

Fig 3D shows ρ_2^{fit} estimates over increasing population sizes. For small population sizes (≤ 20), the ruggedness estimate exhibits higher variability due to noise in the mean population fitness. The effect of mutation rate is analysed in Fig 3E; for small θ, only a small fraction of the fitness decay curve in Fig 3A is captured, leading to an increased variability of ρ_2^{fit}. For very large θ ($\theta \gg 2$), the mean population fitness decays more rapidly to the landscape average (comparable to the top panel of Fig 3A), again leading to a large variability of ρ_2^{fit} since only a few datapoints characterise the decay.

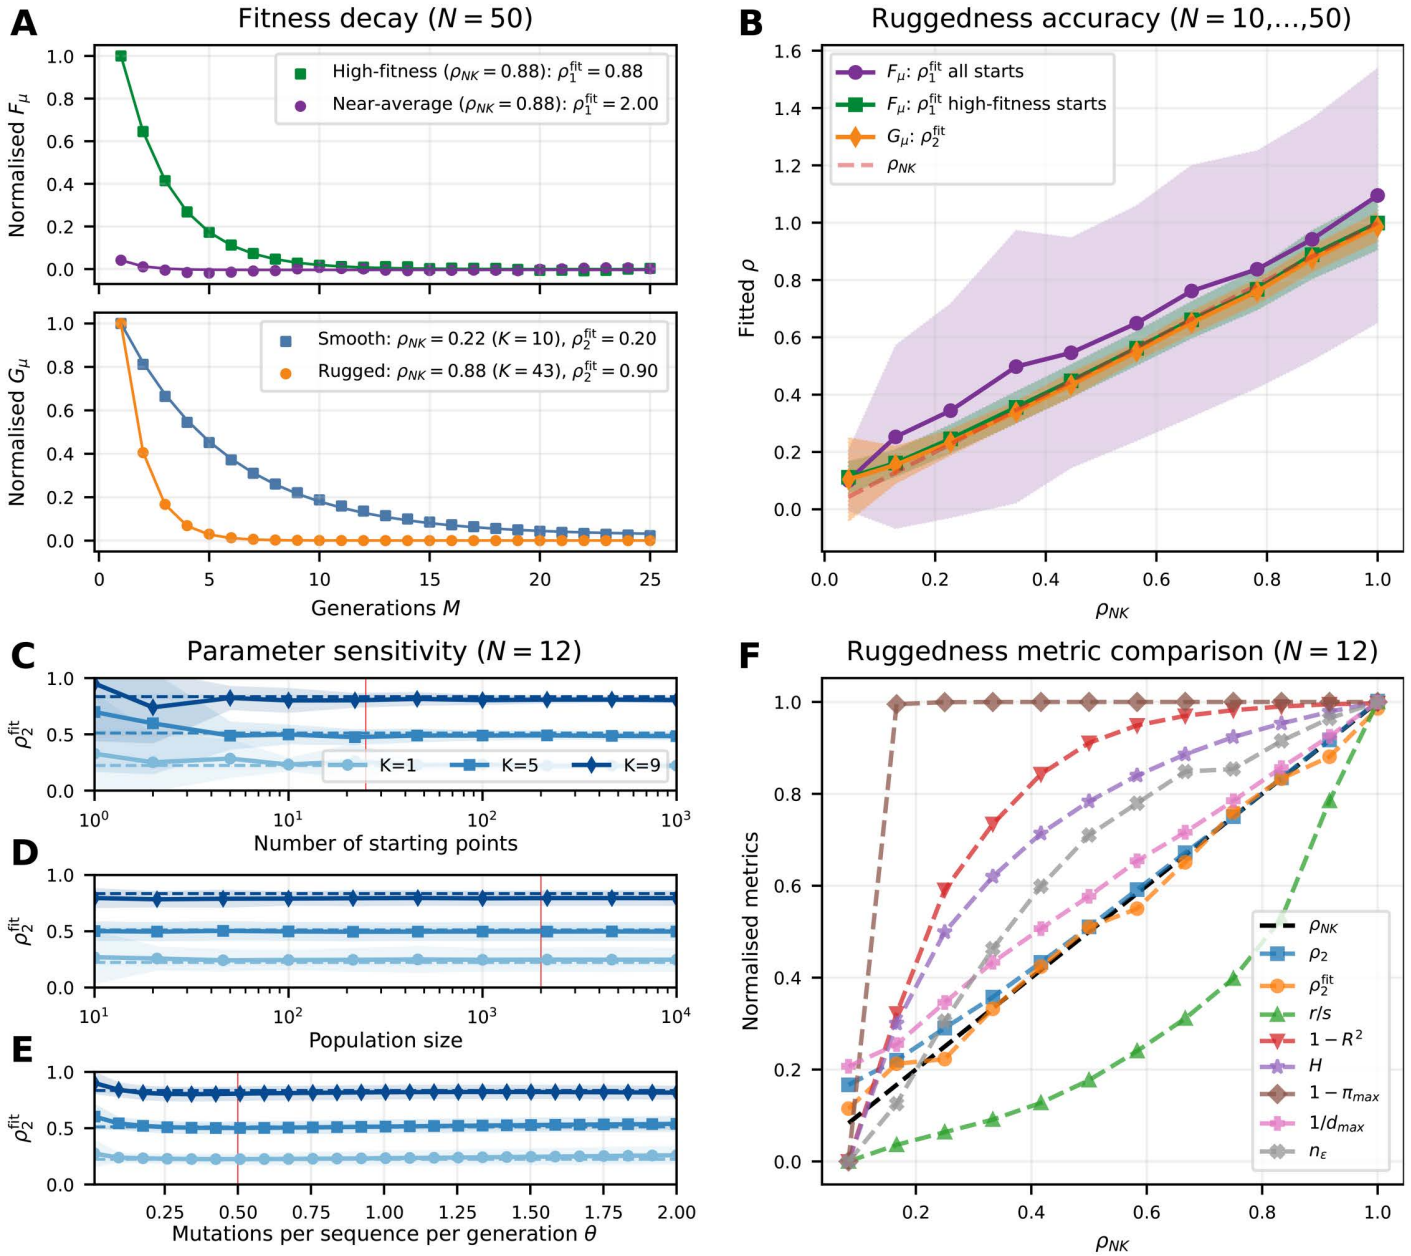


Fig 3. Fitness-decay estimates of landscape ruggedness on *NK* landscapes. A: Example decay trajectories on binary *NK* landscapes. The upper subpanel shows F_μ on a rugged landscape for one near-average-fitness start and one high-fitness start, selected from the starts closest to mean initial fitness and from the top 20% of initial fitness, respectively. The lower subpanel shows G_μ for a smoother landscape and the same rugged landscape; solid lines show exponential fits. **B:** Fitted decay rates across the *NK* grid ($N = 10, \ldots, 50$; 100 *NK* pairs; 25 landscapes per pair; 25 starts per landscape). For F_μ, ρ_1^{fit} is computed either from all starts or from high-fitness starts. For G_μ, ρ_2^{fit} is computed from all starts. Points and bands show bin means and empirical standard deviations (SDs) across estimates in each ρ_{NK} bin. **C–E:** Sensitivity of ρ_2^{fit} for $N = 12$, $K \in \{1, 5, 9\}$, and 25 landscapes per K; red vertical lines denote the fixed parameter values used when varying the other parameters; shaded bands show empirical SDs across landscapes. **F:** Comparison of fitted decay rates and analytical ruggedness metrics on $N = 12$, binary *NK* landscapes with $K = 0, \ldots, 11$. Analytical metrics are averaged over 50 landscapes per K; all ruggedness metrics are normalised for comparison. Unless otherwise stated, diffusion simulations used 25 starts per *NK* landscape, 5 replicate populations per start, mutation rate $\theta = 0.5$ (θ/N per site), and $M = 25$ generations. Replicate trajectories are averaged per start to obtain F_μ; these per-start trajectories are squared and then averaged over starts to obtain G_μ.

https://doi.org/10.1371/journal.pcbi.1014713.g003

In practice, an estimate of the average mutation rate per generation is required, which is assumed to be constant across generations (see Eq (3) in Methods). If the mutation rate per generation θ is not exactly known but estimated as $\tilde{\theta}$, ρ^{fit} will scale with $\theta/\tilde{\theta}$ as long as the mutation rate is constant across generations, e.g., if the actual θ is twice as high as $\tilde{\theta}$ then ρ^{fit} will be twice as large. As such, if the mutation rate error is constant across analysed landscapes, ρ_1^{fit} and ρ_2^{fit} values remain comparable, i.e., the ordering resulting from the metric is preserved under an estimation error.

To further validate the decay rates from Eq (1), Fig 3F provides a comparison between ρ_2 and other ruggedness measures across *NK* landscapes (see Section C of S1 Appendix). The other ruggedness measures are the roughness to slope ratio r/s [4], the landscape R^2 measure [8], the spectral entropy H, number of paths to the global maximum from the opposite position on the landscape π_{max} [7], the distance to the nearest local maximum d_{max} [6], and the local epistasis metric n_ϵ [8]. Fig 3F shows that although all metrics increase with ρ_{NK}, they express either concave (r/s), linear (decay rates and $1/d_{max}$), or convex curves (all other metrics). Because all metrics increase with K, this allows one to compare the ordering (w.r.t. ruggedness) of landscapes across metrics, but the differences in curve shapes prevent comparison of distances (i.e., quantitative ruggedness differences) between landscapes across metrics. However, one can mathematically prove that the decay rates directly correlate with ρ_{NK} (see Eq (S2) in Section D of S1 Appendix), and the same relationship holds regardless of N. Overall we believe this provides a second benefit of our methodology—that it provides an interpretable measure of ruggedness that has immediate meaning in the widely studied *NK* landscapes.

Ruggedness inference on empirical landscapes

The performance of SLIDE on empirical landscapes was assessed *in silico* by testing it on four experimentally measured landscapes: GB1 (4 sites) [2], ParD3 (3 sites) [24], TrpB (4 sites) [25] and TEV (4 sites) [26]. The Fourier spectra of each landscape are shown in Fig 4A, with the dashed vertical lines denoting the frequency centroids computed from the analytical ρ_2 as $\rho_2 \times d/A$. Here, in contrast to *NK* landscapes, the heterogeneity of empirical landscapes or—in other words—the spread of spectra makes it necessary to distinguish between ρ_2^{fit} estimated from one or a few starting points (denoted by ρ_2^{loc}) and ρ_2^{fit} estimated from all or a large set of starts (denoted by ρ_2^{glob}). Fig 4B compares ρ_2^{loc} obtained from single-genotype starting populations (shaded) with the analytical ρ_2 obtained from Eq (1b). On *NK* landscapes, these values broadly correspond, indicating that *NK* landscapes are effectively homogeneous. On empirical landscapes, ρ_2^{loc} and ρ_2 do not correspond. However, the starting population-dependent estimates may still be used to measure the ruggedness of an initial population $\boldsymbol{\nu}$. This is shown in Fig 4C, which shows the SD of ρ_2^{loc} estimates over 5 replicates with the fitness decay computed from a population comprising 10 starting points. Variability between estimates from the same starting populations is much smaller than in Fig 4B.

Practically, in order to obtain a global property from empirical landscapes despite their increased heterogeneity, one must sample multiple starting points. Despite the added labour required to measure multiple starting points, such a configuration remains scalable as it requires random mutagenesis only, something which is commonly conducted in a pooled format for library generation (e.g., using error-prone PCR). Fig 4D shows how the fitted decay rate ρ_2^{fit} varies as the number of starting points is increased. As it approaches the total number of starting points, the decay rate converges from ρ_2^{loc} (which may vary depending on starting populations) towards ρ_2^{glob}. The difference between ρ_2^{glob} and ρ_2 (dotted) depends on the "spread" of the spectrum; if the power is concentrated in a single frequency then the fitted global value coincides with Eq (1b). Generally, ρ_2^{glob} and ρ_2 are related via a first-order approximation of G_μ (see Methods). Fig 4E shows how ρ_2^{fit} varies with population size, and Fig 4F with number of generations sampled to obtain the fitness decay. In Fig 4F, the mutation rate is $\theta = 0.1$ per sequence, so that M generations correspond to $\mu = \theta M$ accumulated mutations. The principal determinant of whether an estimate approaches ρ_2^{glob} is the number of starting points sampled.

Fig 4G compares other ruggedness measures (described in Section C of S1 Appendix) across the empirical landscapes. Local epistasis (n_ϵ), distance to local maximum (d_{max}), and paths to global maximum (π_{max}) agree with the analytical metric ρ_2 in terms of their ordering of all the empirical landscapes tested. The landscape R^2 and roughness to slope

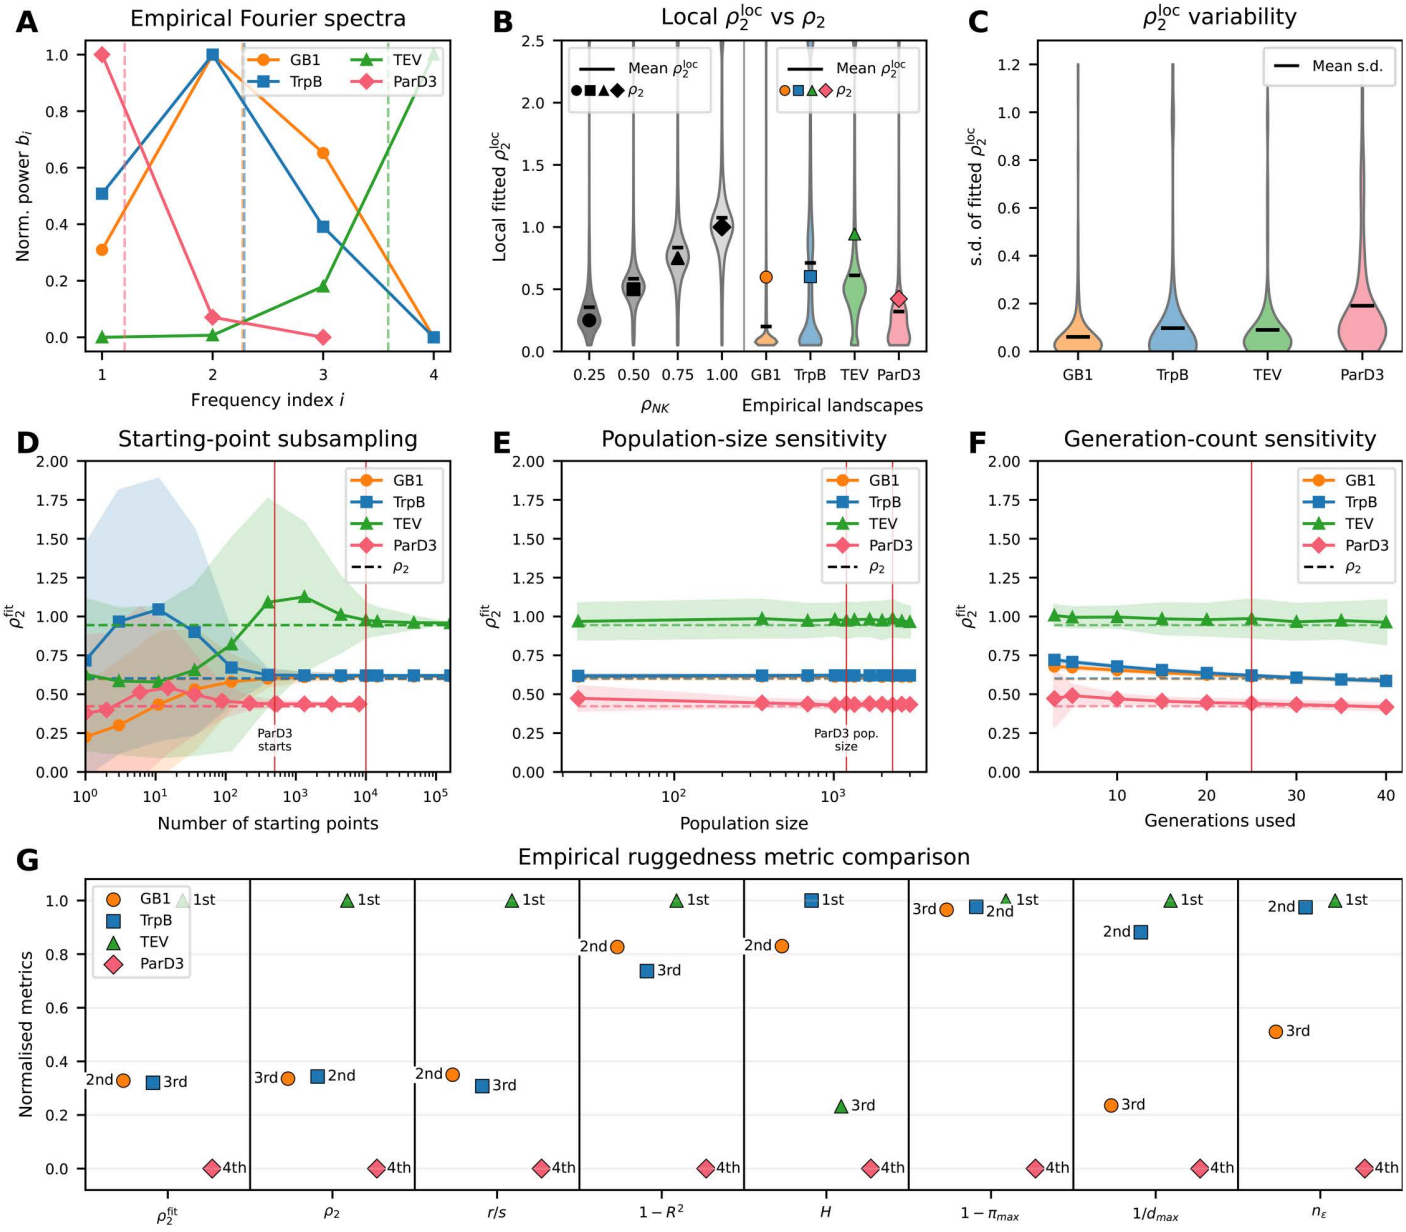


Fig 4. Ruggedness estimation on empirical landscapes. A: Fourier spectra of the complete GB1, TrpB, TEV and ParD3 landscapes computed by GFT; dashed lines indicate the frequency centroid $\rho_2 \times d/A$. **B:** Violin plots showing ρ_2^{loc} obtained by fitting the fitness decay from each individual starting genotype. Markers show the analytical ρ_2. **C:** Violin plots showing replicate variability in ρ_2^{loc}: for each landscape, estimates were repeatedly computed from bootstrap samples of 10 starts and 5 replicates per start (250 bootstrap samples). **D–F:** Sensitivity of ρ_2^{fit}; red vertical lines denote the fixed parameter values used when varying the other parameters; dashed horizontal lines denote ρ_2; shaded bands show SDs across bootstrap samples. **D:** Starting-point subsampling. Estimates were computed from increasing numbers of starting genotypes using 1000 bootstrap samples. **E:** Population-size sensitivity using 250 bootstrap samples per population size. **F:** Generation-count sensitivity using 250 bootstrap samples. **G:** Comparison of decay rates and other ruggedness metrics applied to the empirical landscapes; metrics are normalised for comparison. Unless otherwise stated, diffusion simulations used mutation rate $\theta = 0.1$ per sequence (θ/N per site), 10 replicate populations per start and 40 simulated generations. Global fits use the first 25 generations and, unless varied, 10,000 sampled starts for GB1, TrpB and TEV or 500 starts for ParD3.

https://doi.org/10.1371/journal.pcbi.1014713.g004

ratio (r/s) agree on the two extremes (TEV as the most rugged, ParD3 as the least), but place GB1 and TrpB in a different order. Similarly, ρ_2 groups together TrpB and GB1 landscapes as having a similar, intermediate ruggedness, whereas other methods separate these landscapes, or locate them very close to the TEV landscape. The general agreement between other ruggedness metrics and the ordering produced by ρ_2 on empirical landscapes provides increased confidence in the metric, which has the major advantage over all others of being experimentally measurable without detailed sequencing data (see Section E of S1 Appendix for a biological interpretation of these ruggedness values).

The analyses in Figs 3 and 4 assume uniform, unbiased mutation and are conducted directly on amino-acid landscapes, with $A = 20$ for the empirical landscapes. To assess the robustness of the approach under more realistic conditions, we further analysed ruggedness inference under biased mutation and controlled landscape perturbations, as described in the Methods and shown in Fig 6. The decay-based framework extends to weighted-undirected and directed mutation models, provided that the mutation process converges to a stationary distribution. Across the mutation models considered, the corresponding ruggedness estimates changed only modestly, indicating that the inferred ruggedness was largely robust to the considered mutation spectra. We additionally evaluated ρ_2 after introducing minimum-fitness regions, neutral ridges, and modular structure into *NK* landscapes. The metric responded systematically to each perturbation; for example, increasing the fraction of minimum-fitness genotypes increased the inferred ruggedness. However, because such perturbations generally introduce spatial heterogeneity, estimating a global ruggedness value may require fitness-decay measurements from a larger number of starting points.

Optimising DE outcomes with SLIDE

Inferring the ruggedness of a fitness landscape prior to DE allows informed decisions to be made about the DE strategy [2,3]. Here, as an illustrative application of SLIDE, we interface the decay-rate-based ruggedness estimate with the strategy framework developed in our previous work [10]. In that framework, combinations of two parameters—a fitness-independent selection probability (base chance) and population splitting—were evaluated on *NK* landscapes to determine the optimal strategy as a function of landscape ruggedness. We compare the resulting ruggedness-adapted strategy with two fixed-parameter strategies that can be represented within the same parameter space. The baseline strategy uses no population splitting and selects the top 20% of candidates. The high-exploration (HE) strategy uses high population splitting and a base chance of 20%, producing selection behaviour that approaches an almost random policy. These fixed strategies therefore represent the two extremes of exploitation and exploration, providing benchmarks for assessing the benefit of adapting this balance to the inferred ruggedness.

To explore the relationship between landscape ruggedness and optimal strategy, we evaluated 100 distinct (N,K) combinations at fixed population size and mutation rate (Sections A and B of S1 Appendix). For each combination, all 49 strategy combinations were evaluated for $M = 25$ DE generations on 200 independently generated *NK* landscapes, with 25 simulation replicates per strategy ($10 \leq N \leq 50$). Fig 5A shows the mean optimal base-chance and splitting values as a function of ρ_{NK}, with shaded regions indicating one SD across (N,K) combinations grouped by ρ_{NK}. Note that the optimal strategy is also dependent on the number of generations used in DE—more generations favour more explorative strategies as shown in S4 Fig. Ideally, the SLIDE look-up table should therefore be generated *in silico* using the anticipated number of generations, mutation rate, and population size of the experiment.

Because $\rho_1 \approx \rho_2 \approx \rho_{NK}$ on *NK* landscapes, these results provide a SLIDE look-up table indexed by ruggedness. A fitted estimate, ρ_1^{fit} or ρ_2^{fit}, can therefore be used to select an appropriate combination of DE parameters. Two examples of DE trajectories on *NK* landscapes are displayed in Fig 5B–5C. These figures additionally show how close each selected strategy is to the optimal strategy, the baseline strategy, and the HE strategy. Fig 5B displays a smooth landscape ($\rho_{NK} = 0.04$), where the strategy selected by SLIDE is highly similar to the baseline strategy (which is already near the optimum). Fig 5C, however, displays a more rugged landscape ($\rho_{NK} = 0.58$), on which the baseline strategy is highly sub-optimal and gets trapped at a local optimum within 10 generations. SLIDE, on the other hand, selects the optimal strategy within

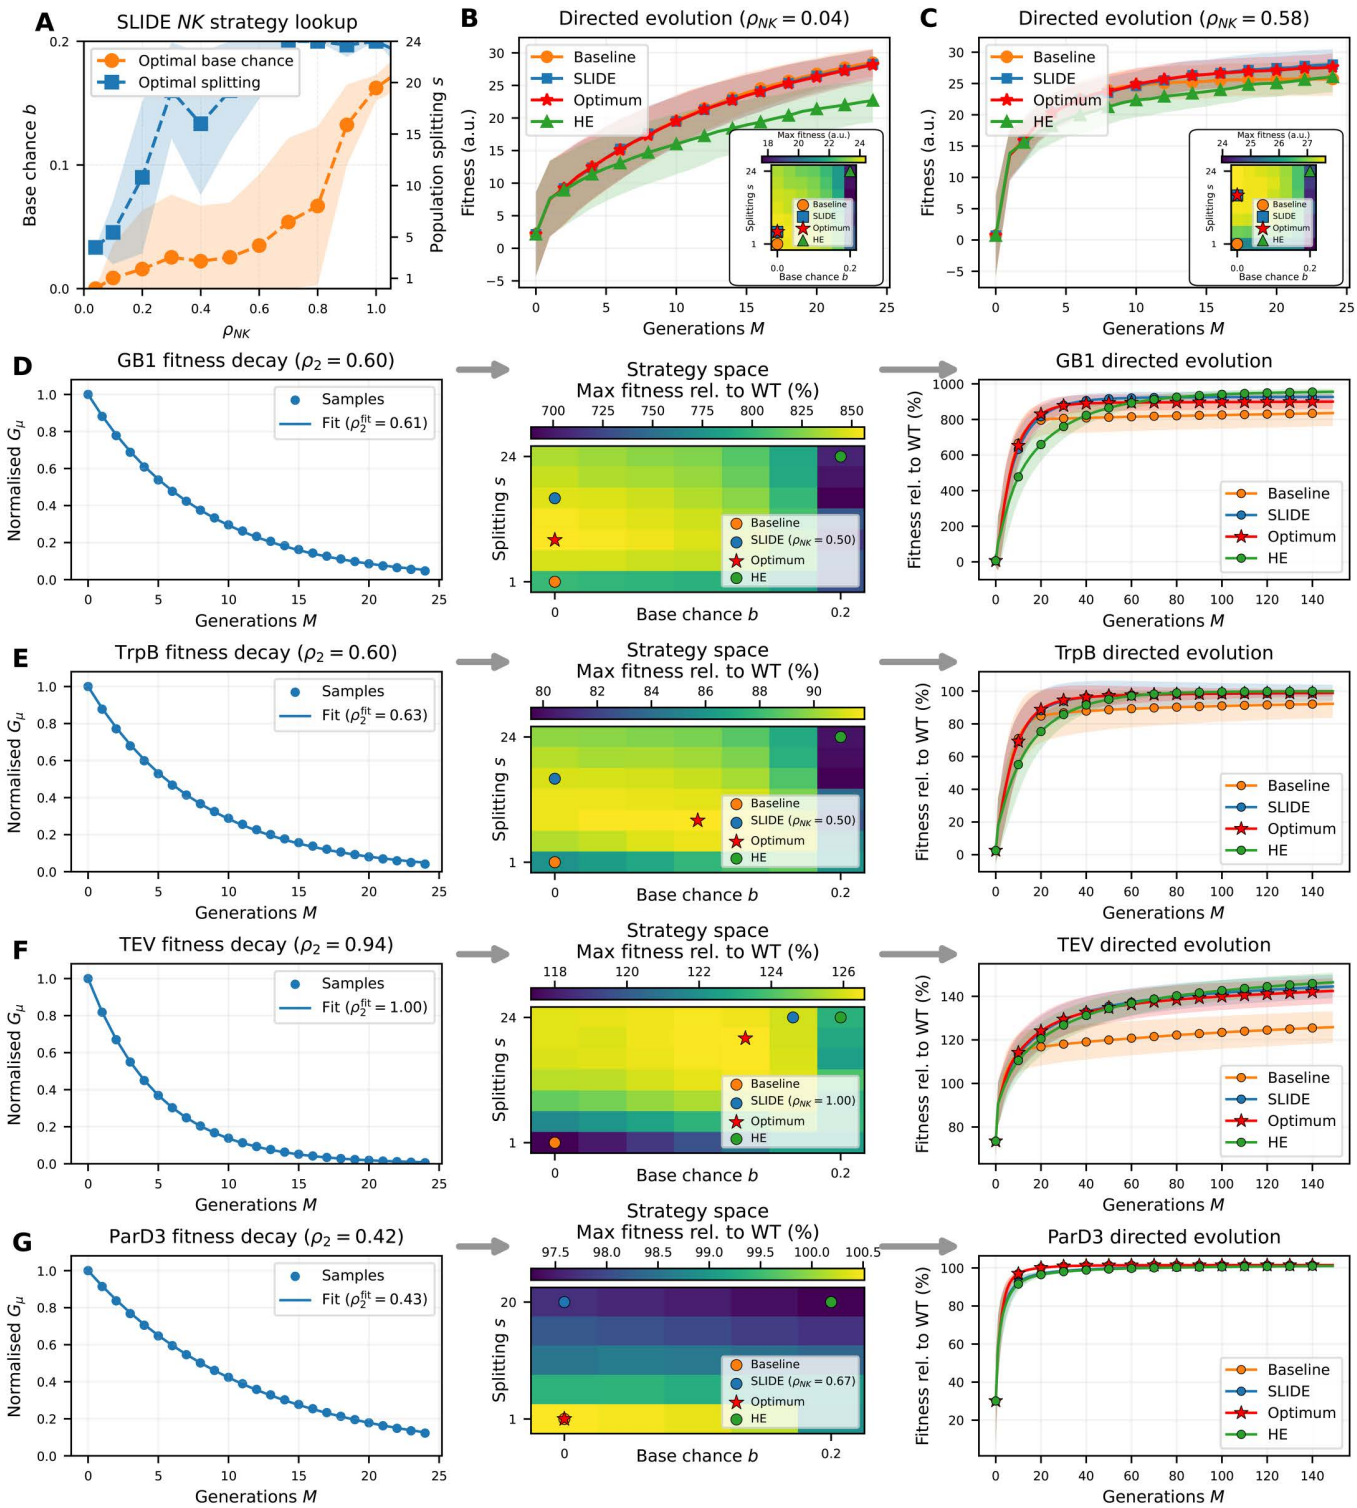


Fig 5. DE using SLIDE on *NK* and empirical landscapes. A: Representation of the SLIDE look-up table for *NK*. Generated from 100 (*N*,*K*) combinations with $10 \leq N \leq 50$. For each combination, all strategies were evaluated for $M = 25$ generations using 200 independently generated landscapes

and 25 replicates per strategy. Base chance options = {0.00, 0.03, 0.06, 0.10, 0.13, 0.16, 0.19}, splitting options = {24, 20, ..., 4, 1}. Lines show the mean optimal parameter values among (*N*,*K*) combinations grouped by ρ_{NK}; shaded regions show one SD across those combinations. **B–C:** SLIDE applied to smooth (*N* = 45,*K* = 1) and rugged (*N* = 45,*K* = 25) *NK* landscapes. Lines show mean fitness over 100 simulations on one landscape; shading indicates ±1 SD. Strategy spaces show mean fitness after 25 generations across 100 random starts and 300 replicates per start, with markers indicating the strategies selected by each approach. Population size = 1200; mutation rate $\theta = 0.1/N$ per site per generation. **D–G:** SLIDE applied to the GB1, TrpB, TEV and ParD3 empirical landscapes. Decay curves were estimated from 10,000 starting genotypes and 10 replicates ($\theta = 0.1/N$ per site; population size = 2500). Strategy spaces show mean fitness after 25 generations across 100 random starts and 300 replicates per start, computed on the corresponding empirical landscape; DE trajectory lines show the corresponding means over 150 generations, with shading indicating ±1 SD across starts. SLIDE strategies were selected using *A* = 20 *NK* look-up tables with *N* = 4. For **G**, *N* = 3, 8,000 decay starts, population size = 60, and splitting options {20,15,10,5,1} were used. HE, high exploration.

https://doi.org/10.1371/journal.pcbi.1014713.g005

the strategy space and achieves > 15 % higher fitness than the baseline strategy after 50 generations. The corresponding results in S1 Fig demonstrate the robustness of this trend across alternative landscape models other than *NK* (Rough Mount Fuji and a stochastic block model).

A similar approach was then applied to the four combinatorially complete empirical landscapes GB1, TrpB, TEV and ParD3 in Fig 5D–5G. The left-hand panels show the decay rate estimation from G_μ using 10,000 starting points. The middle panels evaluate all possible strategies on the corresponding empirical landscape and show the strategy selected via SLIDE; the SLIDE look-up table, as displayed in Fig 5A, was re-generated using *NK* landscapes of the same dimensions as each empirical landscape (*A* = 20, *N* = 4 or *N* = 3 for ParD3) and for *M* = 25 DE generations. SLIDE outperformed the baseline strategy in all examples, apart from ParD3, for which the baseline strategy was optimal (and therefore SLIDE selected the same parameters). After 150 generations on GB1, TrpB and TEV, it achieved relative fitness gains of 24.9% ($p < 0.001$), 22.6% ($p < 0.01$), and 44.2% ($p < 0.001$), respectively, over the baseline approach, which commonly gets trapped at local maxima. These improvements suggest that adapting the DE strategy to the inferred landscape structure can substantially improve optimisation performance. Consequently, integrating the ruggedness estimation framework proposed here with existing DE workflows may yield similar gains compared with fixed-parameter strategies.

Discussion

In this paper, we showed that in the absence of selection pressure, the average fitness decay of a population accumulating mutations can be approximated by an exponential curve whose fitted decay rate ρ_1^{fit} approximates the analytical, genotype-dependent metric ρ_1. The analogous squared-fitness decay rate ρ_2^{fit} approximates the genotype-independent metric ρ_2, which is linked to the normalised Dirichlet energy of the graph Laplacian and quantifies squared fitness variation across mutation neighbours. Under unbiased mutations, the decay rate ρ_2 reduces to the Fourier power-weighted average frequency of the Hamming graph Laplacian. Under biased mutation spectra, the same decay-based construction yields a metric associated with the effective mutation operator rather than the unweighted Hamming graph. Thus, if the mutational or experimental protocol induces non-uniform or selection-coupled dynamics—for example through codon-level mutation biases, viability filtering, repair biases, or growth-coupled sampling—the fitted decay rate should be interpreted as a protocol-dependent dynamical descriptor of the effective landscape explored by the population. For DE control, this effective ruggedness remains relevant, since it characterises the search process experienced under the protocol being used.

We validated the decay rate estimation *in silico* on synthetic *NK* landscapes and empirical datasets. Due to their homogeneity, the estimation procedure on *NK* landscapes is robust to the choice of mutation rate, population size, decay curve length, and number of starting genotypes, and the analytical, fitted, local, and global decay rates broadly correspond. The primary requirements are that the starting genotype is sufficiently far from the landscape average, and that the mutation rate is sufficiently large to induce observable decay, yet not so large that observed fitnesses lose temporal correlation.

We further tested SLIDE on systematically perturbed landscapes containing minimum-fitness regions, neutral ridges, and modular structure. As expected, these perturbations alter the Fourier spectrum and are reflected in the inferred ruggedness: minimum-fitness regions introduce sharp discontinuities and shift spectral power towards higher frequencies, whereas neutral ridges and modularity suppress higher-order variation and shift power towards lower frequencies. These results demonstrate that the inferred ruggedness captures the spectral consequences of diverse structural landscape features. However, because the metric provides a scalar summary of the spectrum, it does not by itself identify which particular landscape feature produced a given ruggedness value.

On more heterogeneous empirical landscapes, the ruggedness estimation requires more starting points to obtain a global estimate. However, representing the complete nucleotide genotype space through the nucleotide-to-amino-acid mapping spreads the landscape spectra over a broader range of frequencies, so that fewer starting genotypes are required to obtain a stable, global estimate at nucleotide level than at amino-acid level. Our simulations showed that for the considered empirical landscapes, estimating ruggedness at the amino-acid level required approximately 500–10,000 starting genotypes (0.25% to 6.25% of the landscape), whereas estimation at the nucleotide level required only 100 starting genotypes ($< 6 \times 10^{-4}$ % of the landscape). At the same time, however, the resulting ruggedness values are more closely spaced, so greater estimation accuracy is required to distinguish reliably between landscapes. By contrast, only 3–5 replicates were required to estimate a local ruggedness metric ρ_1^{fit} or ρ_2^{loc} accurately. Local metrics have not been analysed further here, but could be combined with a downstream DE selection strategy or used to measure the genetic stability of an individual genotype.

Using the decay rate as a ruggedness metric, we introduced SLIDE (Sequence-free Landscape Inference for Directed Evolution), a framework that estimates landscape ruggedness from average fitness decay curves. Conceptually, SLIDE aligns with broader work employing graph-based methods to extract structural properties of fitness landscapes [3,27–29]. Unlike existing methods, SLIDE requires only large-scale phenotypic data alongside a mutation rate estimation—removing the need for large-scale sequence-level information. As long as fitness can be coupled to a measurable phenotype such as fluorescence, SLIDE enables ruggedness estimation from bulk population-level measurements and knowledge of the average mutation rate. This feature makes SLIDE scalable and experimentally accessible.

By using SLIDE's ruggedness estimate to tune DE parameters, we demonstrated up to a 44.2% improvement in DE performance compared to the top-proportion selection baseline, even when ρ_2^{fit} was only within ± 25 % of ρ_2^{glob}. These gains arise because the inferred ruggedness enables the directed evolution strategy to be adapted to the underlying landscape structure, rather than relying on fixed mutation and selection parameters. While this paper applied DE strategies from [10], the ruggedness metric is broadly compatible with a wide range of advanced DE approaches, including AI-driven workflows. For example, in active learning-assisted DE [14], machine learning techniques are used to guide the selection, balancing exploration of unknown phenotypes with exploitation of high-performing variants. Here, the ruggedness metric obtained from SLIDE could be used to modulate the balance of exploration and exploitation. Similarly, in cluster learning-assisted DE [30], SLIDE could be used to tailor clustering granularity.

Although this work focused primarily on the theoretical and computational aspects of SLIDE, the framework is designed to be experimentally implementable. Local ruggedness estimation requires only a small number of replicate trajectories, whereas global estimation remains substantially more demanding because it must average over heterogeneity across the landscape. However, the number of required starting genotypes is significantly reduced at the nucleotide level, and even global SLIDE estimation requires only a fraction of the measurements needed to construct and sequence a combinatorially complete empirical landscape. The key requirement is a biological protocol for introducing mutations without applying selection pressure and with a mutation rate that is approximately constant across generations on average. One example is random *in vitro* mutagenesis, which allows targeted exploration of specific regions of the genotype space. Mutant libraries generated in this way can be expressed in cells, with fitness linked to a measurable phenotype such as fluorescence. By iteratively measuring average fluorescence across generations and estimating mutation accumulation,

or by simultaneously producing samples with increasing levels of mutation, it becomes possible to construct fitness-decay curves and apply SLIDE in a laboratory setting. Such an assay could be conducted in several 96-well plates, or at larger scales using automated liquid-handling systems, and could also be implemented in a pooled format using mother-machine microfluidics. Mother machines are microfluidic chips capable of long-term culture of up to 10^6 individual cells in spatially segregated chambers that eliminate growth-based competition [31,32]. In such a format, one need only introduce the original variant library into the device and observe the cells while inducing mutagenesis, using chemical mutagens, UV [33] or even targeted *in vivo* mutagenesis methods [34].

Several limitations remain. First, the present framework primarily considers fitness measurements that are decoupled from host or variant growth. This requirement could be relaxed by incorporating selection directly into the population dynamics. In general, selection can be represented as an additional operator acting on genotype frequencies—for example, through genotype-dependent growth or explicit enrichment rules—and combined with the mutation operator to yield mutation–selection dynamics. The observed fitness trajectories would then be governed by these combined dynamics, such that SLIDE estimates a ruggedness metric associated with the mutation–selection process rather than with a freely diffusing population. One approximation is to absorb selection into an effective biased mutation spectrum, as described in Methods. When selection induces only weak bias, the resulting metric remains close to the unbiased one; more generally, it quantifies ruggedness along the biased dynamics through which the landscape is actually traversed. This could allow the decay rate to be identified online and used as a feedback signal, either to update the ruggedness estimate or to adapt DE strategy parameters directly. For instance, under truncation selection, the early fitness gradient could be linked computationally and analytically to landscape ruggedness and structure (cf. Fig 5). Such extensions could enable SLIDE to operate *during* DE in settings where mutation and selection act simultaneously, including continuous *in vivo* mutagenesis [34] in small-scale bioreactors.

Second, future work should consider more general mutation operators. In particular, some mutation graphs may not admit a unique stationary distribution, for example when they are reducible or contain sink nodes. In such cases, the long-term fitness average and the inferred decay behaviour may depend on the initial population, and the present interpretation of the ruggedness metric would need to be extended accordingly.

Third, we analysed biased mutation, selection-related effects, and structural landscape perturbations largely in isolation. We did not systematically study their combined effects, for example when biased mutation or selection acts on landscapes containing neutral ridges. These interactions could alter both the effective spectrum and the spatial heterogeneity of the decay process, and therefore the relationship between local and global ruggedness estimates. A unified treatment of mutation, selection, and landscape structure is beyond the scope of the present work but represents an important direction for future study.

Finally, although the empirical landscapes provide combinatorially complete genotype–fitness maps, they cover only three- to four-site protein subspaces, and all validation presented here is computational. The present results therefore provide limited evidence for performance in larger, sparsely sampled, and structurally constrained protein spaces or in experimental DE campaigns. Addressing this limitation will require both experimental validation and theoretical landscape models that can be tuned to reproduce a broader range of real-world properties. Developing such models, potentially using generative approaches or other data-driven methods, represents an important direction for future research.

In summary, SLIDE offers a novel method for estimating the ruggedness of protein fitness landscapes using average phenotypic data alongside mutation rate estimation. These ruggedness estimates can inform experimental strategies in DE, helping to tailor protocols to the underlying structure of the landscape. Next to DE, SLIDE's ruggedness metric could be extended to characterise more complex landscapes, such as those underlying synthetic biological circuits. Large-scale mapping of ruggedness in such systems could provide a quantitative basis for designing circuits with greater genetic stability and functional robustness, complementing recent efforts to understand and engineer the evolutionary potential of synthetic constructs [35].

Methods

Decay rate metric

To show why the fitness decay rates ρ_1 and ρ_2 in Eq (2a) can be used as ruggedness metrics, we give a brief overview of spectral landscape theory. We will represent a population accumulating single-point mutations as a stochastic signal in a graph, modelling mutagenesis and genotype space. For a gene of length N, we denote the genotype space by $\mathcal{G} \stackrel{\text{def}}{=} \mathbb{Z}_A^N$ (the space of positive integers up to A, in a vector of N entries), containing $|\mathcal{G}| = A^N$ genotypes $\boldsymbol{x}_1, ..., \boldsymbol{x}_{A^N}$. Each genotype can be written as $\boldsymbol{x}_i = (\sigma_1^i, ..., \sigma_N^i)^\mathsf{T}$, where $\sigma_j^i \in \mathbb{Z}_A$ is the allele at locus j. For example, if $A = 20$ represents the number of possible amino acids at each position in a protein of length N, the landscape consists of A^N distinct genotypes. These genotypes can be represented as nodes in a graph, with each node being assigned a fitness value $f_i \stackrel{\text{def}}{=} f(\boldsymbol{x}_i)$, where $f : \mathcal{G} \mapsto \mathbb{R}$ characterises the fitness landscape. A population is represented by $\boldsymbol{\nu} = (\nu_1, ..., \nu_{A^N})^\mathsf{T} \in \mathbb{R}^{A^N}$, where each ν_i gives the frequency of genotype $\boldsymbol{x}_i$. Equivalently, $\boldsymbol{\nu}$ can be defined as a probability distribution, in which case $\sum_i \nu_i = 1$ with all $\nu_i \in [0, 1]$. The mean population fitness is then computed as $F = \boldsymbol{f}^\mathsf{T}\boldsymbol{\nu} = \sum_i f_i \nu_i$, where $\boldsymbol{f} = (f_1, ..., f_{A^N})^\mathsf{T} \in \mathbb{R}^{A^N}$ is the stacked vector of fitness values.

The edges of the graph connect genotypes that differ by a single mutation, so that every node is connected to $d \stackrel{\text{def}}{=} N(A-1)$ other nodes; such a graph is also referred to as a Hamming graph [36]. For node transitions, we define the single-mutation operator $\mathcal{D}$ to map an initial population $\boldsymbol{\nu}$ to $\mathcal{D}(\boldsymbol{\nu})$, where every population member has undergone a single, random and equally likely mutation (see Section A of S1 Appendix for details on our mutation model). It can be shown that $\mathcal{D}(\boldsymbol{\nu}) = D\boldsymbol{\nu}$, where $D = D^\mathsf{T} \in \mathbb{R}^{A^N \times A^N}$ is a row-stochastic matrix obtained from multiplying the adjacency matrix of the graph with $1/d$.

We assume that the number of mutations that occurs in one mutagenesis round is Poisson distributed with mean θ, so that the number of mutations X_M occurring between M rounds is Poisson distributed with mean $M\theta \stackrel{\text{def}}{=} \mu$. The population is therefore expected to evolve as

$$\mathbb{E}_{X_M}\left[\mathcal{D}^{X_M}(\boldsymbol{\nu})\right] = \left(\sum_{k=0}^{\infty} \Pr(X_M = k) D^k\right)\boldsymbol{\nu} = \left(\sum_{k=0}^{\infty} e^{-M\theta}\frac{(M\theta)^k}{k!}D^k\right)\boldsymbol{\nu} = \left(\sum_{k=0}^{\infty} e^{-\mu}\frac{\mu^k}{k!}D^k\right)\boldsymbol{\nu} \stackrel{\text{def}}{=} \mathcal{M}^M\boldsymbol{\nu}. \tag{3}$$

Our mutation operator $\mathcal{M}^M$ can be understood as the transition function of a continuous-time Markov chain with a single genotype randomly mutating at rate θ. It can be simplified using matrix exponentials as

$$\mathcal{M}^M = \sum_{k=0}^{\infty} e^{-\mu}\frac{\mu^k}{k!}D^k = e^{-\mu}\sum_{k=0}^{\infty}\frac{(\mu D)^k}{k!} = e^{-\mu}e^{\mu D} = e^{-\mu(I-D)} \stackrel{\text{def}}{=} e^{-\frac{\mu}{d}\Delta}, \tag{4}$$

where we substituted the graph Laplacian $\Delta = \Delta^\mathsf{T} \stackrel{\text{def}}{=} d(I - D)$. The mean fitness F_μ of our population after an average of $\mu = M\theta$ mutations is therefore

$$F_\mu = \boldsymbol{f}^\mathsf{T}\left(\mathcal{M}^M\boldsymbol{\nu}\right) = \boldsymbol{f}^\mathsf{T}e^{-\frac{\mu}{d}\Delta}\boldsymbol{\nu}. \tag{5}$$

If necessary, the dependency on the initial population $\boldsymbol{\nu}$ can be omitted by taking the expectation of Eq (5) over many random starting populations. In particular, if $\boldsymbol{\nu}$ is sampled uniformly from single-genotype populations, so that $\mathbb{E}[\boldsymbol{\nu}\boldsymbol{\nu}^\mathsf{T}] = I/A^N$, we can compute $\mathbb{E}_\nu[F_\mu^2]$ as:

$$G_\mu \stackrel{\text{def}}{=} \mathbb{E}_\nu[F_\mu^2] = \boldsymbol{f}^\mathsf{T}e^{-\frac{\mu}{d}\Delta}\mathbb{E}_\nu[\boldsymbol{\nu}\boldsymbol{\nu}^\mathsf{T}]e^{-\frac{\mu}{d}\Delta^\mathsf{T}}\boldsymbol{f} = \frac{1}{A^N}\boldsymbol{f}^\mathsf{T}e^{-\frac{2\mu}{d}\Delta}\boldsymbol{f}. \tag{6}$$

If $\boldsymbol{\nu}$ is sampled from a smaller subset with $\mathbb{E}[\boldsymbol{\nu}\boldsymbol{\nu}^{\mathsf{T}}] = \Sigma_\nu$, $G_\mu = \boldsymbol{f}^{\mathsf{T}} e^{-\frac{\mu}{d}\Delta}\Sigma_\nu e^{-\frac{\mu}{d}\Delta^{\mathsf{T}}}\boldsymbol{f}$. Eqs (5) and (6) can be further simplified by considering the special structure of a Hamming graph. Its real-valued, symmetric Laplacian admits a unitary Fourier basis as an eigenbasis [37], so that Δ can be factorised as $\Delta = P\Lambda P^*$, where $P \in \mathbb{C}^{A^N \times A^N}$ satisfying $P^*P = I$ contains the eigenvectors, and $\Lambda = \mathrm{diag}(\lambda_1, ..., \lambda_{A^N}) \in \mathbb{R}^{A^N \times A^N}$, with the eigenvalues ordered as $\lambda_1 \leq \lambda_2 \leq ... \leq \lambda_{A^N}$. Analogous to the standard Fourier transform, the eigenvalues capture a frequency, and the columns of P are eigenfunctions, each of which can be interpreted as eigenlandscapes of increasing ruggedness for increasing λ_i. For the Hamming graph Laplacian, there are $N+1$ unique eigenvalues with value $\lambda_i = A \times i$ and multiplicity $\binom{N}{i}(A-1)^i$, $i = 0, 1, ..., N$ [38]. The matrix P may be chosen as $P_{j,i} \stackrel{\text{def}}{=} \chi_j(x_i)/\sqrt{A^N}$, where $\chi_j(x_i) \stackrel{\text{def}}{=} \exp(2\pi \mathrm{i}\langle r_j, x_i\rangle/A)$ is the Fourier basis function, $r_j \in \mathcal{G}$ the frequency index, and $\langle r_j, x_i\rangle$ the dot product defined on $\mathcal{G}$ (dot product modulo A) [37]. For $\lambda_0 = 0$, the eigenlandscape is a constant equal to the landscape average $\bar{f} \stackrel{\text{def}}{=} \mathbf{1}^{\mathsf{T}}\boldsymbol{f}/A^N$, where $\mathbf{1} \stackrel{\text{def}}{=} (1, ..., 1)^{\mathsf{T}}$. For $\lambda_1 = A$, the landscape is linear, i.e., depends only on the allele σ_j at a single locus j (Fig 2C). For larger eigenvalues, the eigenlandscapes are rugged, i.e., they depend on the allele at multiple loci (exactly i different loci for $\lambda_i = Ai$). From these eigenvalue-eigenvector pairs, it also follows that $\lim_{\mu\to\infty} F_\mu = \bar{f}$ and $\lim_{\mu\to\infty} G_\mu = \boldsymbol{f}^{\mathsf{T}}\mathbf{1}\mathbf{1}^{\mathsf{T}}\boldsymbol{f}/(A^N)^2 = \bar{f}^2$. In other words, if we sample uniformly from all single-genotype populations and average the fitness trajectories via $\mathbb{E}_\nu[F_\mu^2]$, the resulting G_μ decays from $G_0 = \boldsymbol{f}^{\mathsf{T}}\boldsymbol{f}/A^N$ exponentially to $\bar{f}^2$ as $\mu \to \infty$.

The special eigenbasis of the Hamming graph allows the mean fitness decay to be broken down into contributions of different frequencies. Substituting $\Delta = P\Lambda P^*$ in Eqs (5) and (6) and defining $\hat{\boldsymbol{f}} \stackrel{\text{def}}{=} P^*\boldsymbol{f}$ and $\hat{\boldsymbol{\nu}} \stackrel{\text{def}}{=} P^*\boldsymbol{\nu}$ yields:

$$F_\mu = \hat{\boldsymbol{f}}^* e^{-\frac{\mu}{d}\Lambda}\hat{\boldsymbol{\nu}} = \sum_{i=0}^{N} e^{-\frac{\mu\lambda_i}{d}}\Big(\sum_{j\in\mathcal{I}_i}\hat{\boldsymbol{f}}_j^*\hat{\boldsymbol{\nu}}_j\Big) \stackrel{\text{def}}{=} \sum_{i=0}^{N} e^{-\frac{\mu\lambda_i}{d}} a_i, \tag{7a}$$

$$G_\mu = \frac{1}{A^N}\hat{\boldsymbol{f}}^* e^{-\frac{2\mu}{d}\Lambda}\hat{\boldsymbol{f}} = \frac{1}{A^N}\sum_{i=0}^{N} e^{-\frac{2\mu\lambda_i}{d}}\Big(\sum_{j\in\mathcal{I}_i}|\hat{\boldsymbol{f}}_j|^2\Big) \stackrel{\text{def}}{=} \frac{1}{A^N}\sum_{i=0}^{N} e^{-\frac{2\mu\lambda_i}{d}} b_i, \tag{7b}$$

where $\mathcal{I}_i$ ($|\mathcal{I}_i| = \binom{N}{i}(A-1)^i$) groups indices associated with the ith eigenvalue and

$$a_i \stackrel{\text{def}}{=} \sum_{j\in\mathcal{I}_i}\hat{\boldsymbol{f}}_j^*\hat{\boldsymbol{\nu}}_j,\ b_i \stackrel{\text{def}}{=} \sum_{j\in\mathcal{I}_i}|\hat{\boldsymbol{f}}_j|^2. \tag{8}$$

Because $\lambda_0 = 0$ and $\lambda_i > 0$ for $i \geq 1$, it holds that

$$\lim_{\mu\to\infty} F_\mu = a_0 = \frac{(\mathbf{1}^{\mathsf{T}}\boldsymbol{\nu})(\mathbf{1}^{\mathsf{T}}\boldsymbol{f})}{A^N} = \bar{f},\ \lim_{\mu\to\infty} G_\mu = \frac{b_0}{A^N} = \frac{(\mathbf{1}^{\mathsf{T}}\boldsymbol{f})^2}{(A^N)^2} = \bar{f}^2. \tag{9}$$

Thus $b_0 = A^N\bar{f}^2$ under the Fourier normalisation used here. Expressions written in Fourier-power units subtract b_0, while equivalent expressions written in genotype-vector units subtract $A^N\bar{f}^2$. The operation $P^*\boldsymbol{f}$ can be understood as a coordinate transformation and the vector $\hat{\boldsymbol{f}}$ as the landscape $\boldsymbol{f}$ represented in the Fourier domain. The coefficients $b_0, ..., b_N$ are obtained from squaring the modulus of Fourier coefficients and therefore represent the power or energy per frequency λ_i. The power spectra of several NK landscapes are shown in Fig 7A. For the less rugged landscapes, the power spectrum peaks at a lower frequency than for the more rugged landscapes.

Eqs (7a) and (7b) show that as the accumulated mutations μ increase, the contribution of each Fourier component to the measured fitness decreases at an exponential rate proportional to λ_i, as shown in Fig 7B. For example, for a linear, zero-mean landscape for which only $b_1 \neq 0$, the mean fitness F_μ will decay at a rate proportional to $\lambda_1 = A$ per accumulated mutation, whereas for a maximally rugged landscape for which only $b_N \neq 0$, it will decay at a rate proportional to

$\lambda_N = AN$, which is N times faster than the linear one. For most landscapes, the decay will be composed of multiple exponentials, such as for the *NK* landscapes in Fig 7B–7C.

Although the coefficients $a_0, ..., a_N$ and $b_0, ..., b_N$ could be estimated from Eqs (7a) and (7b), their estimates can become inaccurate (see following section) if F_μ and G_μ are noisy. Assuming that the coefficients a_i and b_i from Eqs (7a) and (7b) vary significantly in magnitude, the fitness decay curve is dominated by the larger-magnitude exponentials. In the case of a single dominant component, Eqs (5) and (6) can be approximated by $F_\mu \approx a_0 + e^{-\frac{\mu\lambda_{i^*}}{d}} a_{i^*}$ and $G_\mu \approx (b_0 + e^{-\frac{2\mu\lambda_{i^*}}{d}} b_{i^*})/A^N$, where i^* is the index of the dominant component. We therefore seek to approximate Eq (7a) and (7b) using a single exponential as

$$F_\mu \approx C_1 e^{-\mu\rho_1} + c_1, \tag{10a}$$

$$G_\mu \approx C_2 e^{-2\mu\rho_2} + c_2, \tag{10b}$$

where from Eq (9), it holds that $c_1 = \bar{f}$ and $c_2 = \bar{f}^2$. We introduce the decay rates ρ_1 and ρ_2 as scalar ruggedness metrics that approximate the average of the frequencies λ_i weighted by a_i or b_i. This can be shown by expanding the right-hand sides of Eqs (10a) and (10b) as

$$C_1 e^{-\mu\rho_1} + c_1 = C_1(1 - \mu\rho_1 + \frac{1}{2}\mu^2\rho_1^2 + ...) + c_1 = C_1 + c_1 - C_1\mu\rho_1 + \mathcal{O}(\mu^2), \tag{11a}$$

$$C_2 e^{-2\mu\rho_2} + c_2 = C_2(1 - 2\mu\rho_2 + 2\mu^2\rho_2^2 + ...) + c_2 = C_2 + c_2 - 2C_2\mu\rho_2 + \mathcal{O}(\mu^2), \tag{11b}$$

and the right-hand sides of Eqs (7a) and (7b) as

$$\sum_{i=0}^{N} e^{-\frac{\mu\lambda_i}{d}} a_i = \sum_{i=0}^{N} a_i - \frac{\mu}{d}\sum_{i=0}^{N} a_i\lambda_i + \mathcal{O}(\mu^2), \tag{12a}$$

$$\frac{1}{A^N}\sum_{i=0}^{N} e^{-\frac{2\mu\lambda_i}{d}} b_i = \frac{1}{A^N}\sum_{i=0}^{N} b_i - \frac{1}{A^N}\frac{2\mu}{d}\sum_{i=0}^{N} b_i\lambda_i + \mathcal{O}(\mu^2) \tag{12b}$$

Ignoring higher-order terms, comparing coefficients yields $C_1 + c_1 = \sum_{i=0}^{N} a_i$ and $C_1\mu\rho_1 = \mu/d\sum_{i=0}^{N} a_i\lambda_i$ for F_μ, and $C_2 + c_2 = \sum_{i=0}^{N} b_i/A^N$ and $C_2\mu\rho_2 = \mu/d\sum_{i=0}^{N} b_i\lambda_i/A^N$ for G_μ. We also note that since $\lambda_0 = 0$, Eq (7) and Eqs (9)–(10) imply that $c_1 = a_0 = \bar{f}$ for F_μ and $c_2 = b_0/A^N = \bar{f}^2$ for G_μ, so that

$$\rho_1 \approx \frac{1}{d}\frac{\sum_{i=0}^{N} a_i\lambda_i}{\sum_{i=0}^{N} a_i - a_0} = \frac{1}{d}\frac{\hat{\boldsymbol{f}}^* \Lambda \hat{\boldsymbol{\nu}}}{\hat{\boldsymbol{f}}^* \hat{\boldsymbol{\nu}} - a_0} = \frac{1}{d}\frac{\boldsymbol{f}^\mathsf{T}\Delta\boldsymbol{\nu}}{\boldsymbol{f}^\mathsf{T}\boldsymbol{\nu} - \bar{f}}, \tag{13a}$$

$$\rho_2 \approx \frac{1}{d}\frac{\sum_{i=0}^{N} b_i\lambda_i}{\sum_{i=0}^{N} b_i - b_0} = \frac{1}{d}\frac{\hat{\boldsymbol{f}}^* \Lambda \hat{\boldsymbol{f}}}{\hat{\boldsymbol{f}}^* \hat{\boldsymbol{f}} - b_0} = \frac{1}{d}\frac{\boldsymbol{f}^\mathsf{T}\Delta\boldsymbol{f}}{\boldsymbol{f}^\mathsf{T}\boldsymbol{f} - A^N\bar{f}^2}. \tag{13b}$$

The right-hand side of Eq (13b) highlights that the analytical decay rate ρ_2 for G_μ is the Dirichlet energy $\boldsymbol{f}^\mathsf{T}\Delta\boldsymbol{f}$ normalised by $\boldsymbol{f}^\mathsf{T}\boldsymbol{f} - A^N\bar{f}^2$, equivalently $\boldsymbol{f}^\mathsf{T}\boldsymbol{f} - b_0$. Eq (13b) also shows that ρ_2 is related to other spectral measures such as the

R^2 measure, $R^2 \stackrel{\text{def}}{=} \sum_{i=2}^{N} b_i/(\sum_{j=1}^{N} b_j)$, or the spectral entropy, $H \stackrel{\text{def}}{=} -\sum_{i=1}^{N} p_i \log(p_i)$, where $p_i \stackrel{\text{def}}{=} b_i/(\sum_{j=1}^{N} b_j)$. However, note that Eq (13b) explicitly takes the mutation spectrum into account, whereas other measures do not.

From the first-order approximation in Eq (11), the analytical metric ρ_2 characterises the G_μ decay for small μ. Additionally, Jensen's inequality for convex functions can be used to show that

$$G_\mu = \frac{b_0}{A^N} + \frac{B}{A^N} \sum_{i=1}^{N} \frac{b_i}{B} e^{-\frac{2\mu\lambda_i}{d}} \geq \frac{b_0}{A^N} + \frac{B}{A^N} e^{-2\mu\rho_2}, \tag{14}$$

where $B = \sum_{i=1}^{N} b_i$, i.e., the single exponential decaying at rate ρ_2 provides a lower bound.

We note that ρ_1 for F_μ depends on the initial population $\boldsymbol{\nu}$, whereas for a sufficiently large number of samples in Eq (6), ρ_2 for G_μ reflects a global property of the landscape if computed from all starting points. We therefore distinguish between analytical decay rates (ρ_1, ρ_2), fitted decay rates (ρ_1^{fit}, ρ_2^{fit}), and—for ρ_2—its local (ρ_2^{loc}) and global (ρ_2^{glob}) fitted variants. A summary of the notation used is provided in S1 Table, including the decay rates for biased mutation spectra introduced in the following section. If a single component $a_{i^\star}$ or $b_{i^\star}$ is dominating, then the corresponding decay rates satisfy approximately

$$\rho_1 \approx \rho_2 \approx \frac{\lambda_{i^\star}}{d} \leq \frac{A}{A-1}, \tag{15}$$

with fitted decay rates being approximately equal to the one in Eq (15).

The approximation from Eqs (10a) and (10b) can be analysed in more detail for house-of-cards and *NK* landscapes (see Section D of S1 Appendix), where it can be proved that the approximations hold increasingly well in the limit $N \to \infty$ [39].

Biased mutation spectra and landscape perturbations

The single-mutation operator D allows the fitness decay to be related to the Hamming graph Laplacian, which is amenable to a Fourier decomposition. In this scenario, all mutations are equally likely, and each population member evolves according to an unbiased random walk on a regular, undirected graph. However, in practice, the mutation spectra are biased, e.g., due to transition/transversion preferences, context effects, or codon usage. Also natural selection, e.g., due to overexpressed proteins, can be modelled as mutation bias. These effects can result in biased walks on weighted, directed or undirected graphs, whose Laplacians are *not* amenable to a Fourier decomposition. Separately, structural perturbations of the fitness landscape, such as the introduction of minimum-fitness genotypes, neutral ridges or modular structures, alter the landscape Fourier spectrum even when the underlying mutation process remains unbiased. We therefore examine both mutation bias and controlled landscape perturbations in the following.

Biased mutations

Suppose that the graph is weighted but undirected with mutation operator $\tilde{D}$ and symmetric Laplacian $\tilde{\Delta}$, which occurs when, e.g., all transition mutations are preferred over transversion mutations. Since $\tilde{\Delta} = \tilde{\Delta}^\mathsf{T}$, the equations of the previous section remain valid. In particular, the matrix identity from Eq (6), $\exp(-\frac{\mu}{d}\tilde{\Delta})\exp(-\frac{\mu}{d}\tilde{\Delta}^\mathsf{T}) \equiv \exp(-\frac{2\mu}{d}\tilde{\Delta})$ holds, which requires $\tilde{\Delta}$ and $\tilde{\Delta}^\mathsf{T}$ to commute. However, the analytical squared-decay metric $\tilde{\rho}_2$ approximates the normalised Dirichlet energy on the weighted graph instead:

$$\tilde{\rho}_2 \approx \frac{1}{d} \frac{\boldsymbol{f}^\mathsf{T} \tilde{\Delta} \boldsymbol{f}}{\boldsymbol{f}^\mathsf{T} \boldsymbol{f} - \tilde{b}_0}, \tag{16}$$

where $\tilde{b}_0 = b_0 = A^N\bar{f}^2$. Since $\tilde{D}$ is symmetric, row-stochastic, and irreducible, the population converges to the uniform distribution, and $\tilde{G}_\mu$ to the squared uniform landscape average. The metrics ρ_2 and $\tilde{\rho}_2$ are associated with particular mutation operators and may differ in general. A landscape may appear rugged via ρ_2 and smooth via $\tilde{\rho}_2$. Letting $E_{\bar{\Delta}} \overset{\text{def}}{=} \Delta - \tilde{\Delta} = d(\tilde{D} - D)$, the difference $|\rho_2 - \tilde{\rho}_2|$ can be bounded as

$$|\rho_2 - \tilde{\rho}_2| = \frac{1}{d}\frac{|\boldsymbol{f}^\mathsf{T} E_{\bar{\Delta}} \boldsymbol{f}|}{|\boldsymbol{f}^\mathsf{T}\boldsymbol{f} - b_0|} \leq \frac{1}{d}\frac{\|\boldsymbol{f}\|_2^2 \|E_{\bar{\Delta}}\|_2}{|\boldsymbol{f}^\mathsf{T}\boldsymbol{f} - b_0|}. \tag{17}$$

Thus the metrics are close when the mutation kernel perturbation $\|D - \tilde{D}\|_2$ is small.

For a directed graph, let the non-symmetric mutation operator and Laplacian be $\bar{D}$ and $\bar{\Delta}$, respectively. Assume that $\bar{D}$ is row-stochastic and irreducible, so that $\bar{D}\mathbf{1} = \mathbf{1}$, and assume that $\bar{\Delta} = \bar{P}\bar{\Lambda}\bar{P}^{-1}$ is diagonalisable with a simple zero eigenvalue and all other eigenvalues having positive real part. This represents a more general configuration, which can occur, e.g., under selection pressure or using *in vivo* mutagenesis methods [40]. In this configuration, the population converges exponentially to its stationary distribution, and $\bar{G}_\mu$ converges to the square of the stationary-weighted fitness average:

$$\bar{G}_\infty \overset{\text{def}}{=} \lim_{\mu\to\infty} \bar{G}_\mu = \frac{1}{A^N}\boldsymbol{f}^\mathsf{T}\bar{\boldsymbol{d}}\mathbf{1}^\mathsf{T}(\bar{\boldsymbol{d}}\mathbf{1}^\mathsf{T})^\mathsf{T}\boldsymbol{f} = (\boldsymbol{f}^\mathsf{T}\bar{\boldsymbol{d}})^2, \tag{18}$$

where $\bar{\boldsymbol{d}}^\mathsf{T}\bar{D} = \bar{\boldsymbol{d}}^\mathsf{T}$ with $\bar{\boldsymbol{d}}^\mathsf{T}\mathbf{1} = 1$ is the stationary distribution induced by the mutation kernel. Without commutativity of $\bar{\Delta}$ and $\bar{\Delta}^\mathsf{T}$, the expression $\exp(-\frac{\mu}{d}\bar{\Delta}^\mathsf{T})\exp(-\frac{\mu}{d}\bar{\Delta})$ cannot be simplified. Instead, consider the first-order approximation $\exp(-\frac{\mu}{d}\bar{\Delta}) \approx I - \frac{\mu}{d}\bar{\Delta}$, which, when substituted in Eq (6), yields

$$\bar{G}_\mu \approx \frac{1}{A^N}\boldsymbol{f}^\mathsf{T}\left(I - \frac{2\mu}{d}\frac{\bar{\Delta} + \bar{\Delta}^\mathsf{T}}{2}\right)\boldsymbol{f} + \mathcal{O}(\mu^2), \tag{19}$$

where $\bar{\Delta}_{\text{sym}} \overset{\text{def}}{=} (\bar{\Delta} + \bar{\Delta}^\mathsf{T})/2$ is the orthogonal projection (via Frobenius norm) of $\bar{\Delta}$ onto the subspace of symmetric matrices. Comparing coefficients with Eq (11) results in the approximation

$$\bar{\rho}_2 \approx \frac{1}{d}\frac{\boldsymbol{f}^\mathsf{T}\bar{\Delta}_{\text{sym}}\boldsymbol{f}}{\boldsymbol{f}^\mathsf{T}\boldsymbol{f} - \bar{b}_0}, \tag{20}$$

where $\bar{b}_0 = A^N(\boldsymbol{f}^\mathsf{T}\bar{\boldsymbol{d}})^2$ is the weighted, squared landscape average. As for the weighted, undirected setting, the difference $|\rho_2 - \bar{\rho}_2|$ may be upper-bounded using the perturbation $E_{\bar{\Delta}} \overset{\text{def}}{=} \Delta - \bar{\Delta}_{\text{sym}}$ as

$$|\rho_2 - \bar{\rho}_2| \leq \frac{1}{d}\left(\frac{\|\boldsymbol{f}\|_2^2\|E_{\bar{\Delta}}\|_2}{|\boldsymbol{f}^\mathsf{T}\boldsymbol{f} - b_0|} + \frac{|\boldsymbol{f}^\mathsf{T}\bar{\Delta}_{\text{sym}}\boldsymbol{f}||b_0 - \bar{b}_0|}{|\boldsymbol{f}^\mathsf{T}\boldsymbol{f} - b_0||\boldsymbol{f}^\mathsf{T}\boldsymbol{f} - \bar{b}_0|}\right), \tag{21}$$

but the bound now also contains a term proportional to the difference between the uniform and stationary-weighted asymptotes.

To qualitatively assess the impact of biased mutation, the ruggedness analysis from Fig 3B was repeated in Fig 6A–C for different nucleotide-level mutation kernels. The amino-acid *NK* landscapes were represented at the nucleotide level using a codon-to-amino-acid mapping, and the analysis was performed on a representative set of 160,000 codon sequences. Fig 6A uses a uniform mutation kernel, $D = (\mathbf{1}\mathbf{1}^\mathsf{T} - I)/3$, whereas Fig 6C uses a directed, row-stochastic *E. coli* kernel $\bar{D}$ adapted from [41]. Fig 6B uses an undirected, symmetric kernel $\tilde{D} = Q((\bar{D} + \bar{D}^\mathsf{T})/2)Q$, where the diagonal matrix Q is obtained by symmetric Sinkhorn scaling such that $\tilde{D}$ is symmetric and doubly stochastic. Fig 6A–C show that

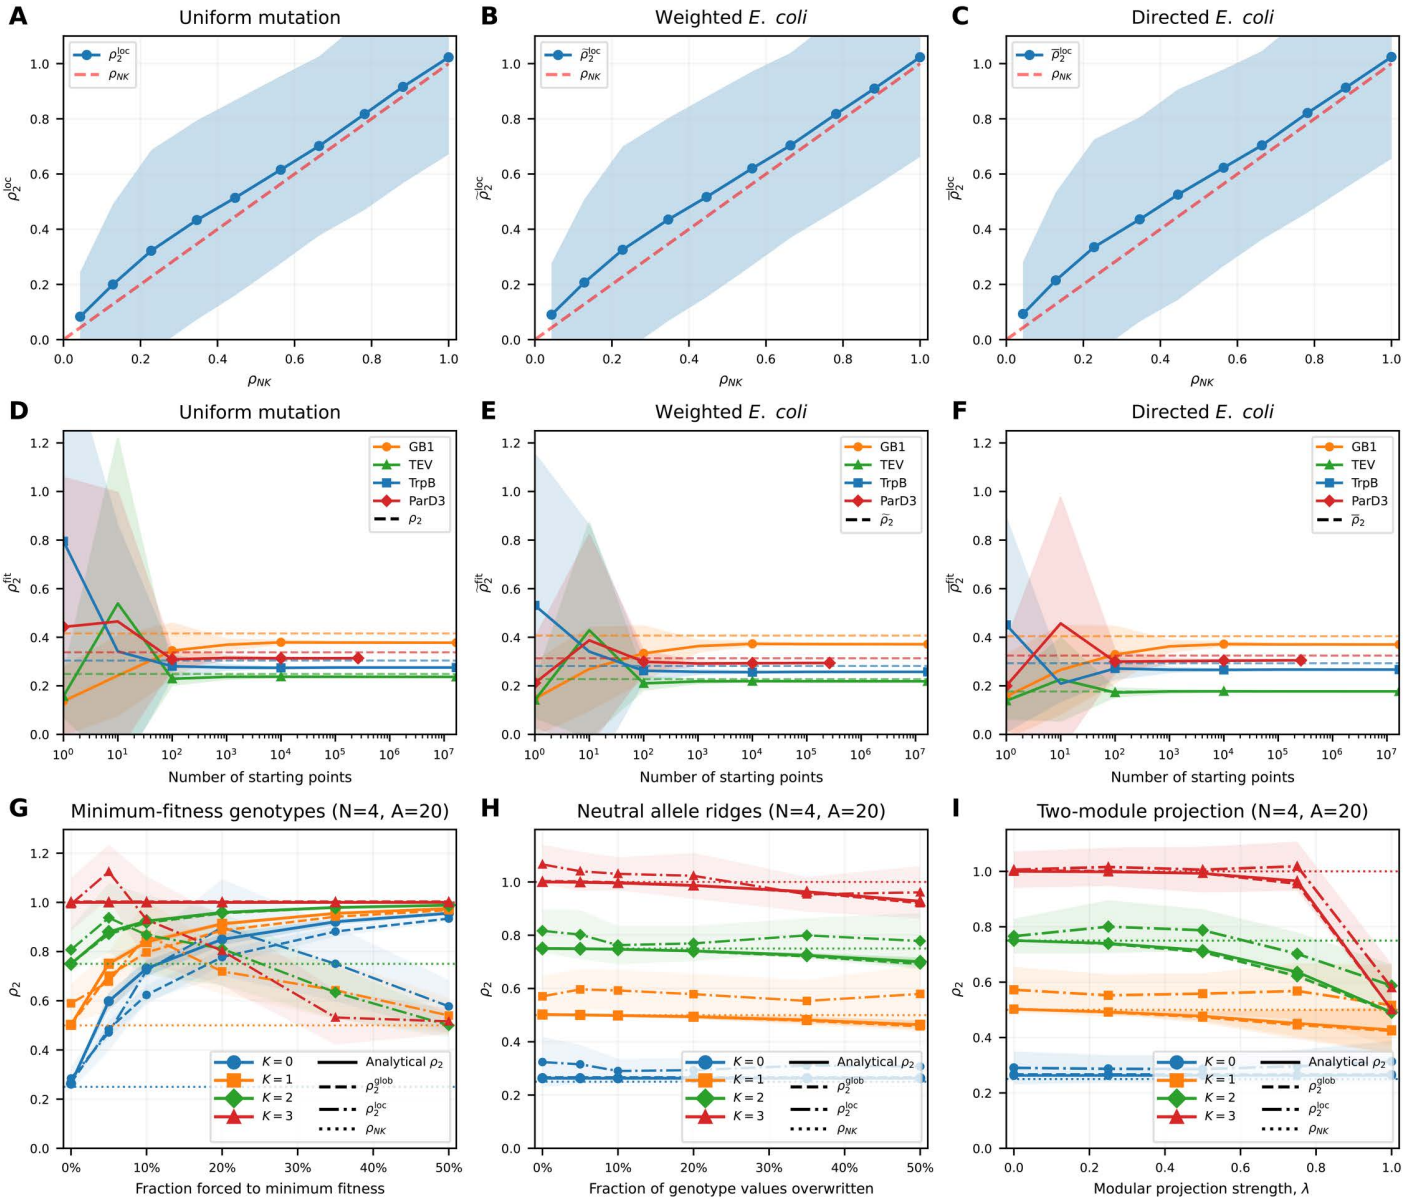


Fig 6. Ruggedness inference under mutation bias and landscape perturbations. A–C: Local decay-rate estimates for nucleotide-encoded *NK* landscapes under uniform **(A)**, weighted-undirected **(B)**, and directed *E. coli* (**C**) mutation kernels. The analysis used a representative set of 160,000 codon sequences for each landscape. Points and shaded regions show the mean and SD across landscapes and starting populations. **D–F:** Global decay-rate estimates for the codon-expanded GB1, TrpB, TEV, and ParD3 landscapes under the corresponding three mutation kernels. All nucleotide genotypes were included. Solid lines and shaded regions show the mean and SD of bootstrap estimates as the number of starting genotypes increases; dashed lines denote the corresponding analytical metrics. **G–I:** Effects of increasing fractions of minimum-fitness genotypes **(G)**, neutral ridges **(H)**, and modular projection strength (**I**) on $N=4$, $A=20$ *NK* landscapes. Solid, dashed, and dash-dotted lines denote analytical ρ_2, fitted ρ_2^{glob}, and fitted ρ_2^{loc}, respectively; shaded regions show the SD across 20 landscapes. Simulation parameters and further details are provided in Section G of S1 Appendix.

https://doi.org/10.1371/journal.pcbi.1014713.g006

the mutation kernels considered have no major effect on ruggedness estimation for *NK* landscapes. The estimates are instead dominated by the homogeneous statistical structure of the *NK* model.

Fig 6D–6F extend this analysis to empirical landscapes by repeating the population-sensitivity simulations from Fig 4D for the three mutation kernels used in Fig 6A–6C. In contrast to the representative sampling used for the *NK* landscapes, these simulations were performed over all 4^{3N} genotypes. The figures show the fitted estimates ρ_2^{fit} as a function of the number of starting points (solid lines), with shaded regions indicating the SD and dashed lines denoting the corresponding analytical squared-decay metrics ρ_2 (**D**), $\tilde{\rho}_2$ (**E**), and $\bar{\rho}_2$ (**F**). The nucleotide-to-amino-acid mapping spreads the landscape spectra across a broader range of frequencies and generally produces a small reduction in the inferred ruggedness. Consequently, fewer starting points are required to obtain a stable global estimate (cf. Fig 4B). However, because the resulting ruggedness values are more closely spaced, greater estimation accuracy is needed to reliably distinguish between ruggedness levels. S2 Fig repeats the analysis using a mutation model for *A. thaliana* and shows the same qualitative behaviour, and all decay rates are summarised in S2 Table. Further details are provided in Section G of S1 Appendix.

Landscape perturbations

We assessed whether the decay-based ruggedness estimate remains informative when the fitness landscape is subjected to controlled structural perturbations while the unbiased mutation process is retained (Fig 6G–6I). We generated synthetic *NK* landscapes ($N=4$, $A=20$, $K \in \{0, 1, 2, 3\}$, with 20 landscapes per value of K) and modified each landscape in three ways. In Fig 6G, an increasing fraction of randomly selected genotypes was assigned the minimum fitness of the corresponding landscape, thereby creating progressively larger minimum-fitness regions. In Fig 6H, neutral ridges were introduced by selecting one site and copying the fitness slice associated with a reference allele to an increasing number of alternative allelic slices at that site, making the corresponding regions identical in fitness. In Fig 6I, modularity was introduced by interpolating between the original landscape $\boldsymbol{f}$ and a two-module landscape, $\boldsymbol{f}_{\text{mod}} \stackrel{\text{def}}{=} \mathbb{E}[\boldsymbol{f}|x_1, x_2] + \mathbb{E}[\boldsymbol{f}|x_3, x_4] - \mathbb{E}[\boldsymbol{f}]$, such that $\boldsymbol{f}_\lambda \stackrel{\text{def}}{=} (1-\lambda)\boldsymbol{f} + \lambda \boldsymbol{f}_{\text{mod}}$, with $\lambda \in [0, 1]$. This construction preserves interactions within each two-site module while removing interactions between the two modules. For all three perturbations, a perturbation strength of zero corresponds to the original *NK* landscape. The analytical value of ρ_2 (solid lines) was recomputed directly from the Fourier spectrum of each perturbed landscape, whereas ρ_2^{glob} (dashed lines) and ρ_2^{loc} (dash-dotted lines) were estimated from simulated fitness trajectories on the same landscape.

Fig 6G shows that increasing the fraction of minimum-fitness genotypes raises ρ_2, consistent with the introduction of sharp fitness discontinuities that redistribute spectral power towards higher frequencies. This effect is stronger for less rugged landscapes than for rugged landscapes, which already contain substantial high-frequency power. Conversely, Fig 6H shows that increasing the extent of the neutral ridges reduces ρ_2, because copying fitness slices suppresses variation along the affected directions and shifts spectral power towards lower frequencies. This effect is more pronounced for rugged landscapes than for less rugged landscapes, whose spectra are already dominated by lower-frequency components. Finally, Fig 6I shows that increasing the modular projection strength λ reduces ρ_2. In particular, ρ_2 decreases from 1.0 to 0.5 for $K=3$ and from 0.75 to 0.491 for $K=2$, reflecting the fact that, at $\lambda = 1$, no interaction can involve more than two sites.

Fig 6G–6I further show that the global estimates, ρ_2^{glob}, consistently reproduce the analytical trends. The local estimates, ρ_2^{loc}, also reproduce the trends for the neutral-ridge and modularity perturbations in Fig 6H and 6I, but not for the minimum-fitness perturbation in Fig 6G, which produces increasingly heterogeneous landscapes as the fraction of minimum-fitness genotypes grows.

Estimating power spectra and ruggedness parameters

Generally, the power spectrum of a fitness landscape is composed of both low-frequency and high-frequency terms, as shown in Fig 7A. Practically, if the landscape f is known, its power spectrum can be computed by reshaping f as an

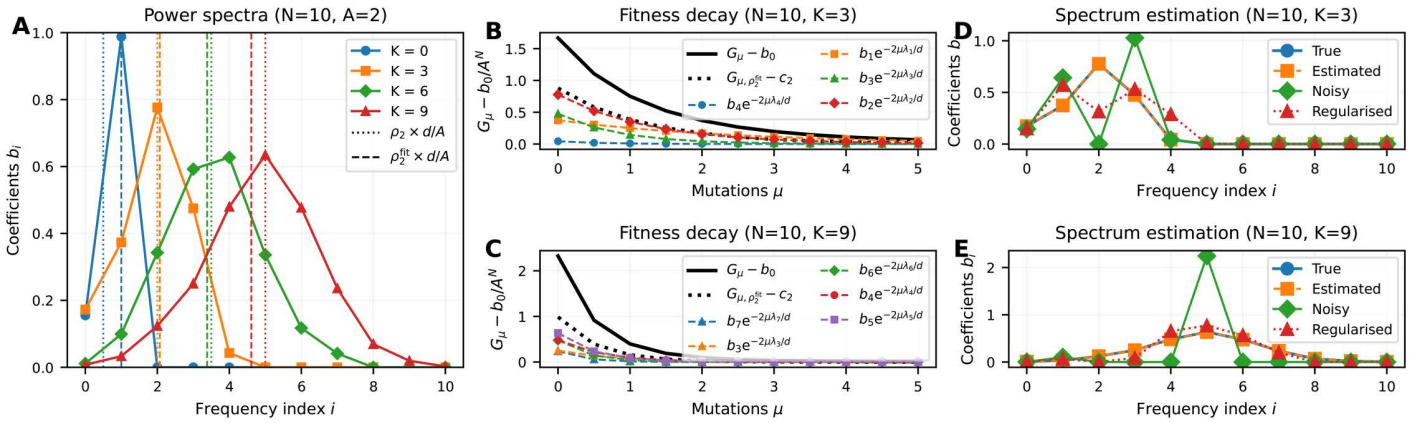


Fig 7. Power-spectrum estimation on *NK* landscapes ($N=10$, $A=2$). A: Power spectra for $K \in \{0, 3, 6, 9\}$. **B–C:** Fitness-decay curves for $K=3$ and $K=9$, respectively. Shown are the non-constant decay $G_\mu - b_0/A^N$, the single-rate fit $G_{\mu,\rho_2^{\text{fit}}} - c_2$ from Eq (26), and the largest non-constant components $b_i e^{-2\mu\lambda_i/d}/A^N$ of the decay decomposition. **D–E:** Power-spectrum estimates for $K=3$ and $K=9$, respectively. The noiseless estimate uses constrained least squares (Eq (22)); noisy estimates use non-negative least squares with or without regularisation (Eq (24)). Regularisation reduces the effect of noise.

https://doi.org/10.1371/journal.pcbi.1014713.g007

N-dimensional cube of size A along each axis, applying an N-dimensional Fast Fourier Transform, and grouping the terms as in Eq (7b).

If the landscape is not known, the power spectral coefficients b_i from Eq (8) can be estimated using the decay curve G_μ from Eq (7b). The procedure for estimating a_i from F_μ is similar. Let $G = (G_0, G_{\mu_1}, ..., G_{\mu_M})^{\mathsf{T}}$ be the squared average fitness measurements of populations accumulating $0, \mu_1, ..., \mu_M$ mutations and define $b = (b_0, ..., b_N)^{\mathsf{T}}$. The following least squares problem can be used to estimate b from fitness measurements G:

$$b = \underset{x \in \mathbb{R}^{N+1}_{\geq 0}}{\arg\min} \|G - \Phi x\|_2^2, \tag{22}$$

where element i,j of the matrix $\Phi \in \mathbb{R}^{(M+1)\times(N+1)}$ is defined as

$$\Phi_{i,j} = e^{-\frac{2\mu_i \lambda_j}{d}}, \tag{23}$$

for $i = 0, ..., M$ and $j = 0, ..., N$. For noise-free observations, problem (22) can be solved using standard techniques to accurately determine the fitness landscape spectrum b (Fig 7D–7E). However, the addition of noise with a signal-to-noise ratio ≤ 10 can already significantly perturb the estimate from Eq (22), because the ill-conditioned Φ amplifies the estimation error. To improve the robustness of the estimate, Eq (22) can be extended with a regularisation term as

$$b = \underset{x \in \mathbb{R}^{N+1}_{\geq 0}}{\arg\min} \|G - \Phi x\|_2^2 + \gamma \|x\|_2^2, \tag{24}$$

where $\gamma \geq 0$ is a regularisation parameter. While the results from Fig 7D–7E show that regularisation can improve the estimation accuracy, it depends on having chosen γ appropriately; choosing a poor γ will lower accuracy. The inference method could be developed further, incorporating more knowledge about the distribution of b and the noise process, which is beyond the scope of this paper.

Instead of estimating the full spectrum, we estimate fitted ruggedness parameters using the approximations from Eq (10). The ρ_2^{fit} obtained from fitting G_μ will in general not equal the analytical ρ_2 from Eq (13b); according to Eq (14), the analytical ρ_2 defines a lower-bound single exponential. If the power spectrum is concentrated in a single frequency (such as for *NK* landscapes), then $\rho_2^{\text{fit}} \approx \rho_2$. The scalar parameters C_2, ρ_2^{fit}, and c_2 are unknown *a priori* and must be estimated from data. As for the power spectrum estimation, the squared average fitness values G are recorded for a population accumulating $\mu_1, ..., \mu_M$ mutations. We then seek to find C_2, ρ_2^{fit}, and c_2 such that

$$G_{\mu_i} \approx C_2 \mathrm{e}^{-2\mu_i \rho_2^{\text{fit}}} + c_2. \tag{25}$$

As for the full spectrum estimation, the mutations are stochastic and the decay curves are noisy, which impacts the parameter estimation process. To simplify the estimation, we assume that G_0 is noise-free, and that the curve fitted via Eq (25) must pass through it, which allows us to discard parameter C_2 by setting $C_2 = G_0 - c_2$. For the remaining data points, we seek to find ρ_2^{fit} and c_2 that minimise the sum of squares error:

$$\min_{\rho_2^{\text{fit}}, c_2 \in \mathbb{R}_{\geq 0}} \sum_{i=1}^{M} \left((G_0 - c_2)\mathrm{e}^{-2\mu_i \rho_2^{\text{fit}}} - (G_{\mu_i} - c_2) \right)^2. \tag{26}$$

The procedure for estimating ρ_1^{fit} from F_μ is similar and requires discarding the factor 2 from exponentials. As long as $c_2 \neq 0$, Eq (26) is a nonlinear optimisation problem that cannot be solved using standard least squares techniques and may have multiple local optima. Here, we are using standard scientific packages for constrained nonlinear optimisation (specifically, the `curve_fit` function from the `scipy.optimize` package [42]), which returns a solution within 1 s on a standard laptop. Initialising the routine with $\rho_2^{\text{fit}}|_{\text{init}} = 0.5$ and $c_2|_{\text{init}} = G_M$ and constraining the optimisation variables by $\rho_2^{\text{fit}} \in [0, A/(A-1)]$ and $c_2 \in [0.5 \times c_2|_{\text{init}}, 1.5 \times c_2|_{\text{init}}]$ (see Code Availability), this procedure is robust to noise and could be extended with uncertainty weights (regularisation). The curve fitting for two *NK* landscapes is shown in Fig 7B–7C, with the fitted decay rates ρ_2^{fit} shown in Fig 7A. For these examples, the power spectrum peaks at a particular frequency, so that Eq (25) results in a good approximation of the true decay composed of multiple exponentials.

Supporting information

S1 Appendix. Supplementary methods and analyses. Contains additional details on mutagenesis and directed-evolution simulations, alternative ruggedness metrics, special cases of the decay-rate metric, empirical fitness landscapes, alternative landscape models, and alternative mutation models [2,4,6–8,10,24–26,39,41,43,44].
(PDF)

S1 Table. Summary of ruggedness and decay-rate notation. Definitions of the analytical and fitted metrics derived from the average fitness decay F_μ and squared average fitness decay G_μ.
(PDF)

S2 Table. Analytical decay rates for empirical landscapes. Decay-rate metrics for the GB1, TrpB, TEV, and ParD3 landscapes under uniform, weighted-undirected, and directed mutation models.
(PDF)

S1 Fig. Comparison of strategy spaces and decay curves across landscape models. Strategy spaces and corresponding fitness-decay curves for *NK*, Rough Mount Fuji, and stochastic block-model landscapes of varying ruggedness.
(PDF)

S2 Fig. Ruggedness inference using an *A. thaliana* mutation model. Replication of Fig 6A-F using weighted-undirected and directed *A. thaliana* mutation kernels on synthetic *NK* and empirical landscapes.
(PDF)

S3 Fig. Decay curves on empirical landscapes under different mutation models. Sampled, fitted, and analytical squared-fitness decay curves for empirical landscapes under uniform, weighted-undirected, and directed mutation kernels.
(PDF)

S4 Fig. Strategy look-up tables for longer directed-evolution experiments. Optimal base-chance and population-splitting parameters for *NK* landscapes over 25, 50, and 100 directed-evolution iterations.
(PDF)

Author contributions

Conceptualization: Sebastian Towers, Jessica James, Harrison Steel, Idris Kempf.

Data curation: Sebastian Towers, Jessica James.

Formal analysis: Sebastian Towers, Jessica James, Idris Kempf.

Investigation: Sebastian Towers, Jessica James, Idris Kempf.

Methodology: Sebastian Towers, Jessica James, Idris Kempf.

Project administration: Sebastian Towers, Jessica James, Idris Kempf.

Software: Sebastian Towers, Jessica James, Idris Kempf.

Supervision: Idris Kempf.

Validation: Sebastian Towers, Jessica James.

Visualization: Sebastian Towers, Jessica James, Idris Kempf.

Writing – original draft: Sebastian Towers, Jessica James, Idris Kempf.

Writing – review & editing: Sebastian Towers, Jessica James, Harrison Steel, Idris Kempf.

References

1. Wright S. The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution. In: Proceedings of the XI International Congress of Genetics, 1932. 209–22.
2. Wu NC, Dai L, Olson CA, Lloyd-Smith JO, Sun R. Adaptation in protein fitness landscapes is facilitated by indirect paths. eLife. 2016;5:e16965. https://doi.org/10.7554/eLife.16965
3. Huang M, Zhou S, Li K. Augmenting Biological Fitness Prediction Benchmarks with Landscapes Features from GraphFLA. 2025. https://arxiv.org/abs/2510.24826
4. Aita T, Iwakura M, Husimi Y. A cross-section of the fitness landscape of dihydrofolate reductase. Protein Eng. 2001;14(9):633–8. https://doi.org/10.1093/protein/14.9.633 PMID: 11707608
5. Stadler PF. Landscapes and their correlation functions. J Math Chem. 1996;20(1):1–45. https://doi.org/10.1007/bf01165154
6. Poelwijk FJ, Tănase-Nicola S, Kiviet DJ, Tans SJ. Reciprocal sign epistasis is a necessary condition for multi-peaked fitness landscapes. J Theor Biol. 2011;272(1):141–4. https://doi.org/10.1016/j.jtbi.2010.12.015 PMID: 21167837
7. Weinreich DM, Watson RA, Chao L. Perspective: Sign epistasis and genetic constraint on evolutionary trajectories. Evolution. 2005;59(6):1165–74. https://doi.org/10.1111/j.0014-3820.2005.tb01768.x PMID: 16050094
8. Szendro IG, Schenk MF, Franke J, Krug J, de Visser JAGM. Quantitative analyses of empirical fitness landscapes. J Stat Mech. 2013;2013(01):P01005. https://doi.org/10.1088/1742-5468/2013/01/p01005
9. Weinberger ED. Fourier and Taylor series on fitness landscapes. Biological Cybernetics. 1991;65(5):321–30. https://doi.org/10.1007/BF00216965
10. James J, Towers S, Foerster J, Steel H. Optimisation strategies for directed evolution without sequencing. PLoS Comput Biol. 2024;20(12):e1012695. https://doi.org/10.1371/journal.pcbi.1012695 PMID: 39700257

11. Arnold FH. Design by Directed Evolution. Acc Chem Res. 1998;31(3):125–31. https://doi.org/10.1021/ar960017f
12. Fox R, Roy A, Govindarajan S, Minshull J, Gustafsson C, Jones JT, et al. Optimizing the search algorithm for protein engineering by directed evolution. Protein Eng. 2003;16(8):589–97. https://doi.org/10.1093/protein/gzg077 PMID: 12968076
13. Wu Z, Kan SBJ, Lewis RD, Wittmann BJ, Arnold FH. Machine learning-assisted directed protein evolution with combinatorial libraries. Proc Natl Acad Sci U S A. 2019;116(18):8852–8. https://doi.org/10.1073/pnas.1901979116 PMID: 30979809
14. Yang J, Lal RG, Bowden JC, Astudillo R, Hameedi MA, Kaur S, et al. Active learning-assisted directed evolution. Nat Commun. 2025;16(1):714. https://doi.org/10.1038/s41467-025-55987-8 PMID: 39821082
15. Li F, Yang J, Johnston K, Gürsoy E, Yue Y, Arnold F. Evaluation of machine learning–assisted directed evolution across diverse combinatorial landscapes. bioRxiv. 2024. https://doi.org/10.1101/2024.10.24.619774
16. Machado FML, Currin A, Dixon N. Directed evolution of the PcaV allosteric transcription factor to generate a biosensor for aromatic aldehydes. J Biol Eng. 2019;13:91. https://doi.org/10.1186/s13036-019-0214-z PMID: 31798685
17. Chen S, Yang Z, Zhong Z, Yu S, Zhou J, Li J, et al. Ultrahigh-throughput screening-assisted in vivo directed evolution for enzyme engineering. Biotechnol Biofuels Bioprod. 2024;17(1):9. https://doi.org/10.1186/s13068-024-02457-w PMID: 38254175
18. Tan Y, Zhang Y, Han Y, Liu H, Chen H, Ma F, et al. Directed evolution of an α1,3-fucosyltransferase using a single-cell ultrahigh-throughput screening method. Sci Adv. 2019;5(10):eaaw8451. https://doi.org/10.1126/sciadv.aaw8451 PMID: 31633018
19. Carpenter AC, Feist AM, Harrison FSM, Paulsen IT, Williams TC. Have you tried turning it off and on again? Oscillating selection to enhance fitness-landscape traversal in adaptive laboratory evolution experiments. Metab Eng Commun. 2023;17:e00227. https://doi.org/10.1016/j.mec.2023.e00227 PMID: 37538933
20. Alpay BA, Desai MM. Effects of selection stringency on the outcomes of directed evolution. PLoS One. 2024;19(10):e0311438. https://doi.org/10.1371/journal.pone.0311438 PMID: 39401192
21. De Ath G, Everson RM, Rahat AAM, Fieldsend JE. Greed is good: exploration and exploitation trade-offs in Bayesian optimisation. ACM Trans Evol Learn Optim. 2021;1(1). https://doi.org/10.1145/3425501
22. Ueda M, Takeuchi N, Kaneko K. Stronger selection can slow down evolution driven by recombination on a smooth fitness landscape. PLoS One. 2017;12(8):e0183120. https://doi.org/10.1371/journal.pone.0183120 PMID: 28809951
23. Kauffman SA, Weinberger ED. The NK model of rugged fitness landscapes and its application to maturation of the immune response. J Theor Biol. 1989;141(2):211–45. https://doi.org/10.1016/s0022-5193(89)80019-0 PMID: 2632988
24. Lite TV, Grant RA, Nocedal I, Littlehale ML, Guo MS, Laub MT. Uncovering the basis of protein-protein interaction specificity with a combinatorially complete library. eLife. 2020;9:e60924. https://doi.org/10.7554/eLife.60924
25. Johnston KE, Almhjell PJ, Watkins-Dulaney EJ, Liu G, Porter NJ, Yang J, et al. A combinatorially complete epistatic fitness landscape in an enzyme active site. Proc Natl Acad Sci U S A. 2024;121(32):e2400439121. https://doi.org/10.1073/pnas.2400439121 PMID: 39074291
26. Tu B, Sundar V, Esvelt KM. An ultra-high-throughput method for measuring biomolecular activities. bioRxiv. 2024. https://doi.org/10.1101/2022.03.09.483646
27. Huang M, Mao P, Li K. Rethinking performance analysis for configurable software systems: A case study from a fitness landscape perspective. Proc ACM Softw Eng. 2025;2. https://doi.org/10.1145/3728954
28. Huang M, Li K. On the Hyperparameter Loss Landscapes of Machine Learning Models: An Exploratory Study. In: Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1, 2025. 555–64. https://doi.org/10.1145/3690624.3709229
29. Huang M, Li K. Exploring Structural Similarity in Fitness Landscapes via Graph Data Mining: A Case Study on Number Partitioning Problems. In: Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, 2023. 5595–603. https://doi.org/10.24963/ijcai.2023/621
30. Qiu Y, Hu J, Wei G-W. Cluster learning-assisted directed evolution. Nat Comput Sci. 2021;1(12):809–18. https://doi.org/10.1038/s43588-021-00168-y PMID: 35811998
31. Wang P, Robert L, Pelletier J, Dang WL, Taddei F, Wright A, et al. Robust growth of Escherichia coli. Curr Biol. 2010;20(12):1099–103. https://doi.org/10.1016/j.cub.2010.04.045 PMID: 20537537
32. Eaton DS, Sánchez C, Gutiérrez-López L, Shenker JQ, Goulev Y, Watson BR, et al. Essentialome-wide multigenerational imaging reveals mechanistic origins of cell growth laws. bioRxiv. 2025. https://doi.org/10.1101/2025.06.10.658525
33. James J, Kempf I, Sechkar K, Wang J, Abrahams G, Towers S, et al. Microscopic PhotoSelection (MiPS) of single cells in mother machine microfluidic devices. 2025. https://doi.org/10.64898/2025.12.03.692065
34. Park H, Kim S. Gene-specific mutagenesis enables rapid continuous evolution of enzymes in vivo. Nucleic Acids Res. 2021;49(6):e32. https://doi.org/10.1093/nar/gkaa1231 PMID: 33406230
35. Castle SD, Grierson CS, Gorochowski TE. Towards an engineering theory of evolution. Nat Commun. 2021;12(1):3326. https://doi.org/10.1038/s41467-021-23573-3 PMID: 34099656
36. Stadler PF. Spectral Landscape Theory. Evolutionary Dynamics: Exploring the Interplay of Selection, Accident, Neutrality, and Function. 2003. 231–71.
37. Neidhart J, Szendro IG, Krug J. Exact results for amplitude spectra of fitness landscapes. J Theor Biol. 2013;332:218–27. https://doi.org/10.1016/j.jtbi.2013.05.002 PMID: 23685065

38. Godsil C, Royle G. Algebraic Graph Theory. 2001.
39. Stadler PF, Happel R. Random field models for fitness landscapes. Journal of Mathematical Biology. 1999;38(5):435–78. https://doi.org/10.1007/s002850050156
40. Molina RS, Rix G, Mengiste AA, Alvarez B, Seo D, Chen H, et al. In vivo hypermutation and continuous evolution. Nat Rev Methods Primers. 2022;2:37. https://doi.org/10.1038/s43586-022-00130-w PMID: 37073402
41. Lynch M. Rate, molecular spectrum, and consequences of human mutation. Proc Natl Acad Sci U S A. 2010;107(3):961–8. https://doi.org/10.1073/pnas.0912629107 PMID: 20080596
42. Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat Methods. 2020;17(3):261–72. https://doi.org/10.1038/s41592-019-0686-2 PMID: 32015543
43. Nam H, Hwang BJ, Choi D-Y, Shin S, Choi M. Tobacco etch virus (TEV) protease with multiple mutations to improve solubility and reduce self-cleavage exhibits enhanced enzymatic activity. FEBS Open Bio. 2020;10(4):619–26. https://doi.org/10.1002/2211-5463.12828 PMID: 32129006
44. Oberer M, Zangger K, Gruber K, Keller W. The solution structure of ParD, the antidote of the ParDE toxin antitoxin module, provides the structural basis for DNA and toxin binding. Protein Sci. 2007;16(8):1676–88. https://doi.org/10.1110/ps.062680707 PMID: 17656583

S1 Appendix for *Sequence-free Landscape Inference for Directed Evolution*

A Simulating mutagenesis and selection

In our simulations, we represent a population of P cells at iteration k as $\boldsymbol{p}_k = (\boldsymbol{x}_k^1, \dots, \boldsymbol{x}_k^P)$, where $\boldsymbol{x}_k^i \in \mathcal{G}$ is the genotype of cell i. Given an empirical or theoretical landscape f, each genotype is associated with a fitness $f_k^i = f(\boldsymbol{x}_k^i)$. When the population $\boldsymbol{p}_k$ is mutagenised, each member $\boldsymbol{x}_k^i$ is exposed to X single mutations to yield $\boldsymbol{x}_{k+1}^i$, where X is an independent and identically distributed Poisson random variable. To obtain the decay curves used for ruggedness parameter estimation, we repeat the mutagenesis process for M iterations and record the average population fitness $F_k = 1/P \sum_{i=1}^P f_k^i$, the squared mean fitness F_k^2, which is subsequently averaged across starting populations to estimate G_μ, and the average accumulated mutations μ_k. This process is repeated and averaged over different starting populations $\boldsymbol{p}_0$ to generate a decay curve such as the one shown in Fig 3A.

For the selection process used in directed evolution (DE) experiments, each iteration is extended with a selection procedure. Here, we use an approach developed in [10] that parameterises selection using probabilistic selection functions. The population members $\boldsymbol{x}_k^i$ are ranked according to their fitness f_k^i. A fraction of cells is then randomly selected from $\boldsymbol{p}_k$ according to the probability distribution shown in Fig 1 (top right). This distribution is parameterised by two parameters: the base chance b and threshold t. If a cell's fitness f_k^i is above the rank threshold t (within the population $\boldsymbol{p}_k$), it is guaranteed to be selected for iteration $k+1$. Otherwise, the selection probability is b, which is independent between cells. Here, the fraction of cells selected at each iteration is kept constant as $\mathrm{f} = 0.2$, fixing t as $t = (1-\mathrm{f})/(1-b)$. Each cell that is not selected is replaced by a new cell with genotype taken at random (uniformly and independently) from the pool of cells that were selected. For example, for $t = 0.8$ and $b = 0$ the top $20\,\%$ of fittest cells are selected for the next iteration. The case $t = 1.0$ and $b = 0.2$ represents a selection-free scenario, with a random $20\,\%$ of cells surviving each generation. All simulations were carried out using a GPU-based DE package published in [10].

B Simulating directed evolution

For simulating DE experiments, the selection and mutagenesis procedures are repeatedly applied to a starting population $\boldsymbol{p}_0$. The selection function is parameterised by the base chance b, while the rank threshold t is fixed via $t = (1-\mathrm{f})/(1-b)$. The fraction of selected cells f is set to $\mathrm{f} = 0.2$, so that at each iteration, approximately $20\,\%$ of the population is selected for the following iteration. The value of b is determined using the decay rate estimate ρ^{fit}. Before starting DE, the population $\boldsymbol{p}_0$ is repeatedly mutagenised without selection for M iterations. The resulting decay curve is used to estimate ρ^{fit}, which is in turn used to select b from a pre-computed map that connects ρ_{NK} grid

values to optimal selection parameters. This map is pre-computed according to the work presented in [10]. In addition to the parameter b, the work from [10] also incorporates a population splitting parameter s that subdivides the population $\boldsymbol{p}_0$ into s smaller populations, yielding s smaller DE experiments that run in parallel. This is illustrated in Fig 1 (bottom right).

The map linking NK ruggedness to the optimal parameter set (b_ρ, s_ρ) is computed using NK landscapes, whose grid values are $\rho_{NK} = (K+1)/N$ via Eq (S2). For fixed N and A, we generate a number of random NK landscapes for $K = 1, \dots, N-1$. For each K, we run DE experiments for different starting populations $\boldsymbol{p}_0$ and different parameter sets (b, s), and store best-performing parameters for each K, defining a map from (N, K) to (b, s). Given N and the fitted estimate ρ^{fit}, the parameter set associated with the value of ρ_{NK} closest to ρ^{fit} is selected. An example is shown in Fig 1 (middle) for $N = 4$, $A = 20$, and $K \in \{1, 2, 3\}$ ($\rho_{NK} \in \{0.5, 0.75, 1\}$). Note that while ρ^{fit} can take on continuous positive values, ρ_{NK} can only take on discrete values for fixed N. Alternatively, the generated maps can be averaged across different values of N.

C Other ruggedness metrics

Various measures have been proposed to quantify landscape ruggedness, most of which require both genotypic and phenotypic information. For our analysis, we compare the analytical decay-rate measure ρ_2 with six other measures, summarised below. Most metrics were originally defined on an allele set of size 2 ($|\mathcal{A}| = 2$), and are extended here to larger sets ($|\mathcal{A}| > 2$). For a more comprehensive discussion of ruggedness measures, we refer the reader to [8].

Roughness to slope ratio (r/s):

Given a genotype $\boldsymbol{x} = (\sigma_1, \dots, \sigma_N)^{\mathsf{T}} \in \mathcal{G}$, a landscape $f_{lin} : \mathcal{G} \to \mathbb{R}$ is *linear* if the contribution of each locus and allele is additive, i.e. $f_{lin}(\boldsymbol{x})$ can be written as $f_{lin}(\boldsymbol{x}) = \sum_{i=1}^{N} \phi_i(\sigma_i)$, where $\phi_i : \mathcal{A} \to \mathbb{R}$ is the fitness contribution of locus i for allele σ_i. The roughness to slope ratio [4] captures how well a given landscape f can be approximated by a linear one: After finding the best linear fit f_{lin} to f, i.e. the $N \times A$ values $\phi_i(\sigma)$, the mean slope is defined as $s \stackrel{\text{def}}{=} \sum_{i=1}^{N} \sum_{\sigma_i \in \mathcal{A}} |\phi_i(\sigma_i)|/(AN)$, the roughness as $r \stackrel{\text{def}}{=} (\sum_{x \in \mathcal{G}} (f(\boldsymbol{x}) - f_{lin}(\boldsymbol{x}))^2 / A^N)^{1/2}$, and the roughness to slope ratio as r/s. For a linear landscape, $r/s = 0$, whereas for more rugged landscapes r/s increases.

Landscape R^2 measure (R^2):

The landscape R^2 measure was originally defined in [8] using landscape Fourier analysis and is directly related to the measure ρ_2 introduced in this work. Suppose that the landscape is represented in vector notation as $\boldsymbol{f} \in \mathbb{R}^{A^N}$ and mapped to the Fourier domain as $\hat{\boldsymbol{f}} = P^* \boldsymbol{f}$. Define $b_i \stackrel{\text{def}}{=} \sum_{j \in \mathcal{I}_i} |\hat{\boldsymbol{f}}_j|^2$, where $\mathcal{I}_i$ refers to the indices of $\hat{\boldsymbol{f}}$ associated with the ith eigenvalue of the adjacency matrix. The R^2 measure is then defined as $R^2 \stackrel{\text{def}}{=} (\sum_{i=2}^{N} b_i)/(\sum_{i=1}^{N} b_i) \in [0, 1]$, i.e., the higher-order epistasis frac-

tion. For a linear landscape, $R^2 = 0$, whereas for a more rugged landscape R^2 increases. The R^2 measure can be computed from decay curves if the full spectrum is estimated via Eqs (22) or (24). However, the difficulties in estimating the full spectrum can make the R^2 measure less robust than the fitted decay rate ρ_2^{fit}. We denote this quantity as "R^2" because it is equivalent to the coefficient of determination.

Spectral entropy (H):

Similarly to the R^2 measure, the spectral or Shannon entropy H uses the power spectral coefficients b_i and is defined as $H \stackrel{\text{def}}{=} -\sum_{i=1}^{N} p_i \log p_i / \log N \in [0, 1]$ with $p_i \stackrel{\text{def}}{=} b_i / \sum_{j=1}^{N} b_j$. The entropy is large when the power is spread uniformly across frequencies, reaching its maximum at 1 when $b_i = b_j \forall i, j \in \{1, \ldots, N\}$, and small if the power is concentrated in a few coefficients, reaching its minimum when $b_j = 0 \forall j \neq i$ for some i.

Local epistasis (n_ϵ):

Epistasis occurs when the fitness contribution of one of the N loci depends on the allele set at other loci, or in terms of mutations, when the effect of one mutation depends on the presence of others. Simple sign epistasis occurs when a mutation is beneficial in the presence of one set of mutations but disadvantageous in another. For example, consider a landscape with $N = 2$ and $\mathcal{A} = \{0, 1\}$, so that $\boldsymbol{x} = (\sigma_1, \sigma_2)^{\mathsf{T}}$ with $\sigma_i \in \mathcal{A}$. Simple sign epistasis occurs if changing σ_1 from 0 to 1 has a positive effect on fitness when $\sigma_2 = 0$ and a negative effect when $\sigma_2 = 1$. Here, we measure local epistasis as the number of genotype pairs with a Hamming distance 2 that exhibit simple sign epistasis.

Distance to local maxima (d_{max}):

A local maximum is defined as a genotype with higher fitness than all its neighbours. The metric d_{max} measures [6] the Hamming distance to the closest local maximum, averaged over all genotypes. A large d_{max} indicates fewer local maxima, hence a smoother landscape.

Number of direct paths to global maximum (π_{max}):

Given a global fitness maximum at genotype $\boldsymbol{x}^\star$, this metric [7] counts the number of fitness-increasing direct (length N) paths π_{max} from the antipodal genotype $\bar{\boldsymbol{x}}^\star$ to $\boldsymbol{x}^\star$. The antipodal genotype $\bar{\boldsymbol{x}}^\star$ is the genotype where, for each locus, the corresponding alleles in $\bar{\boldsymbol{x}}^\star$ and $\boldsymbol{x}^\star$ are different. For $|\mathcal{A}| > 2$, there are multiple antipodal genotypes, so we define π_{max} to be the mean over all $(A-1)^N$ possible $\bar{\boldsymbol{x}}^\star$. A small π_{max} indicates a more rugged landscape.

D Special cases for decay rate metric

House-of-cards landscapes:

For zero-mean house-of-cards landscapes, the fitness values are assumed to be random, independent, and normally distributed as $\boldsymbol{f}_i \sim \mathcal{N}(0,1)$. Due to the orthonormality of $\boldsymbol{P}$, this results in all $\hat{\boldsymbol{f}}_i$ being independent of one another, and $\hat{\boldsymbol{f}}_i \sim \mathcal{N}(0,1)$, i.e. the Fourier representation inherits the same distribution. As a result, the terms b_i in Eq (7b) are chi-squared distributed with mean equal to the multiplicity of λ_i. The eigenvalue multiplicity $\binom{N}{i}(A-1)^i$ is maximised for $i^\star = \text{round}(N(A-1)/A)$, so that $|b_{i^\star}| \gg |b_i|$ for $i \neq i^\star$. To see why this is true, note that the function $\binom{N}{i}(A-1)^i/A^N$ can be interpreted as a binomial distribution, in which case $|i - i_c^\star|/N \leq 3\sqrt{A-1}/(\sqrt{N}A)$ with $i_c^\star = N(A-1)/A$ contains $\sim 99.7\,\%$ of its weight. Substituting the eigenvalue $\lambda_{i^\star} = Ai^\star$ yields the decay rate for house-of-cards landscapes:

$$\rho_{\text{HoC}} = \frac{1}{d} \times \lambda_{i^\star} = \frac{1}{N(A-1)} \times A\frac{N(A-1)}{A} = 1. \tag{S1}$$

Eq (S1) could also be derived differently. Since every fitness value is independent, with an average fitness of 0, every time a cell mutates its new fitness is on average 0. Hence the fitness curve measures how many non-mutated cells are in the population, which is equal to $\mathrm{e}^{-\mu}$.

NK landscapes:

This analysis can be extended to NK landscapes, which can be viewed as a sum of N smaller house-of-cards landscapes of (effective) dimensionality $K+1$ [39]. For each of the smaller landscapes, the eigenvalue multiplicity is maximised for $i^\star = \text{round}((K+1)(A-1)/A)$, which, when substituted in Eq (13), yields

$$\rho_1 \approx \rho_2 \approx \frac{1}{N(A-1)} \times A\frac{(K+1)(A-1)}{A} = \frac{K+1}{N} \stackrel{\text{def}}{=} \rho_{NK}. \tag{S2}$$

Since for NK landscapes the parameter $K \leq N-1$ measures the number of epistatic interactions and therefore ruggedness, Eq (S2) shows that the decay rate definition of ruggedness is closely related to the NK notion of ruggedness. Specifically, the analytical decay-rate metric satisfies $\rho_1 \approx \rho_2 \approx \rho_{NK}$ for these landscapes.

E Empirical landscapes and ruggedness

To interpret the quantitative predictions of each ruggedness metric it is necessary to consider two factors: the intrinsic properties of the protein and the assay used to generate fitness scores. While we speculate here on differences between landscapes and potential causes for their relative ruggedness, we caution that such interpretations require validation through deeper analysis and complementary measurements to disentangle biological from assay-based contributions.

TEV:

The data for the TEV landscape was downloaded from [26], and converted into a $20\times20\times20\times20$ `numpy` array. The TEV protease (the most rugged in all metrics except when measured by the spectral entropy H) recognises a highly specific substrate sequence, and its activity depends acutely on maintenance of its catalytic triad [43], as well as on its substrate-binding pocket and global fold. Catalytic triads in particular are among the least mutation-tolerant motifs in biology—even single substitutions can abolish function—which may contribute to the highly rugged landscape observed. The assay used to collect this data was based on DNA recording with a base editor [26]. This data is inherently noisy due to the added layer of biological complexity, which is likely to contribute significantly to the ruggedness of this landscape. However, of note is that this assay is much more linearly proportional to activity than the growth-based assays used to measure TrpB and ParD3.

ParD3:

The data for the ParD3 landscape was downloaded from [24], and converted into a $20 \times 20 \times 20$ `numpy` array. ParD3 is an antitoxin in which the C-terminal region is intrinsically disordered in the absence of its toxin partner. Such disordered regions typically rely on general physicochemical properties rather than rigid 3D structures, allowing greater mutational tolerance and producing a smoother landscape [44]. From the assay perspective, this fitness landscape was mapped using a growth competition assay [24]. Such an assay is inherently less noisy than that used to evaluate TEV, due to the fact that the measurement is population-level, and the exponential nature of growth also acts to amplify signal at the upper end of fitness. Similarly, growth-based assays of this kind may exhibit "saturation" in the sense that ParD3 variants with toxin binding ability above some threshold at which they can sequester the vast majority of toxin in a cell can thus yield similar (growth-based) fitness scores, even if *in practice* their binding affinity (if it were directly measured) differs significantly. This factor, together with the noise-smoothing properties of this assay, may also contribute to the smoothness of the ParD3 landscape.

TrpB:

The data for the TrpB landscape was downloaded from [25], and converted into a $20 \times 20 \times 20 \times 20$ `numpy` array. Similarly to ParD3, TrpB was evaluated using a growth-based assay [25], allowing us to more confidently compare the two on a biological basis. TrpB is a subunit of tryptophan synthase, which catalyses the synthesis of the amino acid tryptophan. Unlike ParD3, it does not contain any large, intrinsically disordered regions, and thus many residues may be coupled to maintain the folded structure and activity, contributing to epistasis and therefore landscape ruggedness. Meanwhile, in the case of TrpB this growth coupling is achieved through tryptophan synthesis which, following a similar argument to above, may "saturate" in its impact on measured fitness for

high-performing TrpB variants (i.e. once an adequate quantity of tryptophan is produced in the cell, further increase in enzyme efficiency is not rewarded by this assay).

GB1:

The data for the GB1 landscape was downloaded from [2], and converted into a $20 \times 20 \times 20 \times 20$ `numpy` array. GB1 is an immunoglobulin binding protein, whose structure, similarly to TrpB, is globular. Therefore it must also maintain the amino acid interactions that help support folded structure and activity, and consequently contribute to landscape ruggedness. The fitness values in this case were measured using a high-throughput binding assay [2]. The assay achieved a coverage of approximately 100 reads per variant, helping to reduce the noise of the data and therefore assay-related ruggedness, particularly compared to the TEV example.

Given the above arguments, we believe that the ordering of these landscapes by ρ_2 (Fig 4G) is justifiable. The assay used to measure TEV fitness was inherently noisy, alongside the fact that the presence of a catalytic triad in the enzyme could add considerable interaction effects and therefore ruggedness. For this reason, it makes sense that it would be ranked as more rugged than the remaining landscapes, rather than grouped with them—and indeed *all* ruggedness metrics investigated, except for the spectral entropy H, agreed with this conclusion. GB1 and TrpB, both globular proteins that were measured with less noisy assays, have justification to be grouped together (as they were by ρ_2, as well as roughness to slope and landscape R^2) and similarly to be below TEV. ParD3 was positioned at the bottom by all metrics, which aligns with our expectations given both its intrinsically disordered structure and the properties of the growth-based assay used to define its fitness.

Considering again Fig 4G and the metric comparisons in Fig 3F, we can see significant qualitative differences between metrics. In particular, from Fig 4G, the metrics considered can broadly be classified into three groups: Group 1 (decay rate ρ_2, roughness to slope, and landscape R^2) places TEV as the most rugged, ParD3 as the least, and then groups GB1 and TrpB close together at intermediate ruggedness. Group 2 (local epistasis and distance to local max) places ParD3 at the bottom, GB1 at intermediate ruggedness, but then groups TrpB with TEV at the ruggedness maximum. Group 3 (paths to the global maximum) in turn places TEV, GB1, and TrpB at the scale's maximum ruggedness—perhaps not surprising because this metric (as shown in Fig 3F) saturates very quickly as a landscape becomes more rugged. Considering the biological reasoning above we believe our metric (or more broadly those in Group 1) most aligns with the desirable properties of a ruggedness metric (i.e. as a measure of how sensitive a protein's fitness is to mutation): TrpB, whose growth-based measurements are comparatively smooth, falls below TEV, as does GB1. It remains unclear whether GB1 and TrpB *should* exhibit a similar ruggedness value (as seen for all Group 1 metrics)—this would best be answered with further experimental measurements or aug-

mentation of our comparisons with additional protein landscapes as they become available in the future.

F Alternative landscape models

In order to explore the effects of other landscape models on the behaviour of SLIDE, a Rough Mount Fuji model and a stochastic block model were implemented (S1 Fig). For each model, the optimal strategy choice across landscapes of low, mid and high ruggedness was compared. The optimum strategy for all models shifts from low base chance and low splitting to high base chance and high splitting as ruggedness increases, supporting the transferability of SLIDE across landscape models. In addition, the results also confirm that the expected trend in fitness decay curves holds across the different models, with higher ruggedness landscapes exhibiting higher rates of fitness decay in all cases. These results suggest that SLIDE is not specific to the NK model and general trends should hold across other landscape models and empirical landscapes.

G Alternative mutation models

While this work generally assumes an unbiased mutation spectrum, the assumption that transitions between all 20 amino acids are equally probable is biologically unrealistic. For the simulations described in the Biased Mutation Spectra and Landscape Perturbations section, the mutation model is extended. First, for Fig 6A-6C, the model is extended with codon effects, mapping each nucleotide triplet to a codon encoding a specific amino acid or a stop signal. To incorporate this, a landscape of N amino acids is represented by $3N$ nucleotides. We then use codon mapping to translate these sequences back into the amino acid chain, which is subsequently evaluated using the chosen empirical or synthetic landscape. For Fig 6D-6F this process is extended to the full set of 4^{3N} nucleotide genotypes. A critical detail in this process is the treatment of start and stop codons. For this study, we assume a protein becomes completely non-functional if a start or stop codon is modified within its sequence, and assign it the minimum observed fitness. Notably, this mapping does not itself alter the mutation kernel. Instead, it creates a class of landscapes with an inherent constraint: nucleotide configurations that encode the same protein must share identical fitness values. Because our analysis makes no prior assumptions regarding the fitness values within a landscape, it remains applicable to these models.

We considered three nucleotide mutation kernels. The uniform kernel assigns equal probability to each of the three possible substitutions. The directed kernels use species-specific nucleotide transition probabilities for *E. coli* and *A. thaliana*, derived from [41]. For each species, we also constructed a reversible weighted kernel. The directed matrix $\bar{D}$ was first symmetrised and then diag-

onally scaled,

$$\tilde{D} = Q\left(\frac{\bar{D} + \bar{D}^{\mathsf{T}}}{2}\right)Q, \tag{S3}$$

where Q was obtained by symmetric Sinkhorn scaling so that $\tilde{D}$ is symmetric and doubly stochastic. Thus, the weighted kernel retains non-uniform substitution weights while having a uniform stationary distribution. The original directed kernel remains row-stochastic and generally has a non-uniform stationary distribution.

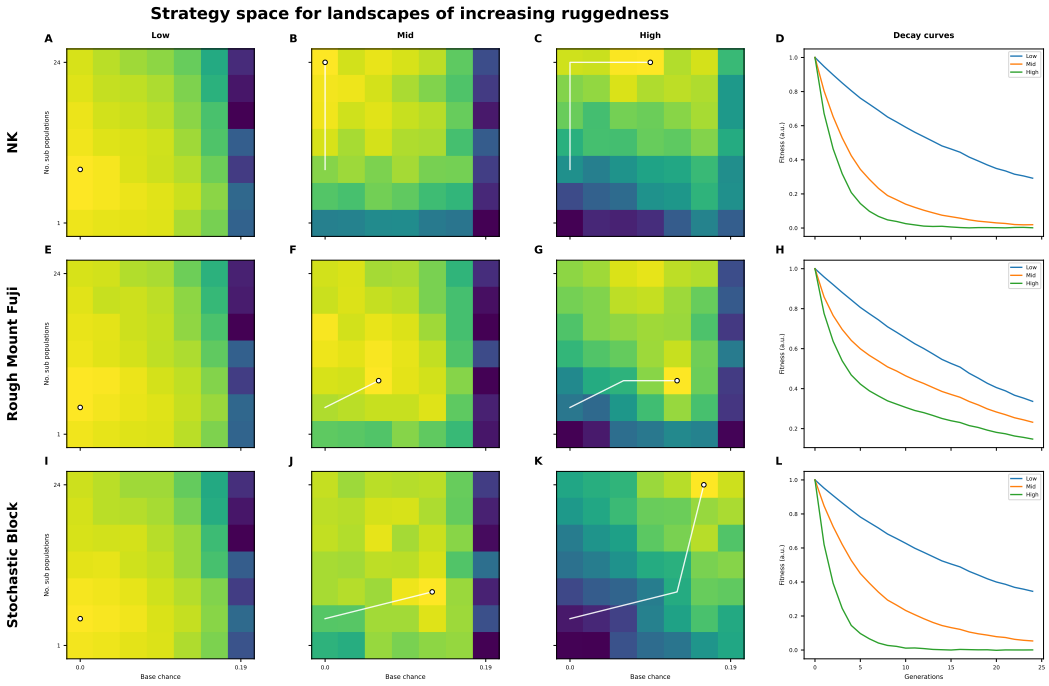
Comparing the starting point sweeps from Fig 4D-4F and Fig 6D-6F, the global fits shown in Fig 6D-6F do not converge as closely to the analytical ρ_2 values as they do for landscapes without codon mapping. To investigate this further, we plot the true decay curves against the theoretical curves using analytical ρ_2 values in S3 Fig. The figure shows that the fitted ρ_2^{fit} values align much more closely with the true decay curves than the theoretical values. This discrepancy arises because the eigenvalues are significantly more dispersed in these landscapes, particularly after codon mapping. Increased eigenvalue dispersion degrades the fit of the analytical ρ_2, which approximates a weighted sum of exponentials by a single exponential. As Eq (14) shows, the analytical ρ_2 yields a decay curve that acts as a strict lower bound, as is visible in the figure.

S1 Table. Summary of ruggedness and decay-rate notation. Column F_μ denotes metrics computed from the average population fitness decay, whereas column G_μ denotes metrics computed from the squared average population fitness decay.

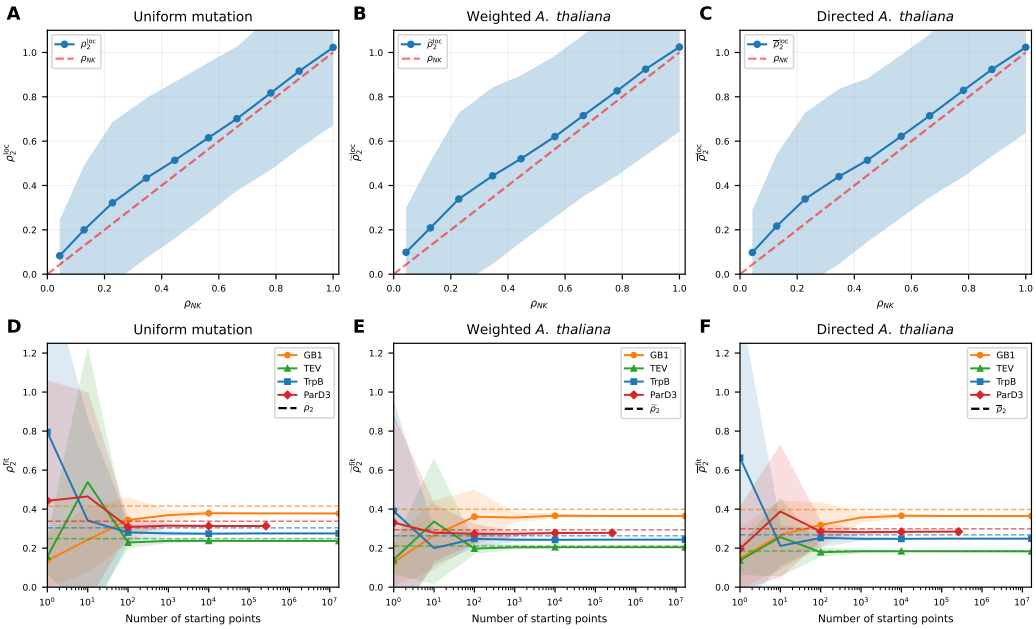
F_μ	G_μ	**Description**
ρ_1	ρ_2	Analytical decay rate.
ρ_1^{fit}	ρ_2^{fit}	Fitted decay rate.
	ρ_2^{loc}	Local fitted decay rate.
	ρ_2^{glob}	Global fitted decay rate.
ρ_{NK}		Closed-form NK ruggedness value.
ρ_{HoC}		Closed-form HoC ruggedness value.
	$\tilde{\rho}_2$	Decay rate for weighted, undirected mutation operator.
	$\bar{\rho}_2$	Decay rate for directed mutation operator.

S2 Table. Analytical decay rates under uniform, weighted, and directed mutation models computed on empirical landscapes.

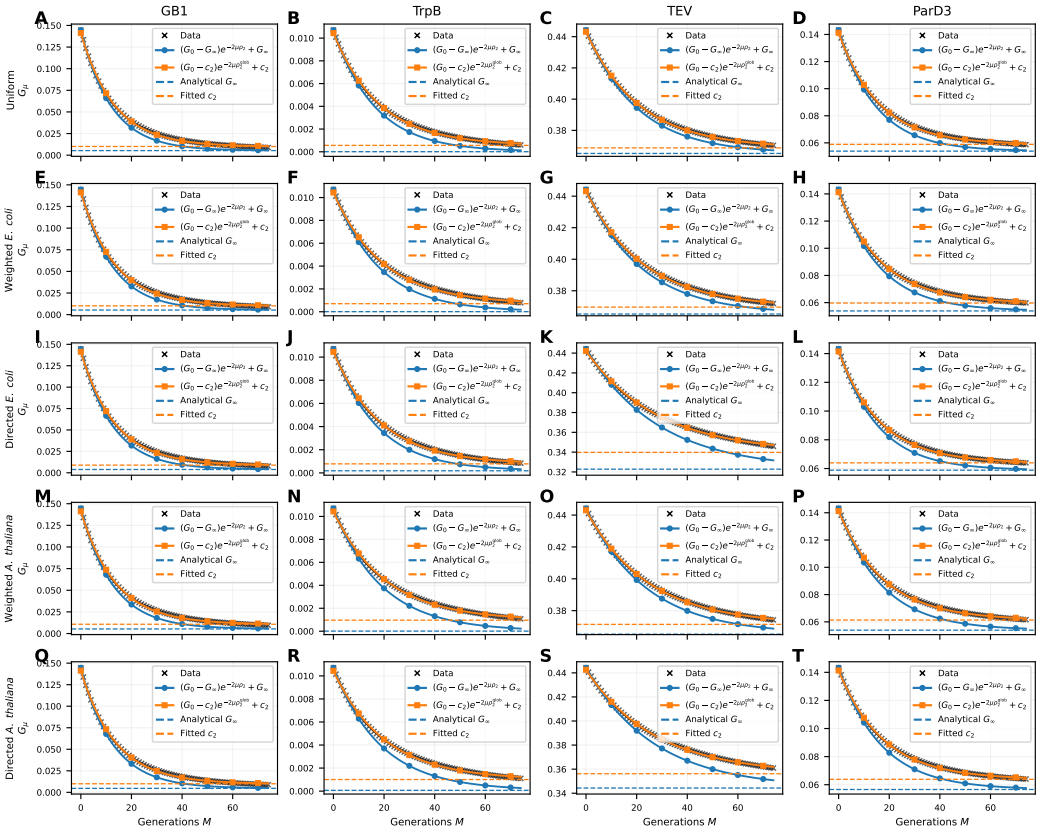
Mutation model	GB1	TrpB	TEV	ParD3
Uniform (ρ_2)	0.4155	0.3034	0.2485	0.3380
$E.\,coli$ weighted ($\tilde{\rho}_2$)	0.4067	0.2811	0.2274	0.3134
Bound (17), $E.\,coli$ weighted	0.3512	0.3386	1.9027	0.5423
$E.\,coli$ directed ($\bar{\rho}_2$)	0.4040	0.2924	0.1764	0.3244
Bound (21), $E.\,coli$ directed	0.3545	0.3424	1.9923	0.5588
$A.\,thaliana$ weighted ($\tilde{\rho}_2$)	0.3988	0.2627	0.2103	0.2939
Bound (17), $A.\,thaliana$ weighted	0.6341	0.6114	3.4360	0.9793
$A.\,thaliana$ directed ($\bar{\rho}_2$)	0.3978	0.2677	0.1850	0.2991
Bound (21), $A.\,thaliana$ directed	0.6208	0.5981	3.4017	0.9646



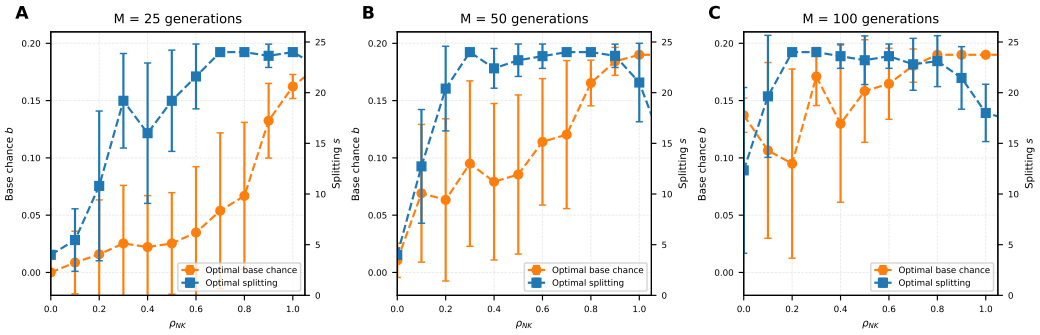
S1 Fig. Comparison of strategy space and decay curve trends across different landscape models. Heat maps show strategy spaces across low-, intermediate-, and high-ruggedness implementations of the NK, Rough Mount Fuji and stochastic block models. The white line tracks the optimal strategy as ruggedness increases. Each heat map represents an average across 300 starting points. Line plots on the right-hand side display the average fitness decay over the same landscapes represented in the heat maps.



S2 Fig. *A. thaliana* mutation-model version of Fig 6A-6F. A–C: Local squared-fitness decay rates inferred from synthetic NK landscapes under uniform mutation (ρ_2^{loc}), weighted *A. thaliana* mutation ($\tilde{\rho}_2^{\mathrm{loc}}$), and directed *A. thaliana* mutation ($\bar{\rho}_2^{\mathrm{loc}}$). Points show binned means across local starts and shaded regions show standard deviations. The dashed red line marks ρ_{NK}. **D–F:** Full-nucleotide empirical landscape estimates for GB1, TEV, TrpB, and ParD3 as a function of the number of starting points. Dashed horizontal lines show the analytical ρ_2 for each landscape and mutation model.



S3 Fig. Decay curves on empirical landscapes. Panels show full-nucleotide squared-fitness decay curves sampled from empirical landscapes, fitted decay curves using ρ_2^{glob}, and idealised analytical decay curves, which are computed as $(G_0 - G_\infty)\mathrm{e}^{-2\rho_2\mu} + G_\infty$. For uniform and weighted symmetric mutation models, G_∞ is computed from the unweighted landscape average. For directed asymmetric mutation models, G_∞ is computed using the stationary distribution of the nucleotide transition matrix. The mutation rate per generation is $\theta = 0.1/N$.



S4 Fig. Look-up table for NK landscapes with higher numbers of DE iterations. Parameters are as in Fig 5A, but with **A:** $M = 25$, **B:** $M = 50$ and **C:** $M = 100$ generations. Landscapes used a binary alphabet ($A = 2$), population size 1,200, mutation rate $\theta = 0.1$, and a 7×7 strategy grid over base chance $b \in \{0, 0.0317, 0.0633, 0.095, 0.1267, 0.1583, 0.19\}$ and splitting $s \in \{24, 20, 16, 12, 8, 4, 1\}$. The look-up was computed for 100 (N, K) pairs: $N \in \{10, 14, 18, 23, 27, 32, 36, 41, 45, 50\}$ with 10 sampled K values spanning low to high ruggedness for each N. For each (N, K) pair and DE horizon, the optimal strategy was selected after averaging over 25 replicate estimates per strategy point; the $M = 25$ grid used 200 random landscapes per (N, K) pair, while the $M = 50$ and $M = 100$ grids used 50. Points show the mean optimal parameter within each rounded $\rho_{NK} = (K + 1)/N$ bin, and error bars show the standard deviation of the optimal parameters across (N, K) pairs in that bin.

4.5 Limitations and Future work

Although this work substantially extends the practical utility of the previous chapter, several limitations remain. These, naturally, include most of the limitations stated in the previous chapter, such as the limitations of the NK model, and the available empirical combinatorially complete landscapes.

4.5.1 Experimental Viability

Likely, the major limitation is the experimental practicality of the experiment. In this work, we propose a calibration procedure to infer the ruggedness of an underlying fitness landscape, and then to use this inferred parameter to guide how a directed evolution experiment should be conducted. Conceptually, this procedure is both simple and elegant. It relies on observing the mean fitness of a population undergoing random, unbiased mutation. As the population mutates randomly, its distribution approaches uniformity over the landscape, and the mean fitness converges to the global mean. The rate of this convergence then determines the ruggedness of the landscape.

Whilst there is a strong case that this quantity is easier to infer than many alternative ruggedness metrics, which often require sequencing large portions of the landscape so that the metric may be computed directly, implementing this procedure in a real laboratory setting reveals several substantial challenges. First, the inferred ruggedness depends linearly on the mutation rate. As a result, if the true mutation rate is twice as large as assumed, the inferred ruggedness will also be twice the true value. In practice, precisely determining mutation rates in biological experiments is extremely difficult, making this a significant and hard-to-correct source of error.

The basic formulation assumes unbiased mutation. We also investigate biased mutation, with encouraging results in the cases tested, although the implications for experimental inference remain less clear. Beyond this, competition effects further complicate the situation. A common method for generating protein variants in directed evolution is to embed the gene encoding the protein into a bacterial host and use the bacteria to express the variants. Even if the inserted gene could be

mutated at the correct rate and with the desired distribution, the host bacteria themselves will continue to mutate and compete with one another. As a result, bacterial lineages that are incidentally more fit may outcompete others, making it extremely challenging to maintain a truly randomly mutating population. This effect is exacerbated if the protein under study has any fitness impact on the host organism, in which case natural selection will inherently bias the exploration of the protein landscape toward certain regions.

Whilst work can be done to make these experimental procedures more accurate, ideally, a ruggedness inference procedure should work with the constraints and features of real biological experiments rather than against them. This likely implies that future ruggedness metrics should sacrifice some theoretical elegance in favour of robustness and experimental feasibility. For example, real-world and simulated data from observing these population drift processes could be used to train a computational model to infer the underlying ruggedness parameters.

It is worth noting that several of these experimental limitations are less severe, or entirely absent, in the light of emerging experimental technologies, such as the MiPS platform [29] introduced in Chapter 1. For example, in the MiPS system cell lineages are physically separated and a constant supply of nutrients is available, eliminating resource competition and the associated complications. Some limitations are still notable when using such a system, notably the importance of calibrating the mutation rate correctly, but the process is overall more feasible with the MiPS platform.

On top of these additional experimental procedures, there is also the possibility of using data from the directed evolution experiment itself. For example, at a very basic level, any setup may be repeated with different strategies to determine which works best. This results in fewer total generations for the optimised strategy than if it were used throughout the experiment, but this may outweigh the costs of using a suboptimal strategy. Such an approach has the advantage of only assuming the ability to compare the results from different directed evolution experiments, and the capacity to run the directed evolution experiment itself. However, it does

assume enough total rounds of directed evolution to get meaningful information after assigning some of these rounds to exploratory purposes. Regardless, using at least some of the information from the directed evolution experiment itself to inform the choice of experimental parameters likely should form a part of future work.

4.5.2 NK landscapes

A second limitation concerns the use of the NK model in strategy selection. Although the Fourier spectrum provides a rich and high-dimensional characterisation of ruggedness, the NK model collapses ruggedness to a single parameter, K. As a result, it is not always clear which choice of (N, K) best corresponds to a given empirical Fourier spectrum, and different choices may lead to different strategy recommendations. This, conveniently, suggests a natural extension to the NK model: by summing together independent NK landscapes with differing ruggedness values, one can construct a model landscape with the desired spectral composition. This can make a class of landscapes that can, to a first approximation, represent all possible Fourier spectra. This is a clear advantage, but it doesn't address issues raised in the previous section, such as the distribution of fitness values, or the fact that the NK landscapes are, in a sense, homogeneous, and solving these issues would require separate alterations.

4.5.3 Ruggedness Metrics

Finally, limitations arise from the Fourier-based measure itself, despite its mathematical elegance and the fact that it is, in many respects, more informative than existing ruggedness measures. For example, the Fourier spectrum describes global structure, whereas evolution is an inherently local phenomenon; there are likely vast regions of the fitness landscape that will never be explored. It appears undesirable for a fitness measure to account for such distant, unreachable regions. While it is possible to augment the Fourier spectrum into a localised measure, the current lack of large-scale empirical landscapes or heterogeneous synthetic landscapes makes it difficult to quantify the advantages of such an approach.

Another issue is that the Fourier spectrum is sensitive to monotonic non-affine transformations, changes that alter the absolute values of the landscape without shifting the rank-order fitness. To illustrate how this becomes an issue, consider the optimisation of a protein to bind to a specific target. In a laboratory setting, a natural quantity to measure is the fraction of proteins bound to the target, P_{bound}. This value is a non-linear function of the binding energy ΔG and the temperature T, expressed by the relation:

$$P_{\text{bound}} = \frac{1}{1 + e^{\Delta G/RT}}$$

where R is the molar gas constant. Because of this non-linearity, the Fourier metric would register a change in the ruggedness of the fitness landscape simply by varying the temperature T, even if the underlying energetics (the ΔG values) remain identical. Consequently, the Fourier measure may capture artefacts of experimental conditions rather than the "intrinsic" biological properties of the landscape.

That being said, it is not immediately clear which transformation of the landscape is the correct one, and if such a notion of correct even applies. This ambiguity underscores a fundamental rift between Fourier-based measures and evolutionary dynamics. The Fourier spectrum is intimately linked to diffusion processes, and indeed diffusion processes are not invariant under monotonic transformations.

Furthermore, it is debatable whether a desirable ruggedness metric should be completely invariant to monotonic transformations. Evolution itself is not entirely invariant to such changes; slightly deleterious mutations can indeed fix in a population due to genetic drift, as the probability of fixation depends on the absolute scale of the fitness difference. In the context of optimisation, absolute scale also matters: the desirability of finding a global optimum depends on its magnitude relative to local optima. Perhaps it is worthwhile adopting a more risky strategy if reaching the global optimum is sufficiently rewarding! It is therefore natural that this sense of scale should feed back into measures of ruggedness.

Ultimately, a broader issue persists: a precise definition of "ruggedness" remains elusive. It is likely not a single property, but rather a collection of characteristics

where different metrics are appropriate for different contexts. This conceptual overlap is evidenced by the fact that various ruggedness metrics in the literature tend to be highly correlated in practice. In many respects, the Fourier measure is a simple and natural metric. If one attempts to construct a ruggedness measure based on the correlation, in the covariance sense, between a genotype and its neighbours across varying distances, the Fourier metric emerges naturally. This makes it an excellent starting point for demonstrating an inferable and functional ruggedness parameter of a fitness landscape.

The analysis of biased mutation provides a starting point, but establishing reliable inference under the biased, noisy mutational processes encountered experimentally remains an important direction for future work. This necessity is further highlighted by new experimental platforms, such as the MiPS platform being developed by the lab. The MiPS platform provides a rich stream of cell-specific data, which the Fourier metric is currently unable to utilise. A future metric capable of effectively ingesting such large-scale, high-resolution data will undoubtedly be more informative regarding the global properties of the landscape.

A man's dignity consists in his capacity to suffer for his ideas.

Thomas Mann *The Magic Mountain*

5 Antibody Shaping

Contents

5.1 Motivation 99
5.2 Summary 100
5.3 Contributions 102
5.4 Paper 102
5.5 Limitations and Future work 123

This chapter presents *ADIOS: Antibody Development via Opponent Shaping*, published in the *Proceedings of the 42nd International Conference on Machine Learning* (2025) [73].

5.1 Motivation

In previous chapters, the focus has been on using computational simulation to better understand how to accelerate evolution, finding better proteins in fewer steps. This chapter instead flips this objective, using similar methods to find ways to instead *inhibit* evolution.

In particular, the work presented in this chapter seeks to study how we may limit the rate of viral evolution. The motive behind limiting the rate of viral evolution is somewhat self-evident: viruses cause immense societal, human, and

economic damage, and this is in large part driven by their ability to rapidly adapt. However, despite their rapid rate of adaptation, current methods focus on treating only the current version of the virus.

This invites an enticing question: is it possible to design a therapy, such that the resulting evolutionary landscape is significantly harder for the virus to traverse, meaning the therapy remains more effective for longer? Is it possible, in theory, to design therapies which are not only resilient to viral evolution, but that actually direct it, guiding the virus to less harmful regions? These are the questions we seek to answer with **ADIOS**.

5.2 Summary

This chapter introduced **ADIOS: Antibody Development via Opponent Shaping**, a framework for designing antibodies that explicitly account for how a virus may evolve in response to them. At the core of ADIOS is a simplified simulation of viral escape: given an antibody, the virus repeatedly adapts to reduce binding, tracing out an evolutionary trajectory. By optimising antibodies against these trajectories rather than only the initial variant, ADIOS aims to identify antibodies that maintain binding in the model even as the simulated virus mutates.

Due to the enormous number of binding estimates needed to simulate an evolutionary trajectory, using high-fidelity binding simulators is infeasible. Instead, we opted to use the Absolut! model, which serves as a rapid, simplified model of antibody-antigen binding. We were able to port this model to the GPU, giving an approximately 10,000-fold speedup.

We can use this computationally efficient method to define fitness in this model via binding-based payoffs. Note there are two slightly distinct notions of fitness, with the viral fitness representing reproductive advantage in the current environment, and the antibody fitness being determined by what the researchers desire. The virus is rewarded for escaping antibody binding while maintaining binding to a host target, representing immune escape and retention of infectivity, respectively. We define this fitness to be based on only the current environment, reflecting the fact

that natural selection is somewhat myopic. The antibody's payoff favours binding to the virus, penalises binding to a human anti-target, and opposes viral binding to the host target. As we desire the antibody to continue to be effective, antibody fitness is the expected average of this payoff over the simulated escape trajectory.

Using this framework, we conduct several studies to understand how antibody design depends on the evaluation horizon used during optimisation. We find a clear trade-off: antibodies optimised only for the starting virus perform well initially but degrade rapidly as the virus adapts. In contrast, antibodies optimised over longer horizons, **shapers**, start with slightly lower binding, yet remain more effective over time. In other words, the "best" design depends on how far ahead one chooses to look, and long-horizon optimisation yields antibodies that are more resilient to viral escape.

A particularly striking result is that these shaper antibodies do more than simply maintain their own efficacy: they also shape the virus into becoming broadly easier to target. Viruses that evolve in the presence of shapers tend to accumulate mutations that make them more susceptible to a wide range of antibodies, when compared to viruses that evolve in the presence of myopic antibodies. This effect suggests that shaping may offer benefits beyond improving the stability of a single therapy.

We further test the robustness of these results by altering the evolutionary environment. When a second, conventional myopic antibody is present, introducing competing selective pressures, the qualitative behaviour remains unchanged: shapers still retain superior long-horizon performance and still guide viruses toward variants that are generally easier to bind in the model.

Overall, these findings provide compelling evidence that myopic or greedy optimisation may be suboptimal in settings where viral adaptation is inevitable. While the present work relies on simplified models of binding and escape, it demonstrates a promising proof of concept. As more accurate and biologically detailed simulators are developed, we believe the ADIOS framework has strong potential to contribute to the design of antiviral therapies that are not only effective against current variants, but resilient to the evolutionary future.

5.3 Contributions

This work is a joint first authorship between Sebastian Towers and Aleksandra Kalisz, under the supervision of Jakob Foerster, who originally came up with the ideas behind this paper. S. Towers provided the mathematical framing of the project, and wrote the initial code, with A. Kalisz's input on the design. In particular, the choice of binding simulator (Absolut!) is due to A. Kalisz. Philippe A. Robert, as the creator of Absolut!, was extremely helpful in the process of porting code over to JAX, and continued to offer extremely useful advice throughout the project. In terms of experiments, the experiments for Figures 3a, 3b and 6 were performed by A. Kalisz, and Figures 3c, 3d and 5 are by S. Towers. For Figure 4, the choice of viruses and their import into our package were performed by A. Kalisz, with the experiments then being run by S. Towers. A. Kalisz wrote the initial version of the paper, which was then jointly iterated on by A. Kalisz and S. Towers, with feedback from all other authors. The code release and blog posts were handled by A. Kalisz, and she is continuing this line of work. Alicia Higueruelo, Francesca Vianello, Ming-Han Chloe Tsai and Harrison Steel provided guidance and feedback throughout the project.

5.4 Paper

ADIOS: Antibody Development via Opponent Shaping

Sebastian Towers [*1 2] **Aleksandra Kalisz** [*1 2] **Philippe A. Robert** [3] **Alicia Higueruelo** [4] **Francesca Vianello** [5]
Ming-Han Chloe Tsai [6] **Harrison Steel** [2] **Jakob Foerster** [1 2]

Abstract

Anti-viral therapies are typically designed to target only the current strains of a virus, a *myopic* response. However, therapy-induced selective pressures drive the emergence of new viral strains, against which the original myopic therapies are no longer effective. This evolutionary response presents an opportunity: our therapies could both *defend against and actively influence viral evolution*. This motivates our method ADIOS: Antibody Development vIa Opponent Shaping. ADIOS is a meta-learning framework where the process of antibody therapy design, the *outer loop*, accounts for the virus's adaptive response, the *inner loop*. With ADIOS, antibodies are not only robust against potential future variants, they also influence, i.e., *shape*, which future variants emerge. In line with the opponent shaping literature, we refer to our optimised antibodies as *shapers*. To demonstrate the value of ADIOS, we build a viral evolution simulator using the Absolut! framework, in which shapers successfully target both current and future viral variants, outperforming myopic antibodies. Furthermore, we show that shapers modify the distribution over viral evolutionary trajectories to result in weaker variants. We believe that our ADIOS paradigm will facilitate the discovery of long-lived vaccines and antibody therapies while also generalising to other domains. Specifically, domains such as antimicrobial resistance, cancer treatment, and others with evolutionarily adaptive opponents. Our code is available at `https://github.com/olakalisz/adios`.

[*]Equal contribution [1]FLAIR, Foerster Lab for AI Research [2]Department of Engineering, University of Oxford, Oxford, UK [3]Department of Biomedicine, University of Basel, Basel, Switzerland [4]Isomorphic Labs, London, UK [5]Exscientia, Oxford, UK [6]Epsilogen Ltd., London, UK. Correspondence to: Sebastian Towers <sebastian.towers@eng.ox.ac.uk>, Aleksandra Kalisz <aleksandra.kalisz@eng.ox.ac.uk>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

1. Introduction

Designing effective therapies to fight off viral pathogens is crucial for limiting their devastating social and economic costs (Nandi & Shet; Orenstein & Ahmed, 2017; Samsudin et al., 2024; Faramarzi et al., 2024). However, traditional design approaches only target the *current* variant of a virus. Although this *myopic* design approach may yield therapies with high initial efficacy, it fails to account for viral adaptation, leaving treatments vulnerable to becoming ineffective over time (Weisblum et al., 2020; Doud et al., 2018; Lee et al., 2019; Dingens et al., 2019; Greaney et al., 2021).

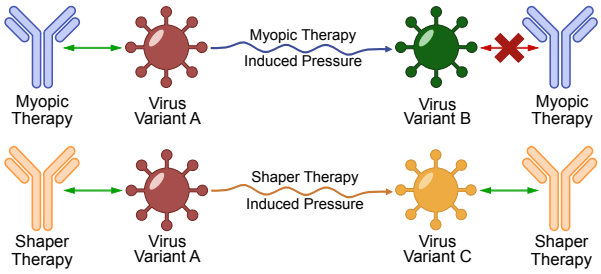


Figure 1. **Myopic Therapies Become Ineffective Over Time.** Initial virus (variant A) evolves in response to the evolutionary pressures induced by therapies, resulting in new variants. In case of traditional *myopic* therapies (top), the new emerging variants (variant B) are often therapy resistant. In contrast, ADIOS designs *shapers* or shaper therapies (bottom) which remain effective and steer the viral evolution towards less harmful variants (variant C).

The COVID-19 pandemic starkly illustrated the challenges of adaptive viruses. While the rapid development of vaccines was a remarkable achievement, concerns quickly arose about their long-term efficacy against new emerging COVID variants (Carabelli et al., 2023; Hu et al., 2021). For example, the B.1.351 variant demonstrated that the vaccine loses its effectiveness against new strains (Madhi et al., 2021). This underscores the need for approaches that consider both the current *and future* efficacy of a designed therapy.

The virus inevitably adapts in response to selective pressures imposed by our therapies, i.e., we *influence* the viral evolution (Chéron et al., 2016; Meijers et al., 2022). Our work turns this influence in our favour, designing therapies that steer the virus towards less dangerous variants, see Figure 1.

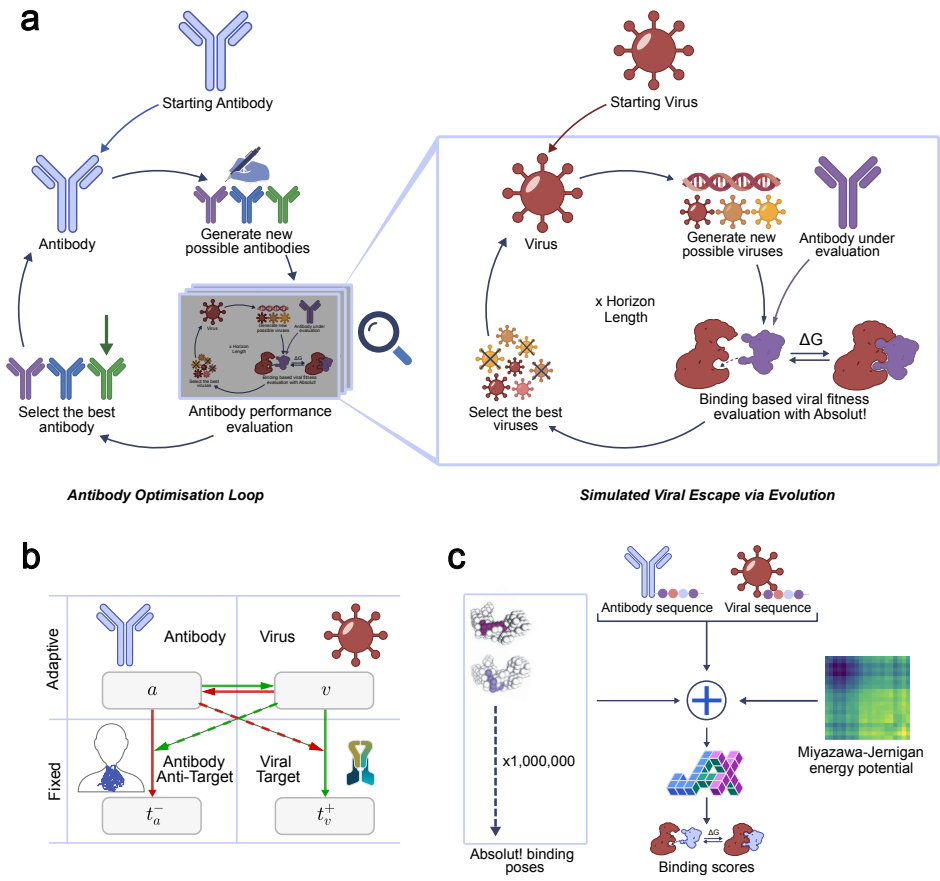


Figure 2. **Main Components of ADIOS. a** The ADIOS framework. In the *Antibody Optimisation Loop* (i.e., *outer loop*), we optimise the antibody to perform well against current and future virus variants; thus influencing the viral evolution. We approximate the future variants through our *Simulated Viral Escape via Evolution* (i.e., *inner loop*) where the viruses evolve to escape from the current antibody over a given horizon length. **b** The payoffs of the antibody and virus. Red arrows indicate binding interactions that players aim to minimise, while green arrows represent those they aim to maximise. The antibody optimises for binding to the virus while avoiding its anti-target. In this zero-sum game, the antibody's optimisation indirectly counters the virus's binding to its target, see Equation 1. **c** Binding simulator. Our JAX (Bradbury et al., 2018) implementation of the binding calculation uses binding poses generated by Absolut! (Robert et al., 2022) and the Miyazawa-Jernigan energy potential matrix (Miyazawa & Jernigan, 1999).

To achieve this we utilise principles from *opponent shaping* (Foerster et al., 2018), a multi-agent reinforcement learning framework that allows agents to both *anticipate* and *influence* the future policies of other agents in their environment. This approach, exemplified by methods such as Learning with Opponent-Learning Awareness (Foerster et al., 2018) and Model-Free Opponent Shaping (Lu et al., 2022), allows agents to consider not only their current performance but also the consequences of their actions on their opponents' future behaviour.

Building on these principles, we introduce ADIOS: Antibody Development vIa Opponent Shaping. Antibodies are immune system proteins that bind to pathogens such as viruses. While naturally produced by the body, it is also possible to design and synthetically produce antibodies as therapies. ADIOS frames the interaction between antibodies and viruses as a *two-player zero-sum game*. In this game, the antibody's payoff is primarily determined by its binding strength to the virus, while the virus has the opposite payoff (Figure 2b). Although our framework can use any binding model in principle, in this work we build on the Absolut! framework (Robert et al., 2022) to estimate the binding strength of protein-protein interactions. To improve computational efficiency, we reimplement parts of Absolut! in JAX (Bradbury et al., 2018), allowing GPU acceleration and a 10,000-fold speedup over the original implementation (Figure 2c).

We use this game to model *viral escape* - the process through which mutations allow a virus to evade a host's immune system (Lucas et al., 2001). Following a meta-learning approach, ADIOS implements two nested optimisation loops, an *inner loop* and an *outer loop* (Figure 2a). In the inner loop, we simulate viral escape via evolution, where the virus adapts to the current antibody by repeatedly finding approx-

imate best responses that decrease binding strength. The outer loop uses a genetic algorithm to optimise the antibodies to be effective across viral evolutionary trajectories, resulting in antibodies we call *shapers*. This is in contrast to only optimising for binding to the initial virus, which results in *myopic* antibodies.

Importantly, our simulations show that shapers not only outperform myopic antibodies in long-term efficacy but also demonstrate the ability to shape viral evolution. Moreover, shapers guide viruses toward variants that are more susceptible to binding to a broad spectrum of antibodies, not just the shaper that induced the given viral evolution, providing insights into the scalability and potential for practical deployment of ADIOS. Our study also explores the trade-offs between the effectiveness of shapers and the computational resources required for their optimisation. Finally, we present an explainability analysis of the key features that distinguish shapers from myopic antibodies.

Our key contributions include:

- ADIOS: A framework that brings opponent shaping to antibody design to address viral escape.
- A GPU-accelerated JAX implementation of the binding simulator Absolut!, achieving a 10,000x speedup.
- An open-source instantiation of ADIOS applied to antibody design for both the dengue virus and three other viruses using our JAX implementation.
- Empirical results showing ADIOS-optimised shapers both significantly outperform myopic antibodies by limiting long-horizon viral escape and guide viral evolution towards variants that can be more easily targeted.
- Analysis of computational trade-offs in shaping horizons, providing practical guidance for deploying ADIOS in compute-constrained settings, e.g. due to more realistic binding simulators.
- Interpretability analysis into how antibody shapers influence viral escape, which could, in principle, provide inspiration to antibody designers.

While our results provide a promising proof of concept, they are based on simplified models of binding and viral escape. However, we believe that as more sophisticated simulators emerge, the ADIOS framework has the potential to significantly impact future antiviral therapy design.

2. Related Work

Antibody Design: Antibodies are essential components of the immune system that bind to unique identifiers (antigens) present on pathogens, including viruses, to identify and neutralise them. While natural antibodies emerge through an immune response, it is also possible to *design* antibodies for use as therapies. Recent work has made significant progress in computational antibody design (Cutting et al., 2024; Zambaldi et al., 2024; Bennett et al., 2024). The common approaches to antibody design utilise energy-based antibody optimisation methods (Li et al., 2014; Adolf-Bryfogle et al., 2018; Pereira et al., 2024), sequence-based language models (Liu et al., 2020; Saka et al., 2021) or structure-based approaches relying on GNNs (Jin et al., 2022) and diffusion models (Martinkus et al., 2024).

In contrast to these works, we are not interested in generating better antibody design methods immediately, but rather in how we should make new methods in the future to account for our effect on evolving viruses.

Predicting Viral Escape: Recent machine learning methods have demonstrated success in predicting future viral strains (Shanker et al., 2024; Wang et al., 2023; Nie et al., 2025). EVEscape (Thadani et al., 2023) decomposes the likelihood of a mutation into three parts: maintaining fitness, accessibility to antibodies, and disrupting binding, demonstrating success through retrospective identification of COVID variants. Han et al. (2023) take a different approach by modelling viral evolution through simulated fitness landscapes. Unlike these methods, ADIOS models the antibody influence on viral evolution, enabling both the simulation of viral escape trajectories and the optimisation of antibodies to minimise viral escape.

3. Background

Antibody Binding Simulators: In our setting, the interaction between antibodies and viruses is characterised by their binding strength $B(\cdot, \cdot)$ - a measure of how strongly the two "attach" to each other through molecular forces. Molecular dynamics simulations offer high accuracy but are computationally intensive (Hollingsworth & Dror, 2018). Sequence-based ML models (Mason et al., 2021; Lim et al., 2022; Ruffolo et al., 2023; Yan Huang et al., 2022) provide faster alternatives but struggle to generalise beyond their training distribution, making them unsuitable for exploring novel viral mutations. To evaluate $B(\cdot, \cdot)$ we use the Absolut! framework (Robert et al., 2022). Absolut! offers a balance between speed and generalisation by modelling binding through discretised protein structures. It focuses on the CDRH3 region of the antibody, the most variable portion that primarily determines binding specificity (VanDyk & Meek, 1992). For each antibody-antigen pair, Absolut! enumerates possible binding poses and computes their energy using the Miyazawa-Jernigan potential (Miyazawa & Jernigan, 1999). The binding strength $B(\cdot, \cdot)$ is then defined as the negative of the lowest binding energy, see Figure 2c and Appendix D for details.

Opponent Shaping: A multi-agent reinforcement learning framework which allows agents to anticipate and influence the future policies of other agents in their environment. Learning with Opponent-Learning Awareness (LOLA) (Foerster et al., 2018) introduced this concept by having LOLA

Algorithm 1 Ev$(\hat{v}, a)$: Simulated Viral Escape (Inner Loop)

Input: virus $\hat{v}$, antibody a, horizon H, population size P, inverse temperature β
Output: Sampled viral trajectory $\hat{\boldsymbol{v}} = [\hat{v}^0, \hat{v}^1, \ldots, \hat{v}^H]$
$\hat{v}_0 \leftarrow \hat{v}$
for $i = 0$ **to** $H - 1$ **do**
  for $k = 1$ **to** P **do**
    $v_k^i \leftarrow \hat{v}^i \oplus$ *Mutation*
  end for
  $p(v) \leftarrow \mathbb{P}(v = v_k^i) \propto \exp(\beta R_{\mathrm{v}}(v_k^i, a))$ *// See Eq. 1*
  $\hat{v}^{i+1} \sim p(v)$
end for
$\hat{\boldsymbol{v}} \leftarrow [\hat{v}_0, \ldots, \hat{v}_H]$
return $\hat{\boldsymbol{v}}$

Algorithm 2 Antibody Optimisation (Outer Loop)

Input: antibody $\hat{a}$,virus $\hat{v}$, horizon H, population size P_a, steps N
Output: Trajectory of antibodies $\hat{\boldsymbol{a}} = [\hat{a}^0, \hat{a}^1, \ldots, \hat{a}^N]$
$\hat{a}_0 \leftarrow \hat{a}$
for $i = 0$ **to** $N - 1$ **do**
  $a_1^i \leftarrow \hat{a}^i$
  for $k = 2$ **to** P_a **do**
    $a_k^i \leftarrow \hat{a}^i \oplus$ *Point Mutation*
  end for
  $\hat{a}^{i+1} \leftarrow \arg\max_k \mathbb{E}\left[F_{\hat{v}}^H(a_k^i)\right]$ *// See Eq. 2*
end for
$\hat{\boldsymbol{a}} \leftarrow [\hat{a}_0, \ldots, \hat{a}_N]$
return $\hat{\boldsymbol{a}}$

agents optimise against anticipated opponent updates rather than static opponent policies. They achieved this through an augmented value function that accounts for the opponent's learning step.

Meta-learning is a set of methods for optimising a learning process itself, "learning to learn". In multi-agent systems, this concept extends to learning about and influencing how other agents learn. Model-Free Opponent Shaping (M-FOS) (Lu et al., 2022) showcases this idea by using gradient-free optimisation to learn meta-policies that accomplish long-horizon opponent shaping. Our approach follows a similar principle, using evolutionary optimisation to shape viral escape trajectories.

4. Method

ADIOS frames antibody design as a two-player game between an antibody shaper agent and a naive virus agent, building on principles from opponent shaping (Figure 2). We present our method in three parts:

In Section 4.1, we introduce the **virus-antibody game**, defining the action spaces and payoffs for both players. Section 4.2 describes our **simulated viral evolution** process, modelling how viruses evolve to escape a given antibody. Finally, Section 4.3 presents our **antibody optimisation** approach, which optimises the antibody shapers in a way that accounts for *future viral mutations* and learns to *influence* viral evolution away from escape.

4.1. Virus-Antibody Game

We formalise the interaction between antibodies and viruses as a *two-player zero-sum game*. In this game, two players – the virus and the antibody – play a game where one player's gain is the other's loss. The game is defined by the set of actions available to each player and their respective payoffs. The players' actions are represented by their amino acid sequences. The sequences are of an antigen protein for the virus and a fragment of a hypervariable region of the heavy chain for the antibody.

We define the set of 20 amino acids as $\mathbb{A}$. Let N_v be the virus sequence length and N_a be the antibody sequence length. So an action of the virus is $v \in \mathbb{A}^{N_v}$, and an action of the antibody is $a \in \mathbb{A}^{N_a}$. Let $B : \mathbb{A}^{N_v} \times \mathbb{A}^{N_a} \to \mathbb{R}$ be our binding function, which measures the strength of the binding between the antibody and the virus with increasing values corresponding to stronger binding[1].

The payoff structure is designed to capture the biological incentives of both players: the antibody aims to bind strongly to the virus while avoiding binding to human proteins (an *anti-target*), whereas the virus seeks to evade antibody binding while maintaining its ability to bind to host cell receptors (a *binding target*). Mathematically, we define the antibody's payoff R_a as:

$$R_a(v, a) = B(v, a) - B(t_a^-, a) - B(v, t_v^+) \qquad (1)$$

where $B(v, a)$ represents the binding strength between the virus v and antibody a, t_a^- is the antibody's anti-target, and t_v^+ is the virus's binding target. The virus's payoff R_v is simply the negative of the antibody's payoff: $R_v(v, a) = -R_a(v, a)$, see Figure 2b. This formulation also ensures that neither player can adopt an overly simplistic strategy: the virus can't become entirely inert without losing its ability to infect host cells, and the antibody can't become universally "sticky" without binding to the human protein, a "false positive", potentially causing the immune system to attack the human body. We give the full Markov Decision Process (MDP) definition $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}^{\mathrm{v}}, \mathcal{A}^{\mathrm{a}}, P, R, \mu \rangle$ in Appendix E.

4.2. Simulated Viral Escape via Evolution

We model the viral escape as a virus $\hat{v}$ naively evolving for H steps in response to some fixed antibody a. The *simulated viral escape via evolution*, see Figure 2a and Algorithm 1, is defined as follows. Given a starting virus $\hat{v}$, the fixed antibody a induces a distribution Ev$(\hat{v}, a)$ over sequences

[1]This is opposite to binding *energies*, which are smaller for stronger binding.

of viruses $\hat{\boldsymbol{v}} = [\hat{v}^0, \hat{v}^1, \hat{v}^2, \dots \hat{v}^H]$, where $\hat{v}^0 = \hat{v}$ and H is the chosen horizon length. We write $\hat{\boldsymbol{v}} \sim \text{Ev}(\hat{v}, a)$ to denote this relationship.

We define the process of generating the escape trajectories inductively. In generation i, we have a virus $\hat{v}^i$. We generate a population of viruses $v_1^i, v_2^i \dots v_P^i$ by duplicating $\hat{v}^i$ P times, then randomly applying mutations such that on average there is one amino acid mutation per viral sequence:

$$v_k^i = \hat{v}^i \oplus \textit{Mutation}$$

In our experiments $P = 15$. For every virus in the population, we evaluate its fitness given by $R_v(v_k^i, a)$. We then sample a new virus $\hat{v}^{i+1}$ based on the fitness values, in particular:

$$\mathbb{P}(\hat{v}^{i+1} = v_k^i) \propto \exp(\beta R_{\text{v}}(v_k^i, a))$$

With duplicates in the population being considered distinct, so that the likelihood of a particular variant increases with the number of duplicates. Furthermore, β is a constant which reflects how random the selection process is, with $\beta \to \infty$ reflecting deterministic max-fitness selection. After H generations, a full escape trajectory $\hat{\boldsymbol{v}} = [\hat{v}^0, \hat{v}^1, \hat{v}^2, \dots, \hat{v}^H]$ has been generated and the simulated viral escape process ends.

4.3. Antibody Optimisation

We define the antibody fitness $F_{\hat{v}}^H(a)$ such that it represents the *true* objective of the antibody, which accounts for the viral escape. Given a horizon H and starting virus $\hat{v}$, the antibody fitness is:

$$F_{\hat{v}}^H(a) = \mathbb{E}_{\hat{\boldsymbol{v}} \sim \text{Ev}(\hat{v}, a)} \left[\frac{1}{H+1} \sum_{i=0}^{H} R_a(\hat{v}^i, a) \right] \quad (2)$$

Note that if $H = 0$ this fitness defaults to a naive antibody payoff that ignores viral escape, i.e., $F_{\hat{v}}^0(a) = R_a(\hat{v}, a)$. We refer to this as the *myopic* objective.

To optimise both shapers and myopic antibodies, we employ Monte Carlo simulations to estimate the antibody fitness, combined with an evolutionary optimisation algorithm. We refer to this process as the *antibody optimisation loop*, see Figure 2a and Algorithm 2. In meta-learning terms, this is the *outer loop* or the *meta-loop*, contrasting to the *inner loop*, which is the *simulated viral escape via evolution*.

Given a starting antibody $\hat{a}$, a starting virus $\hat{v}$ and a viral escape horizon H, the antibody optimisation process generates a trajectory of antibodies $\hat{\boldsymbol{a}} = [\hat{a}^0, \hat{a}^1, \hat{a}^2, \dots, \hat{a}^N]$, where N is the number of antibody optimisation steps (i.e., meta-steps). In the trajectory, $\hat{a}^0 = \hat{a}$ is the starting antibody and $\hat{a}_N$ is the final optimised antibody. This optimisation could, in principle, start from any antibody, but for simplicity we opt to start from purely random antibodies, meaning $\hat{a}$ is random. In most of our experiments $N = 30$.

At the start of antibody optimisation step i, we have an antibody $\hat{a}^i$. We first generate a population of P_a antibodies $[a_1^i, a_2^i \dots a_{P_a}^i]$ by taking both the antibody $\hat{a}^i$ and $P_a - 1$ copies of it, with each copy having exactly a single random mutation in the amino acid sequence of $\hat{a}^i$. For our experiments, $P_a = 40$.

We then sample their fitness values $F_{\hat{v}}^H(a_j^i)$ with a fixed number η of Monte Carlo roll-outs, i.e., we sample η independent viral escape trajectories, each with horizon H viral escape steps. We found $\eta = 5$ to be sufficient. Finally, we select $\hat{a}^{i+1}$ to be the best-performing antibody:

$$\hat{a}^{i+1} = \arg\max_k \mathbb{E}\left[F_{\hat{v}}^H(a_k^i)\right]$$

Once the final optimised antibody $\hat{a}^N$ is generated, a full optimisation trajectory is complete, $\hat{\boldsymbol{a}} = [\hat{a}^0, \hat{a}^1, \hat{a}^2, \dots, \hat{a}^N]$, and the antibody optimisation process finishes.

5. Experimental Setup

5.1. Absolut! Speedup

To meet the computational demands of our opponent shaping approach, which requires rapid evaluation of numerous antibody-virus interactions, we reimplement the core binding calculation of Absolut! (Robert et al., 2022) using JAX (Bradbury et al., 2018), a framework that facilitates GPU-accelerated computation (Figure 2c). Our efficient JAX implementation and the GPU acceleration results in a 10,000-fold speedup compared to the original implementation, see Table 1.

	Absolut!	Absolut! + JAX
Hardware	Apple M2 Max	Nvidia A40
Time/Antigen (s)	1.8	$\mathbf{2.1 \times 10^{-4}}$

Table 1. Comparison of the time taken to compute a single binding query between the original implementation of Absolut! (Robert et al., 2022) and our reimplementation of Absolut! in JAX (Bradbury et al., 2018). The original Absolut! implementation runs on CPU only, hence the difference in evaluation hardware.

5.2. Dengue Virus

We use the antigen protein from the Dengue Virus for our main experiments, specifically, the structure with Protein Data Bank (PDB) code 2R29 (Berman et al., 2000; Lok et al., 2008). First, Absolut! processes this structure to generate binding-relevant information, which is then used by our JAX implementation (details given in Appendix D). In the viral escape step, we mutate only the amino acid sequence of

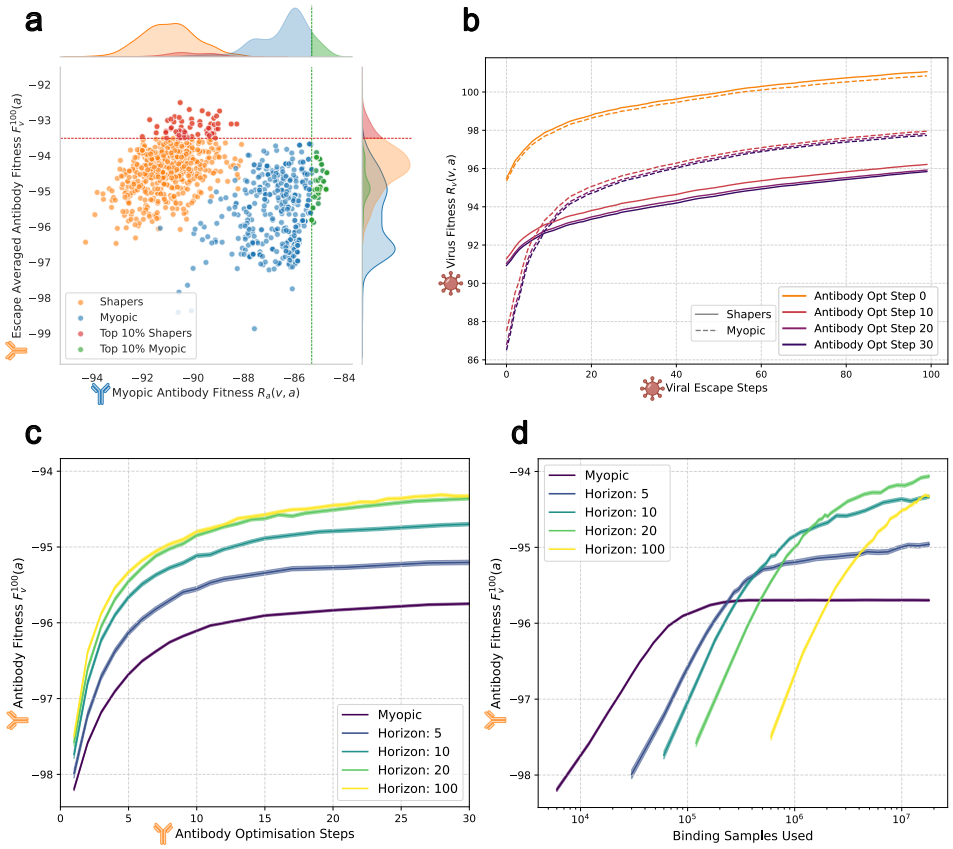


Figure 3. **Shapers Outperform Myopic Antibodies. a** Distribution of antibody shapers optimised with horizon $H = 100$ (orange) vs. myopic antibodies distribution (blue). We highlight the top 10% shapers with respect to $F_v^{100}(a)$ in red, and the top 10% myopic antibodies with respect to $R_a(v, a)$ in green. The x-axis is the myopic antibody fitness $R_a(v, a)$ and the y-axis is the escape averaged antibody fitness for $H = 100$, i.e., $F_{\tilde{v}}^{100}(a)$. Higher values on both axes indicate better performance. **b** Viral escape curves (inner loop performance) for different steps of the antibody optimisation process (outer loop) for antibody shapers optimised with horizon 100 (solid lines) and myopic antibodies (dashed lines). The lighter lines indicate early antibody optimisation steps, and the darker lines show the later steps. The x-axis shows the evolutionary steps of viral escape. The y-axis represents the virus fitness/payoff $R_v(v, a)$, where higher values indicate better virus fitness (and lower values denote better antibody performance, so lower is better for us). **c, d** Antibody optimisation learning curves (outer loop performance) for a varying horizon length. The x-axis shows the antibody optimisation steps, i.e., meta-steps (**c**) or the number of samples from the binding simulator (**d**). The y-axis shows antibody fitness $F_v^{100}(a)$. Error bars correspond to the standard error. Higher values indicate better performance.

the dengue envelope antigen, which is composed of $N_v = 97$ amino acids, and do not consider other components of Dengue Virus. Importantly, we assume that the structure of the antigen does not significantly change over the course of viral escape. All experiments but the ones we discuss in Section 6.3 and Appendix A use the dengue virus.

5.3. Additional Viruses and Bacterium

To demonstrate robustness of the experimental results we achieve with ADIOS on dengue virus we conduct additional experiments with three other viruses and one bacterium. The three viral antigens we use are: West Nile Virus, PDB code 1ZTX (Nybakken et al., 2005); Influenza Neuraminidase Virus, PDB code 4QNP (Wan et al., 2015) and MERS-CoV Virus with PDB code 5DO2 (Li et al., 2015). Furthermore, we show that ADIOS can be easily applied to other pathogens, such as bacteria, too. We perform an extra experiment with the Clostridium Difficile Bacterium, PDB code 4NP4 (Orth et al., 2014).

6. Results

6.1. Shapers vs. Myopic Antibodies

We validate the effectiveness of the antibody shapers in optimising the escape-averaged antibody fitness function $F_v^H(a)$ compared to myopic antibodies that only respond to the current virus v. For our shaper antibodies, we select a long horizon of $H = 100$ to capture extended viral

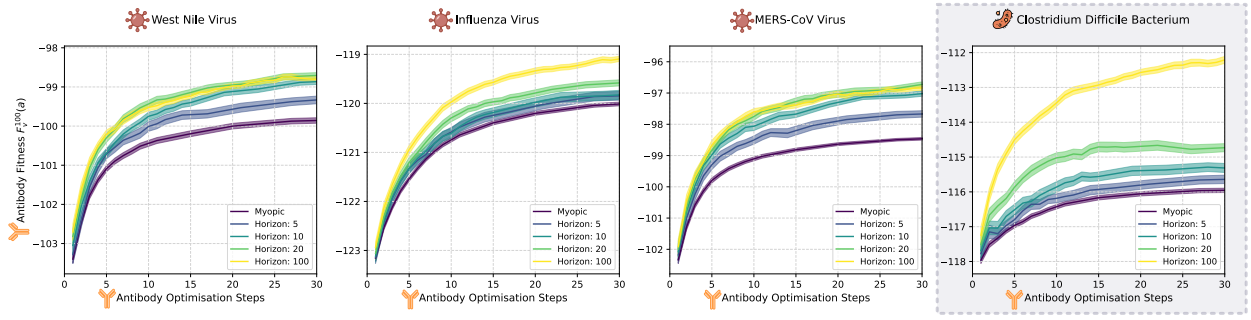


Figure 4. **Shapers Outperform Myopic Antibodies on Other Viruses and a Bacterium.** Antibody optimisation learning curves (outer loop performance) for varying horizon lengths across three viruses - West Nile Virus, Influenza Virus, and MERS-CoV - as well as the bacterium Clostridium Difficile. The x-axis shows the antibody optimisation steps, i.e., meta-steps and the y-axis shows the antibody fitness $F_v^{100}(a)$. Raw fitness values are dependent on Absolut! scores and are relative to the specific antigen, i.e., absolute values should not be compared between different viruses or bacterium but rather the overall trends. Error bars correspond to the standard error. Higher values indicate better performance. Full set of ADIOS results for these four pathogens is provided in Appendix A.

escape trajectories. Both shapers and myopic antibodies are optimised for $N = 30$ steps. Figure 3a presents the performance distributions of shapers and myopic antibodies under both objective functions.

Our results demonstrate a clear advantage of shapers in the escape-averaged objective $F_v^{100}(a)$. The mean of the shapers distribution significantly exceeds that of the myopic distribution, as evident from the marginal density plot in Figure 3a. Notably, none of the myopic antibodies outperform *any* of the top 10% of shapers in this long-term objective. However, there is a trade-off between short-term and long-term optimisation. While shapers do better on the escape-averaged objective, they underperform on the myopic objective $R_a(v, a)$.

We next examine the influence of antibody shapers on viral escape trajectories, comparing $H = 100$ shapers with myopic antibodies, both optimised for $N = 30$ steps. Figure 3b illustrates the viral escape curves induced by both antibody types at different stages of their optimisation process. We first complete the antibody optimisation process, saving antibodies generated at steps $0, 10, 20$, and 30. For each of these optimisation steps, we then simulate viral escape over $H = 100$ evolutionary steps using the corresponding saved antibodies. The presented viral escape curves are averages derived from multiple simulations.

At the outset of the antibody optimisation process (step 0), both the shapers and the myopic antibodies induce similar escape curves, an expected outcome given their initialisation from random antibody sequences. However, as we examine antibodies from later optimisation steps, we observe diverging trends. Myopic antibodies cause the viral fitness to be lower in the initial escape steps, outperforming the shapers. After about 10 escape steps, corresponding to ≈ 10 viral mutations, the two antibody types perform similarly. Beyond that, shapers demonstrate superior results in later escape stages, more effectively preventing viral escape.

These results show that as the antibody optimisation process progresses, shapers learn to influence viral trajectories in a way that minimises long-term viral escape, albeit at the cost of initial performance. While myopic antibodies may offer better immediate control, shapers provide more sustained effectiveness against evolving viral populations.

6.2. Antibody Shapers with Varying Horizons

Finally, we investigate the impact of varying horizons H on the optimisation process of antibody shapers. We optimise myopic antibodies and shapers using horizons $H = \{5, 10, 20, 100\}$ for $N = 30$ steps. To evaluate these antibodies against a consistent "true" objective, we simulate viral escape over $H = 100$ steps for each antibody, regardless of the horizon used during its optimisation. Figure 3c presents these results, demonstrating that shapers optimised with longer horizons H consistently yield better performance throughout all steps of the optimisation process.

However, the number of antibody optimisation steps does not accurately reflect the computational or experimental cost of optimisation. Each simulation of viral escape requires a number of binding samples that increases linearly with the horizon length H. Yet, shorter horizon antibodies optimise an objective that diverges further from our "true" antibody objective $F_v^{100}(a)$. Due to this trade-off, we observe that the optimal training horizon varies depending on the available computational budget.

To illustrate this trade-off, we conduct an additional experiment shown in Figure 3d. Here, instead of fixing the number of optimisation steps N, we constrain the total number of binding samples - queries to our binding strength simulator used to evaluate all antibody and virus payoffs throughout the optimisation process - to be constant across different horizons. This approach provides a performance comparison that accounts for the computational resources

necessary across varying horizon lengths. Interestingly, $H = 20$ shapers perform strongly, nearly matching the performance of those optimised with horizon $H = 100$ for a given number of antibody optimisation steps, and far exceeding it when accounting for the differing computational cost. This suggests that using a cheaper, shorter-horizon proxy for the true antibody objective $F_v^{100}(a)$ can yield substantial benefits.

More generally, we find that the optimisation horizon significantly influences the performance of antibody shapers. While longer horizons lead to better long-term performance, the optimal horizon length is dependent on the available computational resources. Thus, it is important to consider the balance between computational cost and the fidelity of the optimisation objective when designing antibodies for long-term effectiveness against evolving viral populations.

6.3. Antibody Shapers for Other Viruses and Bacterium

To test whether the shaping effects observed on dengue generalise, we evaluate ADIOS on three additional viruses: West Nile, Influenza, and MERS-CoV; as well as the Clostridium Difficile bacterium (Figure 4, Appendix A). In all cases, we observe consistent trends: shapers outperform myopic antibodies in escape-averaged fitness, confirming that ADIOS can successfully achieve shaping across diverse pathogens.

For West Nile virus and MERS-CoV, $H = 100$ shapers appear to perform worse than shorter-horizon $H = 20$ antibodies, see Figure 4. However, given that all antibodies are only optimised for 30 meta-steps, and the compute-normalised (bottom row) plots in Figure A.1 show clearly that $H = 100$ shapers have not yet converged, we hypothesise that $H = 100$ shapers would ultimately yield the best performance if given more optimisation time.

Interestingly, the shaping effect is especially strong on the Clostridium difficile bacterium, where $H = 100$ shapers significantly outperform all other antibody types across all reported metrics; see right-most row in Figure A.1. These results suggests that different antigens can exhibit very different behaviour in a shaping setting. However, ADIOS is able to produce antibodies with effective shaping behaviour across them.

6.4. Attack is the Best Defence

Our previous results demonstrate that antibody shapers, particularly those optimised with longer horizons, manage to effectively minimise viral escape. However, we hypothesise they can achieve this through two distinct strategies: *robustness* or *shaping*. A robustness strategy involves developing antibodies that are inherently resistant to a wide range of potential viral variants — a "good defence" approach. In contrast, a shaping strategy aims to actively influence the evolutionary trajectory of the virus itself, creating evolutionary pressures that guide viral mutations in a direction more favourable to antibody binding — an "attack" approach.

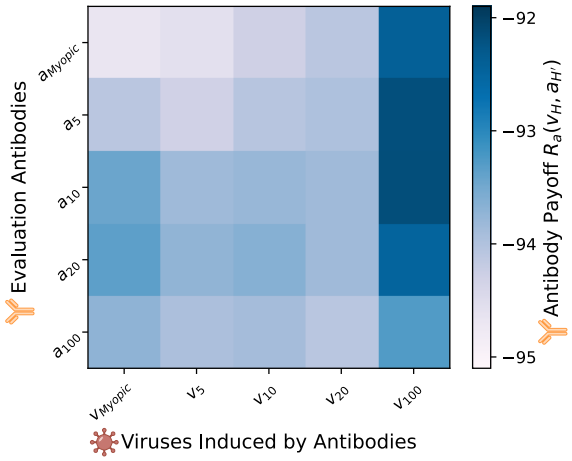


Figure 5. **Robustness vs. Shaping.** We optimise 80 different antibodies a_H across multiple horizons (*Myopic*, $H = 5, H = 10, H = 20, H = 100$), these are represented by the y-axis. We simulate the viral escape to each of these antibodies for 100 steps, and we group the escape viruses v_H by the horizon H of the antibody that induced them; these escape viruses are represented by the x-axis. In colour, we show the mean antibody payoff $R_a(v_H, a_{H'})$ for each group of optimised antibodies $a_{H'}$ against the final escape viral variant v_H induced by other antibodies optimised with horizon H. Darker colours correspond to better antibody payoff.

To disentangle these strategies, we *separately* evaluate the antibodies and the viruses that evolve in response to them. To do this, we compare the viruses against *other* antibodies which *did not* influence the viral evolution. The intuition is that an antibody which is good at shaping (good attack), but less robust (poor defence), will induce viruses which *other* antibodies will perform well against. Specifically, we generate antibodies a_H for each horizon H and simulate viral escape against these antibodies for 100 steps, resulting in viruses v_H. For all pairs of horizons (H, H'), we then cross-evaluate the antibody payoff $R_a(v_H, a_{H'})$. Figure 5 presents the result of this analysis.

Interestingly, viruses v_{100} induced by $H = 100$ shapers are consistently more exploitable by antibodies across all optimisation horizons. This suggests that $H = 100$ shapers actively shape the escape trajectories of the virus in a way that makes the resulting variants more susceptible to antibody binding in general. However, this shaping effect comes at a cost. The $H = 100$ shapers (a_{100}) show slightly lower payoffs compared to the peak performance of shorter-horizon antibodies (a_5 and a_{10}) against the viruses v_{100} induced by the $H = 100$ shapers (see rightmost column of Figure 5). This trade-off indicates that to exert a stronger shaping influence on viral evolution, $H = 100$ shapers sac-

rifice some degree of immediate performance or robustness, that is, their ability to perform well against a wide range of viruses. Therefore, a potential strategy could involve using a mixture of antibodies as therapy, where some are optimised for shaping the virus's evolutionary trajectory, and others are designed for strong immediate binding.

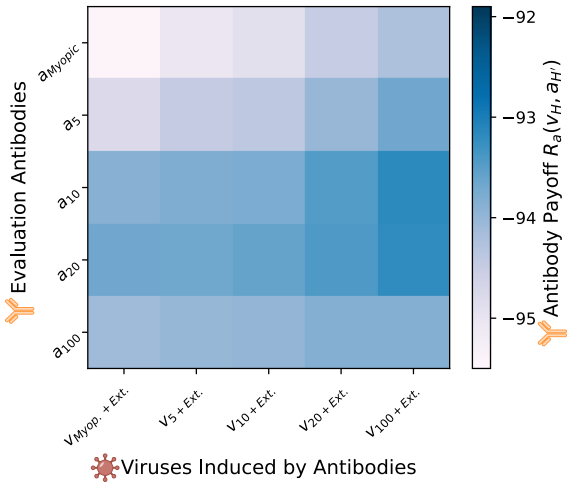


Figure 6. **Shaping with External Pressure.** A similar experiment to Figure 5, but with the additional external pressure of a separate myopic antibody (see Equation 3).

To investigate whether shaping persists in more realistic scenarios with multiple therapeutic pressures, we conduct an additional experiment. Using the same groups of antibodies a_H, we simulate viral escape with external pressure from a myopic antibody a_{Ext} (not included in the original a_{myopic} set). The viruses v_{H+Ext} now evolve according to a modified payoff:

$$R_v^{Ext}(v, a, a_{Ext}) = \frac{1}{2}R_v(v, a) + \frac{1}{2}R_v(v, a_{Ext}) \qquad (3)$$

which represents the scenario where multiple therapies are present in the environment (for example, during COVID-19 when multiple vaccines were available). Figure 6 shows that while the shaping effect is somewhat reduced compared to our original results, it remains clearly visible. This demonstrates that our shaping approach transfers to test regimes where external pressures from other therapies are present.

6.5. Explainability Analysis

To understand what distinguishes shapers from myopic antibodies, we conduct two complementary analyses of amino acid distributions and binding poses (Appendices B and C). Examining amino acid distributions, we find that long-horizon shapers exhibit more uniform distributions, while myopic antibodies tend to cluster around amino acids with extreme binding energies. We hypothesise that by maintaining diversity in their amino acid composition, shapers can preserve robustness against viral mutation since the virus cannot easily escape by avoiding specific, high-binding, parts of the antibody.

Through analysis of binding using pose matrices, representing which parts of the antibody and virus bind with each other, we observe these interaction patterns significantly change as the virus adapts. However, the type of antibody influences the nature of these changes: $H = 100$ shapers actively constrain viral evolution by both preventing unfavourable binding configurations and preserving favourable ones. While these findings are specific to our Absolut! binding simulator, they hint at explainable strategies that shapers use to influence viral evolution, which could inform future antibody design approaches.

7. Conclusion & Future Work

In this work, we introduce ADIOS, a meta-learning framework for designing therapeutic antibodies that not only defend against current viral strains but instead *actively shape viral evolution*. We provide a GPU-accelerated JAX implementation of Absolut!, enabling rapid simulation of viral escape trajectories and outer-loop optimisation. Our results demonstrate that shapers are not only more robust against viral escape, but they also shape viral evolution toward more targetable variants. Lastly, we provide an explainability analysis of how shapers achieve this level of robustness and influence, which we hope will inspire practitioners.

Although dengue virus served as our primary benchmark, we evaluated ADIOS on three other viruses (West Nile, Influenza and MERS-CoV) and on the bacterium Clostridium Difficile. In all four cases ADIOS consistently achieves shaping. These results confirm that ADIOS can generalise across a diverse set of pathogens. More broadly, the same opponent-shaping principle can be transferred to monoclonal-antibody (mAb) therapy for cancer (Zahavi & Weiner, 2020). In that setting, the outer loop would optimise therapeutic mAbs, while the inner loop would simulate the evolution of cancer-cell growth-factor receptors; the goal would be to shape cancer cells into cells that do not proliferate well. Exploring this direction, together with other bacterial or antimicrobial-resistance scenarios, remains an exciting avenue for future work.

While our current implementation uses simplified binding and evolutionary escape models that prevent direct therapeutic application, ADIOS could be integrated with more sophisticated models, like AlphaFold3 (Abramson et al., 2024), to better capture evolving viral and antibody structures. As computational models of protein interactions and evolutionary processes continue to improve, ADIOS has the potential to transform how we develop therapies against viruses, cancers, and other evolving adversaries.

Acknowledgements

We would like to thank Marius Urbonas, Duygu Açıkalın, Michael Matthews, Benjamin Ellis, Benedetta L. Mussati, Mattie Fellows, Lisa Zillig, Ting Lee, Matthew Raybould and Charlotte Deane for their invaluable contributions to this work. Their thoughtful insights helped guide the project direction, their detailed feedback on earlier drafts significantly enhanced the manuscript's clarity and accessibility, and their support during the review process was instrumental in addressing reviewers' concerns. We also thank the anonymous reviewers for their constructive feedback, which led to improvements in the paper. This work was supported by UK Research and Innovation and the European Research Council - Jakob Foerster is partially funded by the UKRI grant EP/Y028481/1, originally selected for funding by the ERC. This work was also supported by Exscientia, which provided Aleksandra Kalisz with a PhD studentship.

Impact Statement

This work introduces a machine learning framework for antibody design that could contribute to developing more effective therapies against evolving pathogens like viruses. While this has a potential positive societal impact through improved disease treatment and pandemic preparedness, we acknowledge that therapeutic development tools carry risks if misused. Our current implementation uses simplified models, and significant additional research and safety validation would be required before any real-world therapeutic applications.

References

Abramson, J., Adler, J., Dunger, J., Evans, R., Green, T., Pritzel, A., Ronneberger, O., Willmore, L., Ballard, A. J., Bambrick, J., Bodenstein, S. W., Evans, D. A., Hung, C.-C., O'Neill, M., Reiman, D., Tunyasuvunakool, K., Wu, Z., Žemgulytė, A., Arvaniti, E., Beattie, C., Bertolli, O., Bridgland, A., Cherepanov, A., Congreve, M., Cowen-Rivers, A. I., Cowie, A., Figurnov, M., Fuchs, F. B., Gladman, H., Jain, R., Khan, Y. A., Low, C. M. R., Perlin, K., Potapenko, A., Savy, P., Singh, S., Stecula, A., Thillaisundaram, A., Tong, C., Yakneen, S., Zhong, E. D., Zielinski, M., Žídek, A., Bapst, V., Kohli, P., Jaderberg, M., Hassabis, D., and Jumper, J. M. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, 630(8016): 493–500, June 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07487-w. URL https://www.nature.com/articles/s41586-024-07487-w.

Adolf-Bryfogle, J., Kalyuzhniy, O., Kubitz, M., Weitzner, B. D., Hu, X., Adachi, Y., Schief, W. R., and Jr, R. L. D. RosettaAntibodyDesign (RAbD): A general framework for computational antibody design. *PLOS Computational Biology*, 14(4):e1006112, April 2018. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1006112. URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1006112. Publisher: Public Library of Science.

Bennett, N. R., Watson, J. L., Ragotte, R. J., Borst, A. J., See, D. L., Weidle, C., Biswas, R., Shrock, E. L., Leung, P. J. Y., Huang, B., Goreshnik, I., Ault, R., Carr, K. D., Singer, B., Criswell, C., Vafeados, D., Sanchez, M. G., Kim, H. M., Torres, S. V., Chan, S., and Baker, D. Atomically accurate de novo design of single-domain antibodies, March 2024. URL https://www.biorxiv.org/content/10.1101/2024.03.14.585103v1. Pages: 2024.03.14.585103 Section: New Results.

Berman, H. M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T. N., Weissig, H., Shindyalov, I. N., and Bourne, P. E. The Protein Data Bank. *Nucleic Acids Research*, 28(1): 235–242, January 2000. ISSN 0305-1048. doi: 10.1093/nar/28.1.235. URL https://doi.org/10.1093/nar/28.1.235.

Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., and Wanderman-Milne, S. JAX: composable transformations of Python+ NumPy programs. 2018. URL https://scholar.google.com/scholar?cluster=8800325452556162710&hl=en&oi=scholarr.

Carabelli, A. M., Peacock, T. P., Thorne, L. G., Harvey, W. T., Hughes, J., de Silva, T. I., Peacock, S. J., Barclay, W. S., de Silva, T. I., Towers, G. J., and Robertson, D. L. SARS-CoV-2 variant biology: immune escape, transmission and fitness. *Nature Reviews Microbiology*, 21(3): 162–177, March 2023. ISSN 1740-1534. doi: 10.1038/s41579-022-00841-7. URL https://www.nature.com/articles/s41579-022-00841-7. Publisher: Nature Publishing Group.

Chéron, N., Serohijos, A. W., Choi, J., and Shakhnovich, E. I. Evolutionary dynamics of viral escape under antibodies stress: A biophysical model. *Protein Science : A Publication of the Protein Society*, 25(7):1332–1340, July 2016. ISSN 0961-8368. doi: 10.1002/pro.2915. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4918420/.

consortium, w. Protein Data Bank: the single global archive for 3D macromolecular structure data. *Nucleic Acids Research*, 47(D1):D520–D528, October 2018. ISSN 0305-1048. doi: 10.1093/nar/gky949. URL https://doi.org/10.1093/nar/gky949.

Cutting, D., Dreyer, F. A., Errington, D., Schneider, C., and Deane, C. M. De novo antibody design with SE(3) diffusion, May 2024. URL http://arxiv.org/abs/2405.07622. arXiv:2405.07622 [q-bio].

Dingens, A. S., Arenz, D., Weight, H., Overbaugh, J., and Bloom, J. D. An Antigenic Atlas of HIV-1 Escape from Broadly Neutralizing Antibodies Distinguishes Functional and Structural Epitopes. *Immunity*, 50(2):520–532.e3, February 2019. ISSN 1074-7613. doi: 10.1016/j.immuni.2018.12.017. URL https://www.cell.com/immunity/abstract/S1074-7613(18)30565-X. Publisher: Elsevier.

Doud, M. B., Lee, J. M., and Bloom, J. D. How single mutations affect viral escape from broad and narrow antibodies to H1 influenza hemagglutinin. *Nature Communications*, 9(1):1386, April 2018. ISSN 2041-1723. doi: 10.1038/s41467-018-03665-3. URL https://www.nature.com/articles/s41467-018-03665-3. Publisher: Nature Publishing Group.

Faramarzi, A., Norouzi, S., Dehdarirad, H., Aghlmand, S., Yusefzadeh, H., and Javan-Noughabi, J. The global economic burden of COVID-19 disease: a comprehensive systematic review and meta-analysis. *Systematic Reviews*, 13(1):68, February 2024. ISSN 2046-4053. doi: 10.1186/s13643-024-02476-6. URL https://doi.org/10.1186/s13643-024-02476-6.

Foerster, J., Chen, R. Y., Al-Shedivat, M., Whiteson, S., Abbeel, P., and Mordatch, I. Learning with Opponent-Learning Awareness. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '18, pp. 122–130, Richland, SC, July 2018. International Foundation for Autonomous Agents and Multiagent Systems.

Greaney, A. J., Starr, T. N., Gilchuk, P., Zost, S. J., Binshtein, E., Loes, A. N., Hilton, S. K., Huddleston, J., Eguia, R., Crawford, K. H. D., Dingens, A. S., Nargi, R. S., Sutton, R. E., Suryadevara, N., Rothlauf, P. W., Liu, Z., Whelan, S. P. J., Carnahan, R. H., Crowe, J. E., and Bloom, J. D. Complete Mapping of Mutations to the SARS-CoV-2 Spike Receptor-Binding Domain that Escape Antibody Recognition. *Cell Host & Microbe*, 29(1):44–57.e9, January 2021. ISSN 1931-3128. doi: 10.1016/j.chom.2020.11.007. URL https://www.cell.com/cell-host-microbe/abstract/S1931-3128(20)30624-7. Publisher: Elsevier.

Han, W., Chen, N., Xu, X., Sahil, A., Zhou, J., Li, Z., Zhong, H., Gao, E., Zhang, R., Wang, Y., Sun, S., Cheung, P. P.-H., and Gao, X. Predicting the antigenic evolution of SARS-COV-2 with deep learning. *Nature Communications*, 14(1):3478, June 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-39199-6. URL https://www.nature.com/articles/s41467-023-39199-6.

Hollingsworth, S. A. and Dror, R. O. Molecular dynamics simulation for all. *Neuron*, 99(6):1129–1143, September 2018. ISSN 0896-6273. doi: 10.1016/j.neuron.2018.08.011. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6209097/.

Hu, J., Peng, P., Wang, K., Fang, L., Luo, F.-y., Jin, A.-s., Liu, B.-z., Tang, N., and Huang, A.-l. Emerging SARS-CoV-2 variants reduce neutralization sensitivity to convalescent sera and monoclonal antibodies. *Cellular and Molecular Immunology*, 18(4):1061–1063, April 2021. ISSN 1672-7681. doi: 10.1038/s41423-021-00648-1. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7905196/.

Jin, W., Barzilay, R., and Jaakkola, T. Antibody-Antigen Docking and Design via Hierarchical Equivariant Refinement, July 2022. URL http://arxiv.org/abs/2207.06616. arXiv:2207.06616 [q-bio].

Lee, J. M., Eguia, R., Zost, S. J., Choudhary, S., Wilson, P. C., Bedford, T., Stevens-Ayers, T., Boeckh, M., Hurt, A. C., Lakdawala, S. S., Hensley, S. E., and Bloom, J. D. Mapping person-to-person variation in viral mutations that escape polyclonal serum targeting influenza hemagglutinin. *eLife*, 8:e49324, August 2019. ISSN 2050-084X. doi: 10.7554/eLife.49324. URL https://doi.org/10.7554/eLife.49324. Publisher: eLife Sciences Publications, Ltd.

Li, T., Pantazes, R. J., and Maranas, C. D. OptMAVEn – A New Framework for the de novo Design of Antibody Variable Region Models Targeting Specific Antigen Epitopes. *PLOS ONE*, 9(8):e105954, August 2014. ISSN 1932-6203. doi: 10.1371/journal.pone.0105954. URL https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0105954. Publisher: Public Library of Science.

Li, Y., Wan, Y., Liu, P., Zhao, J., Lu, G., Qi, J., Wang, Q., Lu, X., Wu, Y., Liu, W., Zhang, B., Yuen, K.-Y., Perlman, S., Gao, G. F., and Yan, J. A humanized neutralizing antibody against MERS-CoV targeting the receptor-binding domain of the spike protein. *Cell Research*, 25(11):1237–1249, November 2015. ISSN 1748-7838. doi: 10.1038/cr.2015.113. URL https://www.nature.com/articles/cr2015113. Publisher: Nature Publishing Group.

Lim, Y. W., Adler, A. S., and Johnson, D. S. Predicting antibody binders and generating synthetic antibodies us-

ing deep learning. *mAbs*, 14(1):2069075, 2022. ISSN 1942-0870. doi: 10.1080/19420862.2022.2069075.

Liu, G., Zeng, H., Mueller, J., Carter, B., Wang, Z., Schilz, J., Horny, G., Birnbaum, M. E., Ewert, S., and Gifford, D. K. Antibody complementarity determining region design using high-capacity machine learning. *Bioinformatics*, 36(7):2126–2133, April 2020. ISSN 1367-4803. doi: 10.1093/bioinformatics/btz895. URL https://doi.org/10.1093/bioinformatics/btz895.

Lok, S.-M., Kostyuchenko, V., Nybakken, G. E., Holdaway, H. A., Battisti, A. J., Sukupolvi-Petty, S., Sedlak, D., Fremont, D. H., Chipman, P. R., Roehrig, J. T., Diamond, M. S., Kuhn, R. J., and Rossmann, M. G. Binding of a neutralizing antibody to dengue virus alters the arrangement of surface glycoproteins. *Nature Structural & Molecular Biology*, 15(3):312–317, March 2008. ISSN 1545-9985. doi: 10.1038/nsmb.1382.

Lu, C., Willi, T., Witt, C. A. S. D., and Foerster, J. Model-Free Opponent Shaping. In *Proceedings of the 39th International Conference on Machine Learning*, pp. 14398–14411. PMLR, June 2022. URL https://proceedings.mlr.press/v162/lu22d.html.

Lucas, M., Karrer, U., Lucas, A., and Klenerman, P. Viral escape mechanisms – escapology taught by viruses. *International Journal of Experimental Pathology*, 82(5):269–286, 2001. ISSN 1365-2613. doi: 10.1046/j.1365-2613.2001.00204.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1046/j.1365-2613.2001.00204.x.

Madhi, S. A., Baillie, V., Cutland, C. L., Voysey, M., Koen, A. L., Fairlie, L., Padayachee, S. D., Dheda, K., Barnabas, S. L., Bhorat, Q. E., Briner, C., Kwatra, G., Ahmed, K., Aley, P., Bhikha, S., Bhiman, J. N., Bhorat, A. E., du Plessis, J., Esmail, A., Groenewald, M., Horne, E., Hwa, S.-H., Jose, A., Lambe, T., Laubscher, M., Malahleha, M., Masenya, M., Masilela, M., McKenzie, S., Molapo, K., Moultrie, A., Oelofse, S., Patel, F., Pillay, S., Rhead, S., Rodel, H., Rossouw, L., Taoushanis, C., Tegally, H., Thombrayil, A., van Eck, S., Wibmer, C. K., Durham, N. M., Kelly, E. J., Villafana, T. L., Gilbert, S., Pollard, A. J., de Oliveira, T., Moore, P. L., Sigal, A., Izu, A., NGS-SA Group, and Wits-VIDA COVID Group. Efficacy of the ChAdOx1 nCoV-19 Covid-19 Vaccine against the B.1.351 Variant. *The New England Journal of Medicine*, 384(20):1885–1898, May 2021. ISSN 1533-4406. doi: 10.1056/NEJMoa2102214.

Martinkus, K., Ludwiczak, J., Cho, K., Liang, W.-C., Lafrance-Vanasse, J., Hotzel, I., Rajpal, A., Wu, Y., Bonneau, R., Gligorijevic, V., and Loukas, A. AbDiffuser: Full-Atom Generation of in vitro Functioning Antibodies, March 2024. URL http://arxiv.org/abs/2308.05027. arXiv:2308.05027 [q-bio].

Mason, D. M., Friedensohn, S., Weber, C. R., Jordi, C., Wagner, B., Meng, S. M., Ehling, R. A., Bonati, L., Dahinden, J., Gainza, P., Correia, B. E., and Reddy, S. T. Optimization of therapeutic antibodies by predicting antigen specificity from antibody sequence via deep learning. *Nature Biomedical Engineering*, 5(6):600–612, June 2021. ISSN 2157-846X. doi: 10.1038/s41551-021-00699-9. URL https://www.nature.com/articles/s41551-021-00699-9.

Meijers, M., Ruchnewitz, D., Łuksza, M., and Lässig, M. Vaccination shapes evolutionary trajectories of SARS-CoV-2, July 2022. URL http://arxiv.org/abs/2207.09329. arXiv:2207.09329 [q-bio].

Miyazawa, S. and Jernigan, R. L. An empirical energy potential with a reference state for protein fold and sequence recognition. *Proteins: Structure, Function, and Bioinformatics*, 36(3):357–369, 1999. ISSN 1097-0134. doi: 10.1002/(SICI)1097-0134(19990815)36:3⟨357::AID-PROT10⟩3.0.CO;2-U. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/%28SICI%291097-0134%2819990815%2936%3A3%3C357%3A%3AAID-PROT10%3E3.0.CO%3B2-U.

Nandi, A. and Shet, A. Why vaccines matter: understanding the broader health, economic, and child development benefits of routine vaccination. *Human Vaccines & Immunotherapeutics*, 16(8):1900–1904. ISSN 2164-5515. doi: 10.1080/21645515.2019.1708669. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7482790/.

Nie, Z., Liu, X., Chen, J., Wang, Z., Liu, Y., Si, H., Dong, T., Xu, F., Song, G., Wang, Y., Zhou, P., Gao, W., and Tian, Y. A unified evolution-driven deep learning framework for virus variation driver prediction. *Nature Machine Intelligence*, 7(1):131–144, January 2025. ISSN 2522-5839. doi: 10.1038/s42256-024-00966-9. URL https://www.nature.com/articles/s42256-024-00966-9. Publisher: Nature Publishing Group.

Nybakken, G. E., Oliphant, T., Johnson, S., Burke, S., Diamond, M. S., and Fremont, D. H. Structural basis of West Nile virus neutralization by a therapeutic antibody. *Nature*, 437(7059):764–769, September 2005. ISSN 1476-4687. doi: 10.1038/nature03956. URL https://www.nature.com/articles/nature03956. Publisher: Nature Publishing Group.

Orenstein, W. A. and Ahmed, R. Simply put: Vaccination saves lives. *Proceedings of the National Academy of Sciences*, 114(16):4031–4033, April 2017. doi: 10.1073/pnas.1704507114. URL https://www.pnas.org/doi/10.1073/pnas.1704507114. Publisher: Proceedings of the National Academy of Sciences.

Orth, P., Xiao, L., Hernandez, L. D., Reichert, P., Sheth, P. R., Beaumont, M., Yang, X., Murgolo, N., Ermakov, G., DiNunzio, E., Racine, F., Karczewski, J., Secore, S., Ingram, R. N., Mayhood, T., Strickland, C., and Therien, A. G. Mechanism of Action and Epitopes of *Clostridium difficile* Toxin B-neutralizing Antibody Bezlotoxumab Revealed by X-ray Crystallography. *Journal of Biological Chemistry*, 289(26):18008–18021, June 2014. ISSN 0021-9258. doi: 10.1074/jbc.M114.560748. URL https://www.sciencedirect.com/science/article/pii/S0021925820405149.

Pereira, P., Minoux, H., Walczak, A. M., and Mora, T. Energy-based generative models for monoclonal antibodies, November 2024. URL http://arxiv.org/abs/2411.13390. arXiv:2411.13390 [q-bio].

Robert, P. A., Akbar, R., Frank, R., Pavlović, M., Widrich, M., Snapkov, I., Slabodkin, A., Chernigovskaya, M., Scheffer, L., Smorodina, E., Rawat, P., Mehta, B. B., Vu, M. H., Mathisen, I. F., Prósz, A., Abram, K., Olar, A., Miho, E., Haug, D. T. T., Lund-Johansen, F., Hochreiter, S., Haff, I. H., Klambauer, G., Sandve, G. K., and Greiff, V. Unconstrained generation of synthetic antibody–antigen structures to guide machine learning methodology for antibody specificity prediction. *Nature Computational Science*, 2(12):845–865, December 2022. ISSN 2662-8457. doi: 10.1038/s43588-022-00372-4. URL https://www.nature.com/articles/s43588-022-00372-4. Publisher: Nature Publishing Group.

Ruffolo, J. A., Chu, L.-S., Mahajan, S. P., and Gray, J. J. Fast, accurate antibody structure prediction from deep learning on massive set of natural antibodies. *Nature Communications*, 14(1):2389, April 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-38063-x.

Saka, K., Kakuzaki, T., Metsugi, S., Kashiwagi, D., Yoshida, K., Wada, M., Tsunoda, H., and Teramoto, R. Antibody design using LSTM based deep generative model from phage display library for affinity maturation. *Scientific Reports*, 11(1):5852, March 2021. ISSN 2045-2322. doi: 10.1038/s41598-021-85274-7. URL https://www.nature.com/articles/s41598-021-85274-7. Publisher: Nature Publishing Group.

Samsudin, E. Z., Yasin, S. M., Ruslan, N.-H., Abdullah, N. N., Noor, A. F. A., and Hair, A. F. A. Socioeconomic impacts of airborne and droplet-borne infectious diseases on industries: a systematic review. *BMC Infectious Diseases*, 24(1):93, January 2024. ISSN 1471-2334. doi: 10.1186/s12879-024-08993-y. URL https://doi.org/10.1186/s12879-024-08993-y.

Shanker, V. R., Bruun, T. U. J., Hie, B. L., and Kim, P. S. Unsupervised evolution of protein and antibody complexes with a structure-informed language model. *Science*, 385(6704):46–53, July 2024. doi: 10.1126/science.adk8946. URL https://www-science-org.ezproxy-prd.bodleian.ox.ac.uk/doi/10.1126/science.adk8946. Publisher: American Association for the Advancement of Science.

Thadani, N. N., Gurev, S., Notin, P., Youssef, N., Rollins, N. J., Ritter, D., Sander, C., Gal, Y., and Marks, D. S. Learning from prepandemic data to forecast viral escape. *Nature*, 622(7984):818–825, October 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06617-0. URL https://www.nature.com/articles/s41586-023-06617-0. Publisher: Nature Publishing Group.

VanDyk, L. and Meek, K. Assembly of IgH CDR3: mechanism, regulation, and influence on antibody diversity. *International Reviews of Immunology*, 8(2-3):123–133, 1992. ISSN 0883-0185. doi: 10.3109/08830189209055568.

Wan, H., Yang, H., Shore, D. A., Garten, R. J., Couzens, L., Gao, J., Jiang, L., Carney, P. J., Villanueva, J., Stevens, J., and Eichelberger, M. C. Structural characterization of a protective epitope spanning A(H1N1)pdm09 influenza virus neuraminidase monomers. *Nature Communications*, 6(1):6114, February 2015. ISSN 2041-1723. doi: 10.1038/ncomms7114. URL https://www.nature.com/articles/ncomms7114. Publisher: Nature Publishing Group.

Wang, G., Liu, X., Wang, K., Gao, Y., Li, G., Baptista-Hon, D. T., Yang, X. H., Xue, K., Tai, W. H., Jiang, Z., Cheng, L., Fok, M., Lau, J. Y.-N., Yang, S., Lu, L., Zhang, P., and Zhang, K. Deep-learning-enabled protein–protein interaction analysis for prediction of SARS-CoV-2 infectivity and variant evolution. *Nature Medicine*, 29(8):2007–2018, August 2023. ISSN 1546-170X. doi: 10.1038/s41591-023-02483-5. URL https://www.nature.com/articles/s41591-023-02483-5. Publisher: Nature Publishing Group.

Weisblum, Y., Schmidt, F., Zhang, F., DaSilva, J., Poston, D., Lorenzi, J. C., Muecksch, F., Rutkowska, M., Hoffmann, H.-H., Michailidis, E., Gaebler, C., Agudelo, M., Cho, A., Wang, Z., Gazumyan, A., Cipolla, M.,

Luchsinger, L., Hillyer, C. D., Caskey, M., Robbiani, D. F., Rice, C. M., Nussenzweig, M. C., Hatziioannou, T., and Bieniasz, P. D. Escape from neutralizing antibodies by SARS-CoV-2 spike protein variants. *eLife*, 9:e61312, October 2020. ISSN 2050-084X. doi: 10.7554/eLife.61312. URL https://doi.org/10.7554/eLife.61312. Publisher: eLife Sciences Publications, Ltd.

Yan Huang, Ziding Zhang, and Yuan Zhou. AbAgIntPre: A deep learning method for predicting antibody-antigen interactions based on sequence information. *Frontiers in Immunology*, 2022. doi: 10.3389/fimmu.2022.1053617.

Zahavi, D. and Weiner, L. Monoclonal Antibodies in Cancer Therapy. *Antibodies*, 9(3):34, July 2020. ISSN 2073-4468. doi: 10.3390/antib9030034. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7551545/.

Zambaldi, V., La, D., Chu, A. E., Patani, H., Danson, A. E., Kwan, T. O. C., Frerix, T., Schneider, R. G., Saxton, D., Thillaisundaram, A., Wu, Z., Moraes, I., Lange, O., Papa, E., Stanton, G., Martin, V., Singh, S., Wong, L. H., Bates, R., Kohl, S. A., Abramson, J., Senior, A. W., Alguel, Y., Wu, M. Y., Aspalter, I. M., Bentley, K., Bauer, D. L. V., Cherepanov, P., Hassabis, D., Kohli, P., Fergus, R., and Wang, J. De novo design of high-affinity protein binders with AlphaProteo, September 2024. URL http://arxiv.org/abs/2409.08022. arXiv:2409.08022 [q-bio].

A. Experimental Results on Other Viruses and Bacterium

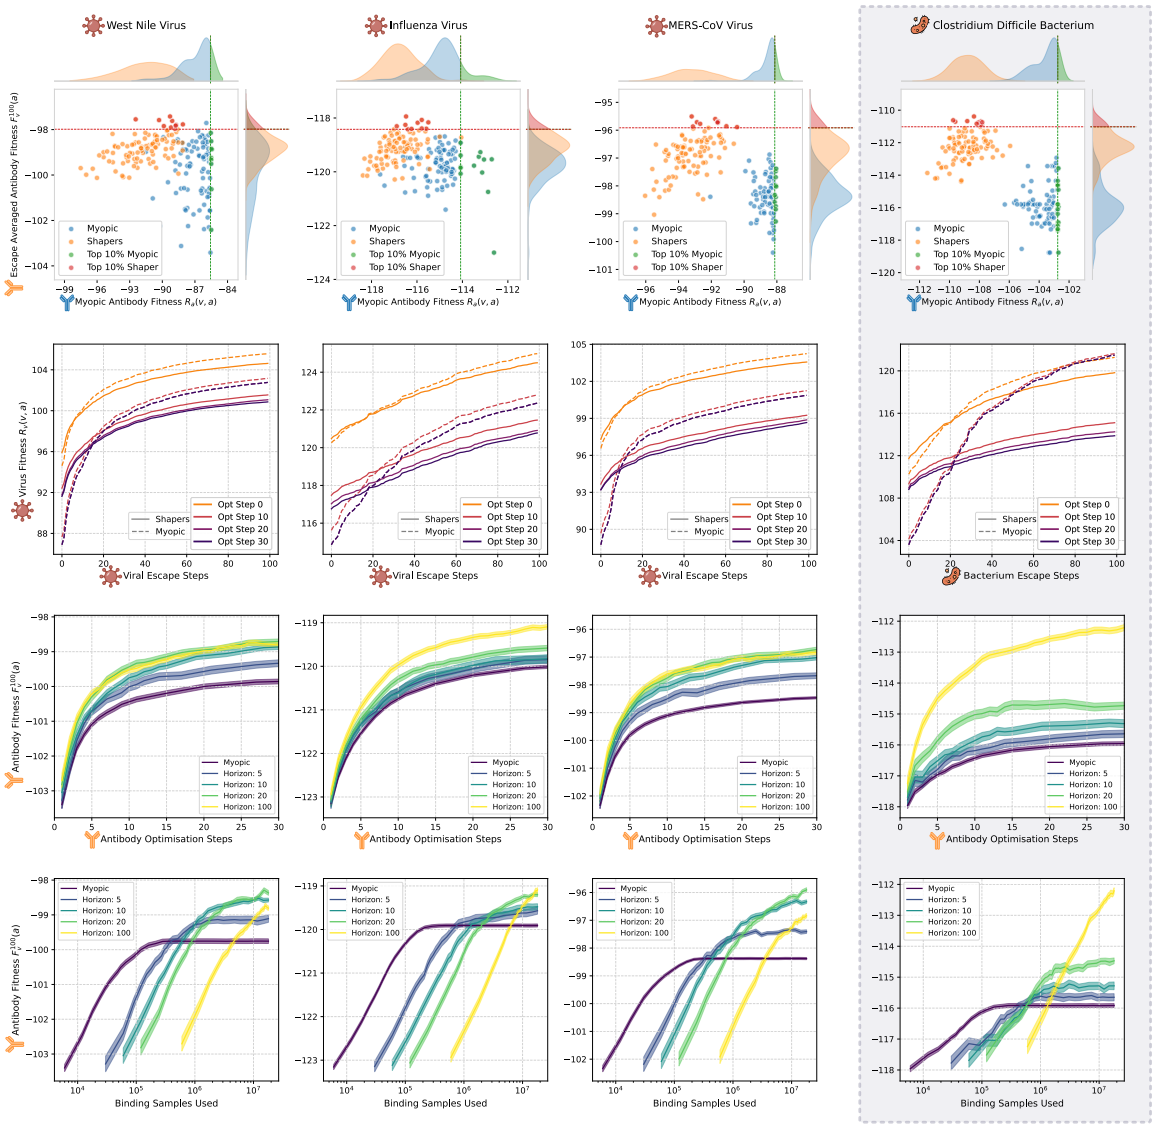


Figure A.1. **Shapers Outperform Myopic Antibodies on Other Viruses and a Bacterium.** The figure shows results across four different pathogens: West Nile Virus, Influenza Virus, MERS-CoV Virus, and Clostridium Difficile Bacterium (left to right columns). **First row** Distribution of antibody shapers optimised with horizon $H = 100$ (orange) vs. myopic antibodies distribution (blue). We highlight the top 10% shapers with respect to $F_v^{100}(a)$ in red, and the top 10% myopic antibodies with respect to $R_a(v, a)$ in green. The x-axis is the myopic antibody fitness $R_a(v, a)$ and the y-axis is the escape averaged antibody fitness for $H = 100$, i.e., $F_v^{100}(a)$. Higher values on both axes indicate better performance. **Second row** Viral escape curves (inner loop performance) for different steps of the antibody optimisation process (outer loop) for antibody shapers optimised with horizon $H = 100$ (solid lines) and myopic antibodies (dashed lines). The lighter lines indicate early antibody optimisation steps, and the darker lines show the later steps. The x-axis shows the evolutionary steps of viral escape. The y-axis represents the virus fitness/payoff $R_v(v, a)$, where higher values indicate better virus fitness (and lower values denote better antibody performance, so lower is better for us). **Third row** Antibody optimisation learning curves (outer loop performance) for varying horizon lengths. The x-axis shows the antibody optimisation steps, i.e., meta-steps, and the y-axis shows antibody fitness $F_v^{100}(a)$. **Forth row** Antibody optimisation learning curves accounting for computational cost. The x-axis shows the number of samples from the binding simulator, and the y-axis shows antibody fitness $F_v^{100}(a)$. Higher values indicate better performance. Error bars correspond to the standard error. In all these results, raw payoff/fitness values are dependent on Absolut! scores and are relative to the specific antigen, i.e., absolute values should not be compared between different viruses or bacterium but rather the overall trends.

B. Amino Acid Distribution in Shapers and Myopic Antibodies

To further understand the performance differences between myopic antibodies and shapers, we analyse how the amino acid distributions of the antibodies change with the optimisation horizon H. Within our model, each amino acid is solely characterised by its binding strength to other amino acids as defined by the Miyazawa-Jernigan energy potential matrix (Miyazawa & Jernigan, 1999). However, despite this simplicity, we still see interesting patterns in the amino acid distribution. Figure B.1 showcases the results of our experiment.

Antibodies optimised with longer horizons, especially the $H = 100$ shapers, exhibit a more uniform distribution of amino acids, while those with shorter horizons show a tendency to cluster around amino acids associated with either high or low binding energies. The flatter distribution of long-horizon shapers suggests a more diverse and balanced approach to viral antigen binding. We hypothesise that this strategy helps to preserve robustness against viral mutations. By maintaining a more even distribution across energy levels, these antibodies may be less susceptible to viral escape.

In contrast, the clustering behaviour we observe in shorter-horizon antibodies indicates a more specialised strategy. By concentrating on amino acids at the extremes of the binding energy spectrum, these antibodies may achieve strong immediate binding but potentially at the cost of long-term robustness. However, while this analysis hints at the robustness of long-term shapers, it does not fully explain the shaping behaviour we observed in our previous results. In the next section, we investigate the distribution of amino acids within specific binding poses.

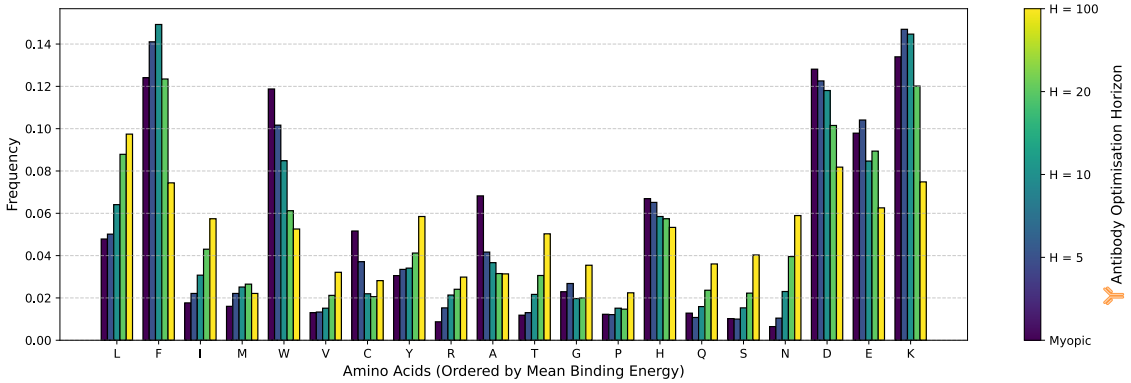


Figure B.1. **Distribution of amino acids in myopic antibodies and shapers.** The antibodies are optimised for $N = 30$ steps using different viral escape horizons H. Longer horizon shapers push the amino acid distribution closer to a uniform distribution.

C. Influence of Antibody Shapers on Binding Poses

In the Absolut! framework (Robert et al., 2022), binding poses are defined as sets of interacting residue pairs between the antibody and the antigen. The binding energy of a pose is calculated by populating these residue locations with the amino acid sequences of both the antibody and the virus and then summing the pairwise interaction energies defined by (Miyazawa & Jernigan, 1999). Absolut! considers a vast number of possible poses (on the order of 10^6) and determines the overall interaction energy as the energy of the minimum pose, refer to Appendix D for more details. Importantly, only a small part of the viral sequence contributes to this minimum energy pose.

As both the virus and the antibody mutate during our optimisation process, the lowest energy pose can change. To capture these dynamics, we introduced the concept of a pose matrix: a 20×20 matrix with one entry for each possible pair of amino acids. One dimension corresponds to the antibody amino acids, and the other to the viral amino acids. The entries in this matrix represent the number of interactions between the specific amino acid pairs in the lowest energy pose for the binding configuration between an antibody and an antigen.

Figure C.1a presents average pose matrices from multiple optimisation runs of both myopic antibodies and long-horizon shapers. We observe two key trends. First, as viral escape steps increase (top row vs bottom row), the pose matrices become more "diffused". This is expected, as the virus explores more "pose possibilities" through mutations during escape. Second, as the horizon of antibody optimisation increases, the poses also become more "diffused". This is particularly interesting, as all antibodies have the same number of mutations regardless of the horizon, suggesting that this diffusion might relate to the increased robustness of shapers.

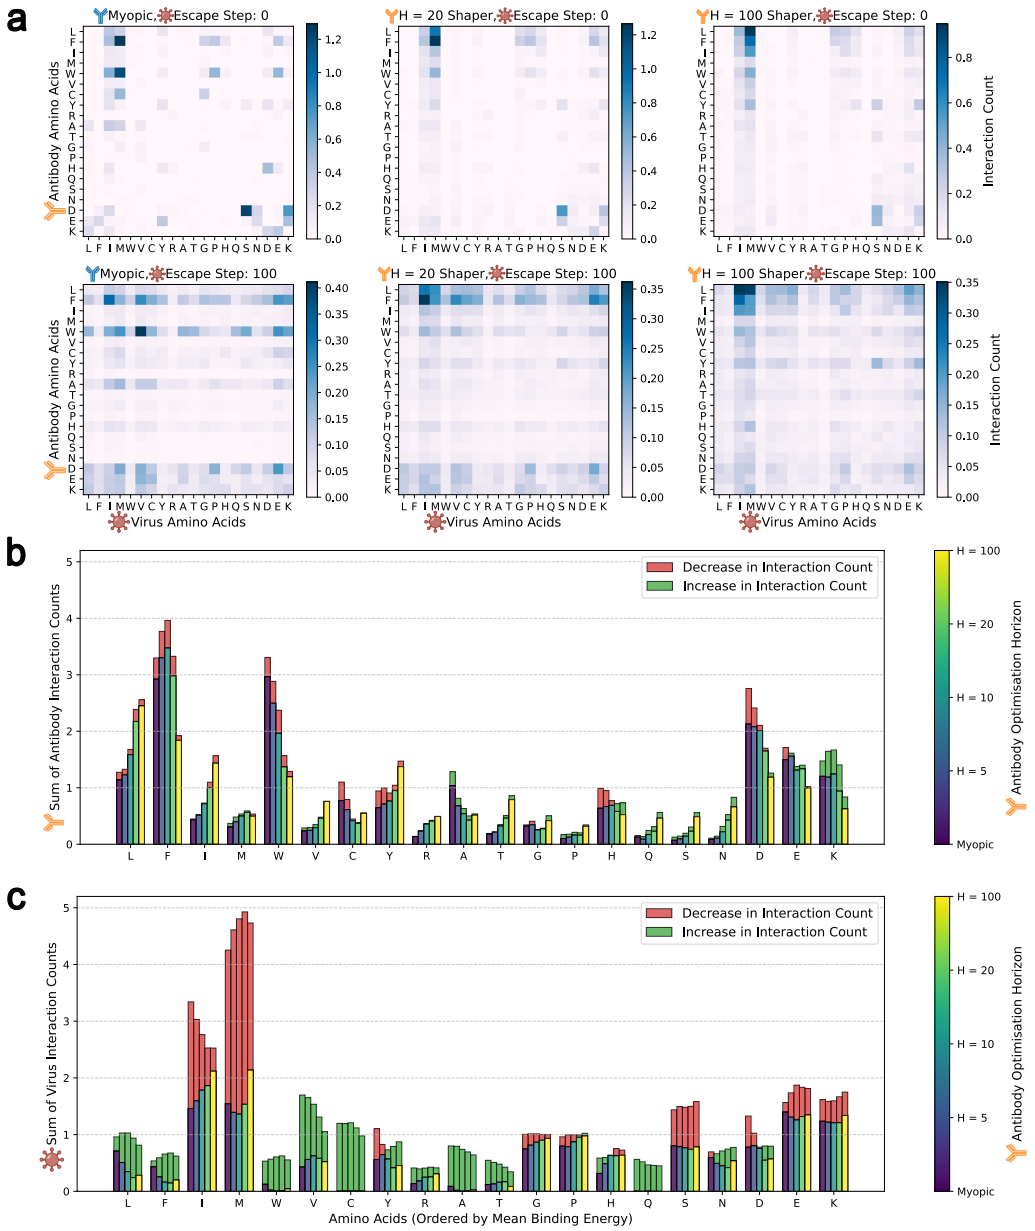


Figure C.1. **Influence of antibody shapers on binding poses. a** Average pose matrices between antibodies optimised using different horizons and the virus at various stages of its escape. The escape steps increase from left to right, and the horizon increases from top to bottom. The full grid of matrices with more antibody horizons and virus escape steps is available in the Supplementary Information, Figure C.2. **b, c** Aggregated sum of pose matrices w.r.t the antibody axis (**b**) and w.r.t the virus axis (**c**). The plots show a change in the interaction counts in the poses from the viral escape step 0 to 100. Red indicates a decrease in the interaction count and green an increase.

To further understand these pose dynamics, we aggregate the pose matrices along the antibody axis (Figure C.1b) and the virus axis (Figure C.1c). These figures show the change in interaction counts between viral escape steps 0 and 100. Figure C.1b shows that as the virus escapes it includes more of the antibody's lowest binding amino acids (particularly K, Lysine) in the pose. Notably, long-horizon shapers, especially $H = 100$ shapers, are most effective at preventing this increase in K (Lysine) interactions. Furthermore, Figure C.1c shows another viral escape strategy, where the virus removes its high-binding amino acids I (Isoleucine) and M (Methionine) from the pose. Again, $H = 100$ shapers are most successful in counteracting this trend, although they cannot completely prevent it.

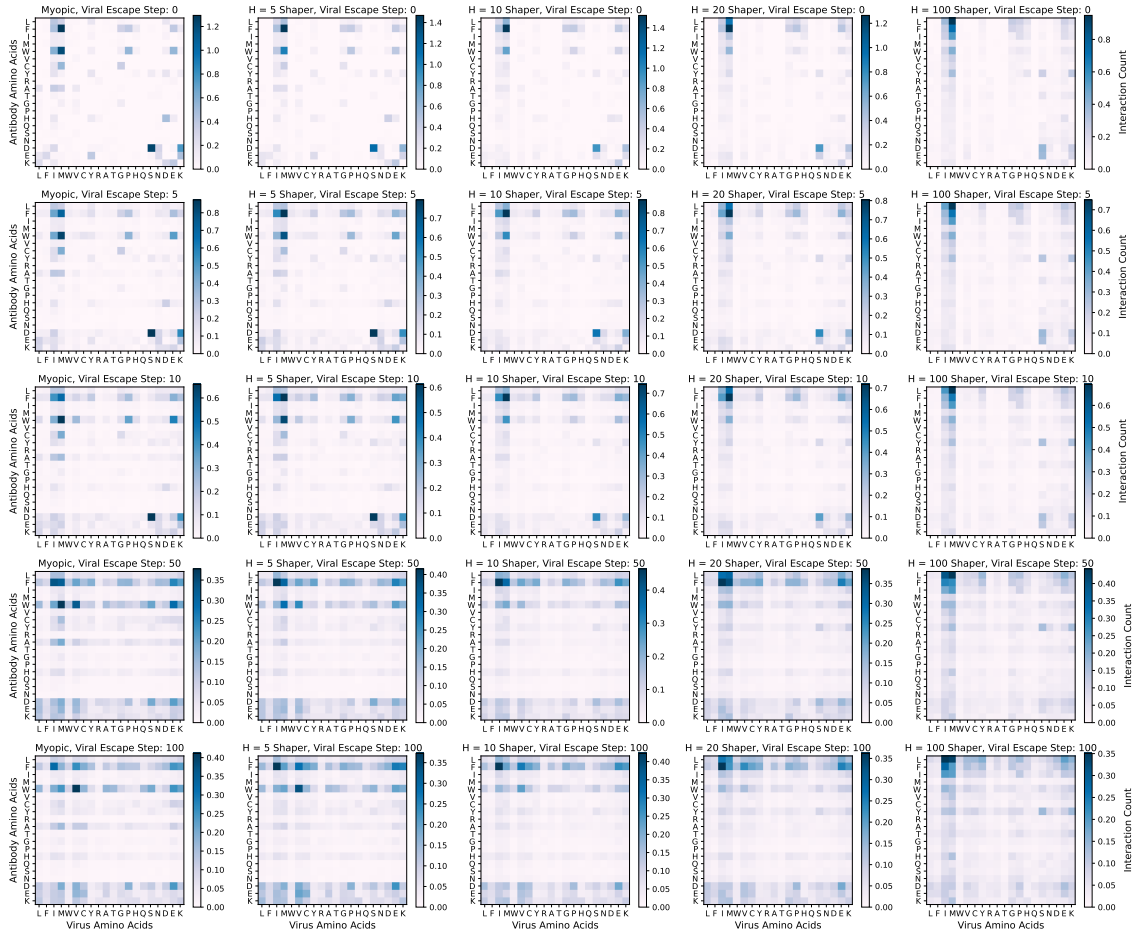


Figure C.2. Full grid of pose matrices. They represent the average pose interactions between myopic antibodies or antibody shapers optimised with different horizons and viruses at different stages of their escape.

Based on these observations, we hypothesise that the shaping ability of $H = 100$ shapers relies on two main mechanisms. Preventing the virus from including the antibody's lowest binding amino acids in the pose, and inhibiting the virus from removing its own high-binding amino acids from the pose. These strategies constrain the viral escape trajectories, making the resulting viral variants more susceptible to antibody binding in general. While these results are specific to our Absolut! binding simulator, they demonstrate that the behaviour of antibody shapers is both explainable and intuitive. This work serves as a proof of concept, showing that opponent shaping techniques can optimise antibodies to more effectively prevent viral escape.

D. Binding Function

In general, ADIOS is independent of the choice of the binding function $B : \mathbb{A}^{N_v} \times \mathbb{A}^{N_a} \to \mathbb{R}$. In our work, we rely on the Absolut! framework (Robert et al., 2022) to implement the binding function. In this section, we mathematically formalise the binding energy calculation that Absolut! uses. For further explanation, readers are recommended to refer to the original Absolut! paper (Robert et al., 2022).

For two given protein structures, there are many possible joint configurations. Each of these joint configurations yields an energy. The configurations which are associated with lower energy will require more external energy to cause the system to leave that state, meaning in turn that they are more stable. If the configuration is sufficiently stable, this may be referred to as a binding pose.

In Absolut, poses are represented as pairs of residues[2] which are adjacent to each other in that pose. In particular, the pairs may be from the antigen to the antibody, or from the antibody to itself. We define the space of possible poses Φ:

$$\Phi = 2^{N_v \times N_a} \times 2^{N_a \times N_a}$$

Where N_v and N_a are taken to be the set of integers up to N_v and N_a, respectively.

The energy of a complex of a virus $v \in \mathbb{A}^{N_v}$ and an antibody $a \in \mathbb{A}^{N_a}$, in a given pose $(\phi^{v\times a}, \phi^{a\times a}) \in \Phi$ is defined by sum of the energies of each adjacent residues. The energy between a residue pair is determined by which two amino acids it contains, given by a symmetric *interaction matrix* $M : \mathbb{A} \times \mathbb{A} \to \mathbb{R}$, which is determined experimentally (Miyazawa & Jernigan, 1999).

We then define the energy of a single pose to be:

$$\hat{E}(a, b; (\phi^{v\times a}, \phi^{a\times a}), M) = \sum_{(i,j)\in\phi^{v\times a}} M(v_i, a_j) + \sum_{(i,j)\in\phi^{a\times a}} M(a_i, a_j)$$

Finally, given a set of poses $S \subseteq \Phi$, the binding strength is:

$$B(v, a) = -E(v, a; S, M) = -\min_{\phi\in S} \hat{E}(v, a; \phi, M) \tag{4}$$

Absolut! generates S through a two-step process. First, Absolut! discretises a given structure of the virus v (or any antigen) which is taken from the PDB (consortium, 2018). Second, Absolut! does a brute-force search over possible (discretised) poses for an antibody a to join to the viral structure. The exact details are not necessary for this paper, we refer interested readers to the original paper.

However, we find that Absolut! generates more poses than we require. Since the energy function, E is a minimum over poses, certain poses contribute far more than others. In particular, if a pose ϕ tends to yield higher energies, so $\hat{E}(a, b; \phi, M)$ is relatively large, it will have little impact on the result of B.

To give a more concrete example, for this paper, we use the dengue virus antigen (Lok et al., 2008). Absolut! gives ≈ 1.5 million poses for this structure. Absolut! also comes with ≈ 20 million real-world antibody sequences. When using the base dengue sequence as the antigen, across the 20 million binding calculations only 1027 binding poses are ever the minimum. Furthermore, the relevance of each pose drops exponentially. The most relevant pose accounts for 20% of binding configurations, and by using the top 100 poses one would get the exact same result for binding in 95% of antibodies. This gives us a way to make the computation 1000 times faster[3] for a negligible accuracy drop for this particular antigen sequence.

However, this leads to more errors as soon as we change the viral antigen sequence. Looking at the particular poses which lead to binding reveals another way to cut down on the total number of poses: all of the poses contain at least 18 pairs of residues. As the interaction matrix M is strictly negative, having more pairs of residues always makes the binding energy of a pose lower, meaning it is more likely to be where binding occurs. Out of the original 1.5 million poses, only approximately 37 thousand (1 in 40) contain 18 or more pairs of residues. When using only these residues, we see no differences across any of the evolutionary simulations. It is possible that a pose with 17 or less pairs is the dominant one for some antigen v with antibody a, but if so, then such pairs appear to be extremely rare.

Using these methods of pruning poses gives us two subsets of the original set of poses, a larger one which almost exactly matches performance, and another which sometimes differs, but is much faster to compute. We refer to these as the *high-resolution* and *low-resolution* binding simulators, respectively. Note that for the low-resolution binding simulator, the more mutations the virus undergoes, the less accurate it becomes. Furthermore, we also do binding to the antibody anti-target, t_a^-. To account for this, we compute the relevant poses for this anti-target too.

When running dengue virus experiments, we always train with the low-resolution binding simulator, then perform "verification" with the high-resolution one, and these are the results we report throughout the paper. For the other viruses and the bacterium we run "verification" in a differently initialised instantiation of the low-resolution simulator. The reason is

[2] A single amino acid position on a protein

[3] In practice, the difference is closer to $10,000$, likely due to the GPU having to move less data.

twofold. Firstly this enables us to run many more evolution experiments. Secondly, this mimics the real-life process of transferring out of simulation to the real world. By showing we transfer from the low-resolution binding simulation to the slower, high-resolution binding simulation, we demonstrate that our results are not extremely specific to the exact simulation we use and that any result will not disappear as soon as a more accurate simulation is used. We emphasise that Absolut! does not represent an accurate model of antibody binding. It is instead a toy simulation to demonstrate our methodology. For example, we do not expect our framework *when used with this simulation model* to yield highly effective, superior antibodies for real-world applications.

E. MDP for the Virus–Antibody Game

We formalise a *single interaction round* between an antibody and a virus as the finite–horizon Markov Decision Process $\mathcal{M}$:

$$\mathcal{M} = \langle \mathcal{S},\, \mathcal{A}^{\mathrm{v}},\, \mathcal{A}^{\mathrm{a}},\, P,\, R,\, \mu \rangle.$$

- **State space**: $\mathcal{S} = \mathbb{A}^{N_v} \times \mathbb{A}^{N_a}$, where a state $s = (v, a)$ is the pair of viral (v) and antibody (a) amino-acid sequences.
- **Action spaces** (chosen simultaneously):

$$\mathcal{A}^{\mathrm{v}}(s) = \mathbb{A}^{N_v}, \qquad \mathcal{A}^{\mathrm{a}}(s) = \mathbb{A}^{N_a}.$$

- **Transition kernel** (episode terminates after one step):

$$P\big(s' \mid s, a^{\mathrm{v}}, a^{\mathrm{a}}\big) = \mathbf{1}_{s'=s_\star}, \quad s_\star \text{ is terminal.}$$

- **Reward vector** $R = (R_{\mathrm{v}}, R_{\mathrm{a}})$. Given joint action $\big(a^{\mathrm{v}}, a^{\mathrm{a}}\big)$ in state s,

$$R_{\mathrm{a}}\big(s, a^{\mathrm{v}}, a^{\mathrm{a}}\big) = B\big(a^{\mathrm{v}}, a^{\mathrm{a}}\big) - B\big(t_a^-, a^{\mathrm{a}}\big) - B\big(a^{\mathrm{v}}, t_v^+\big), \qquad R_{\mathrm{v}} = -R_{\mathrm{a}}.$$

- **Initial-state distribution** μ puts mass on the wild-type virus and the initial antibody candidate: $\mu\big(v_0, a_0\big) = 1$.

Because this is a single-step MDP, the return equals the immediate reward R and the discount factor is irrelevant.

5.5 Limitations and Future work

The Absolut! framework is sufficiently efficient to support the large-scale evolutionary trajectories required for this study; however, it is a simplified lattice-based model that does not fully capture the nuanced dynamics of protein-protein interactions. In contrast, full Molecular Dynamics (MD) simulations provide a much higher degree of realism but at a prohibitive cost. A single MD simulation for a specific binding pose can require several hours on a high-end GPU, representing a computational demand approximately 10^7 times greater per evaluation than Absolut!.

This disparity introduces uncertainty regarding the magnitude and robustness of the observed effects in *in vivo* biological systems. Consequently, validating the ADIOS framework using more sophisticated simulators remains a critical direction for future research. Recent advances in deep learning-based structural biology offer a potential middle ground. Protein structure prediction tools, such as AlphaFold3 or RoseTTAFold-All-Atom, could provide complex structures for subsequent binding-affinity estimation using an additional scoring model. Structure prediction alone can take minutes per complex, depending on system size and hardware, before subsequent affinity scoring [74]. While this remains orders of magnitude slower than lattice-based models, necessitating the use of Absolut! for this initial stage of work, it represents a feasible next step for validating these strategies on a more biologically representative scale.

A second limitation is the lack of real-world validation. Ideally, one would test ADIOS in a simple, experimentally tractable system, such as a bacteria–phage system, where it can be both computationally simulated with reasonable fidelity and easily manipulated in a wet lab. However, designing and executing such experiments is highly non-trivial. An alternative approach would be to validate the framework using existing empirical datasets. This avoids the need for lengthy experimental cycles and may allow analysis of major evolutionary events, such as the emergence of SARS-CoV-2 variants, through the lens of ADIOS. Careful experimental design would still be required, but even a preliminary form of real-world validation would significantly strengthen the framework.

Finally, ADIOS itself is a rich setting, and the analyses in this chapter only scratch the surface of the possible questions. In our simulations, a single antibody is selected and then held fixed throughout the evolutionary process. In practice, deploying multiple therapies or adapting therapies dynamically in response to viral mutations may yield substantial benefits. How to search efficiently over the combinatorial space of dual (or multi-agent) therapies, particularly when limited observations of the virus are available, remains an open question. Likewise, this chapter begins to characterise shaper antibodies, but the analysis remains preliminary. Identifying simple, predictive features that reliably distinguish shapers from non-shapers would be especially valuable, and could open the possibility that these insights transfer to more realistic biological settings.

Our desires cut across one another, and in this confused existence it is rare for happiness to coincide with the desire that clamoured for it.

Marcel Proust *In Search of Lost Time*

6 Conclusion

Contents

6.1 Summary 125
6.2 Limitations and Future work 127
6.3 Future Direction 129
6.4 Final Remarks 130

This thesis presents research on simulating and steering evolutionary systems. This concluding chapter summarises the contributions, their limitations, and potential future directions.

6.1 Summary

The goal of understanding how to best steer evolutionary systems is divided into two parts. The first part, directed evolution, can be seen as the basic co-operative problem, where we desire to accelerate an evolutionary system to reach the best optimum possible in the smallest number of iterations. The second part, antibody shaping, can be instead seen as the zero-sum alternative, where the primary goal is to limit the rate of an evolutionary system. The chapters are summarised via the initial questions they seek to answer.

How can we frame directed evolution as an optimisation problem? Chapter 3, **Optimisation strategies for directed evolution without sequencing**, frames selection rules and population allocation as choices in an optimisation problem that can be compared through simulation. It shows how a naive greedy strategy may be suboptimal depending on the properties of the underlying landscape (such as ruggedness), and the parameters of the evolutionary system (such as mutation rate). In particular, we introduce **base chance**, some fixed probability of cells being selected regardless of fitness, and **population splitting**, dividing your sample budget across several independent runs, as potential ways to enhance the results of a directed evolution experiment in a relatively simple manner.

How can we infer useful properties of a fitness landscape from population-level measurements? Chapter 4, **Sequence-free landscape inference for directed evolution**, shows how the dynamics of mean population fitness under random mutation can be used to estimate a scalar measure of landscape ruggedness without sequencing. This naturally builds on Chapter 3 by providing information that can guide the choice of selection strategy. The choice of ruggedness metric is lifted from **spectral landscape theory**, chosen primarily due to it being inferable from a landscape in a relatively simple, sequencing-free procedure, and it having a strong connection to the mathematical NK model. We compare this Fourier ruggedness metric to other ruggedness metrics from the literature, showing they all strongly correlate, and also providing statistical insights into how best to estimate this metric from noisy data.

Can we apply principles from opponent shaping to viral evolution? Chapter 5 demonstrates this approach in simulation through **ADIOS: Antibody Development via Opponent Shaping**, which optimises antibodies against the viral escape trajectories they induce. This chapter considers the setting of an evolving virus, in which the evolutionary objectives of the virus and those of

humanity are in opposition. We model the system as zero-sum, with our sole control being the choice of antibody deployed. Note that, technically, the interaction is not strictly zero-sum in the real world, as evolution optimises reproductive fitness, whereas our primary concern is limiting viral lethality. This model is enabled using Absolut!, a basic but extremely fast binding simulator that is necessary for the huge number of binding samples required to perform enough simulations to get any level of statistical significance. Even in this simple model of binding, the system is shown to exhibit rich and interesting properties, with different antibodies being best (from those tested) depending on the evolutionary time horizon that is being optimised for. The sequences, too, are analysed, showing how the particular amino acids that tend to accumulate in the sequences depend on if the antibody is short-horizon (myopic) or long-horizon (**shapers**). Further than just the amino acids, the antibodies and evolved viruses are tested against each other, showing that the viruses that result from shaper antibodies yield higher antibody payoffs across the tested groups, suggesting shapers are not only robust, that is individually able to deal with the results of adversarial evolution, but actually guide the virus towards variants that are easier for other antibodies to bind in the model.

6.2 Limitations and Future work

Overall, the two primary limitations of this body of work are the lack of empirical validation and the constraints inherent in the synthetic landscapes and models employed.

For the directed evolution component, empirical validation is currently bottlenecked by the small size of combinatorially complete landscapes. Because the curse of dimensionality suggests that such landscapes are unlikely to grow significantly larger in the near future, the most viable path forward is direct laboratory experimentation. This is the primary motivation for utilising Microscopic PhotoSelection (**MiPS**) in mother-machine microfluidic devices. As a platform capable of automatically observing and manipulating $\sim 10^6$ cell lineages, MiPS provides a robust benchmark for demonstrating the advantages of complex selection methods. This line of work

has been developed with the MiPS platform in mind, as its capacity for fine-grained control and high-throughput data acquisition demands more sophisticated selection and optimisation frameworks. However, since the MiPS platform is being built by the lab almost entirely from scratch, both hardware and software, there have been unexpected difficulties in development. This is only to be expected with such sophisticated bleeding-edge technologies; delays are likely to occur. It would have been ideal to be able to test these works on the MiPS platform for this thesis, but fate has decided otherwise.

For virus shaping, empirical validation is more nuanced, as the work centres on evolutionary counterfactuals and assumes prior knowledge of the underlying system. The fundamental question is whether shaping is possible, and what level of systemic understanding is required to achieve it. An ideal empirical validation would involve identifying a simplified real-world system of adversarial evolution and attempting to implement shaping within it. A more accessible, albeit weaker, form of validation would be to demonstrate that a complex binding simulator can leverage information regarding circulating antibodies to improve predictions of viral evolution. This would validate that antibody-driven selective pressure can be used to forecast evolutionary trajectories, even if it does not fully confirm the ultimate "shaping" hypothesis.

Regarding directed evolution, the use of NK landscapes represents a significant limitation. With only two primary parameters, the NK model, while elegant, lacks the complexity required for data-driven simulation applications. A richer space of synthetic landscapes would allow for better alignment between synthetic and real-world landscape properties, enabling more efficient computational calibration. However, a challenge remains: while the deficiencies of current synthetic landscapes are clear, it is difficult to determine which metrics would define a “superior” synthetic model. One could train methodologies on synthetic landscapes and evaluate their transferability to empirical ones; however, since currently available combinatorially complete landscapes are small and potentially biased (often $N \leq 4$), there is a significant risk of overfitting to these limited samples. Nevertheless, developing models that surpass the simplicity of the NK baseline is a necessary step forward.

Finally, for ADIOS, the limitations of the simple binding simulator (Absolut!) are more straightforward to address. Utilising more accurate, albeit computationally intensive, binding simulators is entirely feasible. While running ADIOS on such platforms would represent a major increase in computational cost, it remains within the realm of practical feasibility and would significantly enhance the fidelity of the results.

6.3 Future Direction

For this section, I speak more freely regarding what I believe to be the future of the directions I have investigated for my thesis.

Using simulation to understand evolutionary systems has a long history, and I believe this thesis represents a natural continuation of it to modern hardware and questions. Having methodologies which leverage compute is incredibly valuable, as computation tends to increase rapidly over time, especially when compared with the cost of gathering real-world data.

However, this only works as long as what is being simulated is relevant to reality. There is a trade-off here: that which is more computationally efficient to calculate lends itself towards a more extensive set of computational experiments, but each of those experiments will be less representative of reality. For example, Absolut! exists on one end of this spectrum and a full molecular dynamics simulation exists on the other end.

There is also a question of what "accuracy" means for these proxies. The NK model clearly doesn't represent any individual protein; instead, it is meant to help us understand the nature of natural selection in the abstract. Similarly, Absolut! is not designed to actually say how well a given real-world antibody will bind, but rather to help us understand the binding fitness landscape.

I believe filling out this space of fitness landscape proxies is of great interest. We may not understand exactly what the fitness landscapes are, and due to their high dimensionality we possibly never will, but we have some understanding of what is reasonable, and what is not, and this should be captured in our synthetic landscapes.

This is relatively easy to begin with, as the NK landscape is so extremely simple and basic that it seems trivial to make improvements to it, or at least to have a greater diversity of landscapes to prevent overfitting to any one. The challenge then becomes understanding how to test what is biologically relevant, which may change depending on the objective in question.

With this, it is easier to use computation to probe what our current understanding of fitness landscapes implies, rather than gather new data. After all, raw computation will likely always be vastly cheaper than gathering more data. Having a well-established prior is also useful in getting more leverage out of the data we do gather. The directed evolution work in this thesis demonstrates this nicely: there are general features of landscapes, so you can train in simulation, then gather just enough data to identify parameters.

So, overall, I would push for more computationally tractable models, more simulation to understand the properties of said models, and then more ML inference to connect real-world settings to said models.

6.4 Final Remarks

Biology is entering a transformative phase. While biological systems are inherently noisy and complex, modern machine learning methods have demonstrated a remarkable ability to extract underlying patterns from significant noise. This shift is fundamentally data-driven, facilitated by the emergence of massively parallel experimental platforms, such as the MiPS platform discussed previously, which provide the high-resolution data necessary for sophisticated modelling.

The work presented in this thesis is, by its nature, somewhat theoretical (or forward-looking, to be more generous); it explores the theoretical landscape of how we might leverage these advanced platforms and predictive frameworks to control and anticipate evolutionary trajectories. As the synergy between high-throughput experimentation and predictive modelling continues to mature, it will be of great interest to see how these frameworks transition from theoretical benchmarks to transformative tools in the laboratory and beyond.

The first kind of intellectual and artistic personality belongs to the hedgehogs, the second to the foxes ...

— Sir Isaiah Berlin [75]

References

[1] Michael L. Metzker. "Sequencing technologies — the next generation". In: *Nature Reviews Genetics* 11.1 (Jan. 2010). Number: 1 Publisher: Nature Publishing Group, pp. 31–46. URL: https://www.nature.com/articles/nrg2626 (visited on 10/27/2023).

[2] Anthony Rhoads and Kin Fai Au. "PacBio Sequencing and Its Applications". In: *Genomics, Proteomics & Bioinformatics* 13.5 (Oct. 2015), pp. 278–289.

[3] Ashraful Haque et al. "A practical guide to single-cell RNA-sequencing for biomedical research and clinical applications". In: *Genome Medicine* 9.1 (Aug. 18, 2017), p. 75.

[4] Jeremy J. Agresti et al. "Ultrahigh-throughput screening in drop-based microfluidics for directed evolution". In: *Proceedings of the National Academy of Sciences of the United States of America* 107.9 (Mar. 2, 2010), pp. 4004–4009.

[5] Gökcen Eraslan et al. "Deep learning: new computational modelling techniques for genomics". In: *Nature Reviews Genetics* 20.7 (July 2019). Number: 7 Publisher: Nature Publishing Group, pp. 389–403. URL: https://www.nature.com/articles/s41576-019-0122-6 (visited on 10/27/2023).

[6] Philipp Eulenberg et al. "Reconstructing cell cycle and disease progression using deep learning". In: *Nature Communications* 8.1 (Sept. 6, 2017). Number: 1 Publisher: Nature Publishing Group, p. 463. URL: https://www.nature.com/articles/s41467-017-00623-3 (visited on 10/27/2023).

[7] Alex Zhavoronkov et al. "Deep learning enables rapid identification of potent DDR1 kinase inhibitors". In: *Nature Biotechnology* 37.9 (Sept. 2019). Number: 9 Publisher: Nature Publishing Group, pp. 1038–1040. URL: https://www.nature.com/articles/s41587-019-0224-x (visited on 10/27/2023).

[8] Andrew W. Senior et al. "Improved protein structure prediction using potentials from deep learning". In: *Nature* 577.7792 (Jan. 2020). Number: 7792 Publisher: Nature Publishing Group, pp. 706–710. URL: https://www.nature.com/articles/s41586-019-1923-7 (visited on 10/27/2023).

[9] Wujie Wang, Simon Axelrod, and Rafael Gómez-Bombarelli. *Differentiable Molecular Simulations for Control and Learning*. Feb. 26, 2020.

[10] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. "ImageNet Classification with Deep Convolutional Neural Networks". In: *Advances in Neural Information Processing Systems*. Vol. 25. 2012.

[11] Mark S. Friedrichs et al. “Accelerating Molecular Dynamic Simulation on Graphics Processing Units”. In: *Journal of Computational Chemistry* 30.6 (2009), pp. 864–872.

[12] James Bradbury et al. *JAX: composable transformations of Python+NumPy programs*. Version 0.3.13. 2018. URL: http://github.com/jax-ml/jax.

[13] Samuel S. Schoenholz and Ekin D. Cubuk. *JAX, M.D.: A Framework for Differentiable Physics*. arXiv:1912.04232 [physics]. Dec. 2020. URL: http://arxiv.org/abs/1912.04232 (visited on 01/08/2026).

[14] Yin Li et al. *pmwd: A Differentiable Cosmological Particle-Mesh N-body Library*. arXiv:2211.09958 [astro-ph]. Nov. 2022. URL: http://arxiv.org/abs/2211.09958 (visited on 01/08/2026).

[15] Michael S. Packer and David R. Liu. “Methods for the directed evolution of proteins”. In: *Nature Reviews Genetics* 16.7 (July 2015). Number: 7 Publisher: Nature Publishing Group, pp. 379–394. URL: https://www.nature.com/articles/nrg3927 (visited on 10/27/2023).

[16] Frances H. Arnold. “Design by Directed Evolution”. In: *Accounts of Chemical Research* 31.3 (Mar. 1, 1998), pp. 125–131. URL: https://pubs.acs.org/doi/10.1021/ar960017f (visited on 09/08/2023).

[17] Cynthia H. Collins, Jared R. Leadbetter, and Frances H. Arnold. “Dual selection enhances the signaling specificity of a variant of the quorum-sensing transcriptional activator LuxR”. In: *Nature Biotechnology* 24.6 (June 2006). Number: 6 Publisher: Nature Publishing Group, pp. 708–712. URL: https://www.nature.com/articles/nbt1209 (visited on 10/28/2023).

[18] Troy E. Sandberg et al. “Evolution of Escherichia coli to 42 °C and subsequent genetic engineering reveals adaptive mechanisms and novel mutations”. In: *Molecular Biology and Evolution* 31.10 (Oct. 2014), pp. 2647–2662.

[19] Brooke J. Bevis and Benjamin S. Glick. “Rapidly maturing variants of the Discosoma red fluorescent protein (DsRed)”. In: *Nature Biotechnology* 20.1 (Jan. 2002). Number: 1 Publisher: Nature Publishing Group, pp. 83–87. URL: https://www.nature.com/articles/nbt0102-83 (visited on 10/28/2023).

[20] Justin R. Meyer et al. “Repeatability and contingency in the evolution of a key innovation in phage lambda”. In: *Science (New York, N.Y.)* 335.6067 (Jan. 27, 2012), pp. 428–432.

[21] Jens Frickel, Michael Sieber, and Lutz Becks. “Eco-evolutionary dynamics in a coevolving host–virus system”. In: *Ecology Letters* 19.4 (2016). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ele.12580, pp. 450–459. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/ele.12580 (visited on 10/28/2023).

[22] Armita Nourmohammad, Jakub Otwinowski, and Joshua B. Plotkin. “Host-Pathogen Coevolution and the Emergence of Broadly Neutralizing Antibodies in Chronic Infections”. In: *PLOS Genetics* 12.7 (July 21, 2016). Publisher: Public Library of Science, e1006171. URL: https://journals.plos.org/plosgenetics/article?id=10.1371/journal.pgen.1006171 (visited on 03/24/2025).

[23] Hua-Xin Liao et al. "Co-evolution of a broadly neutralizing HIV-1 antibody and founder virus". In: *Nature* 496.7446 (Apr. 2013). Publisher: Nature Publishing Group, pp. 469–476. URL: https://www.nature.com/articles/nature12053 (visited on 03/24/2025).

[24] Arindam Nandi and Anita Shet. "Why vaccines matter: understanding the broader health, economic, and child development benefits of routine vaccination". In: *Human Vaccines & Immunotherapeutics* 16.8 (), pp. 1900–1904. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7482790/ (visited on 01/29/2025).

[25] Walter A. Orenstein and Rafi Ahmed. "Simply put: Vaccination saves lives". In: *Proceedings of the National Academy of Sciences* 114.16 (Apr. 2017). Publisher: Proceedings of the National Academy of Sciences, pp. 4031–4033. URL: https://www.pnas.org/doi/10.1073/pnas.1704507114 (visited on 01/29/2025).

[26] Ely Zarina Samsudin et al. "Socioeconomic impacts of airborne and droplet-borne infectious diseases on industries: a systematic review". In: *BMC Infectious Diseases* 24.1 (Jan. 2024), p. 93. URL: https://doi.org/10.1186/s12879-024-08993-y (visited on 01/29/2025).

[27] Ahmad Faramarzi et al. "The global economic burden of COVID-19 disease: a comprehensive systematic review and meta-analysis". In: *Systematic Reviews* 13.1 (Feb. 2024), p. 68. URL: https://doi.org/10.1186/s13643-024-02476-6 (visited on 01/29/2025).

[28] Olivia Gallup et al. "Computational Synthetic Biology Enabled through JAX: A Showcase". In: *ACS Synthetic Biology* 13.9 (Sept. 2024). Publisher: American Chemical Society, pp. 3046–3050. URL: https://doi.org/10.1021/acssynbio.4c00307 (visited on 01/08/2026).

[29] Jessica James et al. *Microscopic PhotoSelection (MiPS) of single cells in mother machine microfluidic devices*. en. ISSN: 2692-8205 Pages: 2025.12.03.692065 Section: New Results. Dec. 2025. URL: https://www.biorxiv.org/content/10.64898/2025.12.03.692065v1 (visited on 01/07/2026).

[30] Jessica James et al. *Microscopic PhotoSelection (MiPS) of single cells in mother machine microfluidic devices*. Reviewed Preprint, version 1. Aug. 28, 2026. URL: https://doi.org/10.7554/eLife.112391.1.

[31] Ken A. Dill and Justin L. MacCallum. "The Protein-Folding Problem, 50 Years On". In: *Science* 338.6110 (Nov. 2012). Publisher: American Association for the Advancement of Science, pp. 1042–1046. URL: https://www.science.org/doi/10.1126/science.1219021 (visited on 10/31/2025).

[32] John Jumper et al. "Highly accurate protein structure prediction with AlphaFold". In: *Nature* 596.7873 (Aug. 2021). Number: 7873 Publisher: Nature Publishing Group, pp. 583–589. URL: https://www.nature.com/articles/s41586-021-03819-2 (visited on 10/28/2023).

[33] Gregory R. Bowman. "AlphaFold and Protein Folding: Not Dead Yet! The Frontier Is Conformational Ensembles". eng. In: *Annual Review of Biomedical Data Science* 7.1 (Aug. 2024), pp. 51–57.

[34] Scott A. Hollingsworth and Ron O. Dror. "Molecular Dynamics Simulation for All". In: *Neuron* 99.6 (Sept. 2018), pp. 1129–1143. URL: https://www.sciencedirect.com/science/article/pii/S0896627318306846 (visited on 11/21/2025).

[35] Outi M. H. Salo-Ahen et al. "Molecular Dynamics Simulations in Drug Discovery and Pharmaceutical Development". en. In: *Processes* 9.1 (Dec. 2020). Company: Multidisciplinary Digital Publishing Institute Distributor: Multidisciplinary Digital Publishing Institute Institution: Multidisciplinary Digital Publishing Institute Label: Multidisciplinary Digital Publishing Institute Publisher: publisher. URL: https://www.mdpi.com/2227-9717/9/1/71 (visited on 11/21/2025).

[36] Samuel Genheden and Ulf Ryde. "The MM/PBSA and MM/GBSA methods to estimate ligand-binding affinities". In: *Expert Opinion on Drug Discovery* 10.5 (May 2015), pp. 449–461. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC4487606/ (visited on 11/21/2025).

[37] Niels Hansen and Wilfred F. van Gunsteren. "Practical Aspects of Free-Energy Calculations: A Review". In: *Journal of Chemical Theory and Computation* 10.7 (July 2014). Publisher: American Chemical Society, pp. 2632–2647. URL: https://doi.org/10.1021/ct500161f (visited on 11/21/2025).

[38] Richard Evans et al. *Protein complex prediction with AlphaFold-Multimer*. en. Pages: 2021.10.04.463034 Section: New Results. Mar. 2022. URL: https://www.biorxiv.org/content/10.1101/2021.10.04.463034v2 (visited on 11/21/2025).

[39] Patrick Bryant. "Deep learning for protein complex structure prediction". In: *Current Opinion in Structural Biology* 79 (Apr. 2023), p. 102529. URL: https://www.sciencedirect.com/science/article/pii/S0959440X23000039 (visited on 11/21/2025).

[40] Kenneth M. Murphy and Casey Weaver. *Janeway's Immunobiology*. 9th ed. New York, NY: Garland Science, Taylor & Francis Group, 2017. URL: https://www.ncbi.nlm.nih.gov/nlmcatalog/101674267.

[41] S. Tonegawa. "Somatic generation of antibody diversity". eng. In: *Nature* 302.5909 (Apr. 1983), pp. 575–581.

[42] Markus Engstler et al. "Hydrodynamic Flow-Mediated Protein Sorting on the Cell Surface of Trypanosomes". In: *Cell* 131.3 (2007), pp. 505–515. URL: https://doi.org/10.1016/j.cell.2007.08.046.

[43] Mark L. Chiu et al. "Antibody Structure and Function: The Basis for Engineering Therapeutics". In: *Antibodies* 8.4 (Dec. 2019), p. 55. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC6963682/ (visited on 10/31/2025).

[44] Ahmed Rohaim et al. “Engineering of a synthetic antibody fragment for structural and functional studies of K+ channels”. In: *Journal of General Physiology* 154.4 (Mar. 2022), e202112965. URL: https://doi.org/10.1085/jgp.202112965 (visited on 10/31/2025).

[45] Jin Cheng et al. “A new era of antibody discovery: an in-depth review of AI-driven approaches”. In: *Drug Discovery Today* 29.6 (June 2024), p. 103984. URL: https://www.sciencedirect.com/science/article/pii/S1359644624001090 (visited on 10/31/2025).

[46] Frances H. Arnold. “Directed Evolution: Bringing New Chemistry to Life”. In: *Angewandte Chemie International Edition* 57.16 (2018). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/anie.201708408, pp. 4143–4148. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/anie.201708408 (visited on 10/31/2025).

[47] Dan Li et al. “Directed evolution of a picomolar-affinity, high-specificity antibody targeting phosphorylated tau”. In: *The Journal of Biological Chemistry* 293.31 (Aug. 2018), pp. 12081–12094. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC6078456/ (visited on 10/27/2025).

[48] Bradley S. Heater et al. “Directed evolution of a genetically encoded immobilized lipase for the efficient production of biodiesel from waste cooking oil”. In: *Biotechnology for Biofuels* 12.1 (June 2019), p. 165. URL: https://doi.org/10.1186/s13068-019-1509-5 (visited on 10/27/2025).

[49] S. B. Jennifer Kan et al. “Directed Evolution of Cytochrome c for Carbon–Silicon Bond Formation: Bringing Silicon to Life”. In: *Science (New York, N.Y.)* 354.6315 (Nov. 2016), pp. 1048–1051. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC5243118/ (visited on 10/27/2025).

[50] Kevin K. Yang, Zachary Wu, and Frances H. Arnold. “Machine-learning-guided directed evolution for protein engineering”. en. In: *Nature Methods* 16.8 (Aug. 2019). Publisher: Nature Publishing Group, pp. 687–694. URL: https://www.nature.com/articles/s41592-019-0496-6 (visited on 10/27/2025).

[51] Surojit Biswas et al. “Low-N protein engineering with data-efficient deep learning”. en. In: *Nature Methods* 18.4 (Apr. 2021). Publisher: Nature Publishing Group, pp. 389–396. URL: https://www.nature.com/articles/s41592-021-01100-y (visited on 10/27/2025).

[52] Lutz Becks and Aneil F. Agrawal. “The Evolution of Sex Is Favoured During Adaptation to New Environments”. In: *PLOS Biology* 10.5 (2012), e1001317. URL: https://doi.org/10.1371/journal.pbio.1001317.

[53] S. Wright. “The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution”. In: *Proceedings of the Sixth International Congress of Genetics*. Vol. 1. Ithaca, NY, 1932, pp. 356–366.

[54] Frank J. Poelwijk et al. “Reciprocal sign epistasis is a necessary condition for multi-peaked fitness landscapes”. In: *Journal of Theoretical Biology* 272.1 (Mar. 2011), pp. 141–144. URL: https://www.sciencedirect.com/science/article/pii/S0022519310006703 (visited on 10/30/2025).

[55] Daniel M. Weinreich, Richard A. Watson, and Lin Chao. “Perspective: Sign epistasis and genetic constraint on evolutionary trajectories”. eng. In: *Evolution; International Journal of Organic Evolution* 59.6 (June 2005), pp. 1165–1174.

[56] T. Aita, M. Iwakura, and Y. Husimi. “A cross-section of the fitness landscape of dihydrofolate reductase”. eng. In: *Protein Engineering* 14.9 (Sept. 2001), pp. 633–638.

[57] Ivan G. Szendro et al. “Quantitative analyses of empirical fitness landscapes”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2013.01 (Jan. 2013). arXiv:1202.4378 [q-bio], P01005. URL: http://arxiv.org/abs/1202.4378 (visited on 10/30/2025).

[58] Jasper Franke et al. “Evolutionary Accessibility of Mutational Pathways”. en. In: *PLOS Computational Biology* 7.8 (Aug. 2011). Publisher: Public Library of Science, e1002134. URL: https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1002134 (visited on 10/27/2025).

[59] J. F. C. Kingman. “A Simple Model for the Balance between Selection and Mutation”. In: *Journal of Applied Probability* 15.1 (1978). Publisher: Applied Probability Trust, pp. 1–12. URL: https://www.jstor.org/stable/3213231 (visited on 10/27/2025).

[60] Julien Berestycki, Éric Brunet, and Zhan Shi. “The number of accessible paths in the hypercube”. In: *Bernoulli* 22.2 (2016). Publisher: [Bernoulli Society for Mathematical Statistics and Probability, International Statistical Institute (ISI)], pp. 653–680. URL: https://www.jstor.org/stable/43864013 (visited on 10/27/2025).

[61] Jakob Foerster et al. “Learning with Opponent-Learning Awareness”. In: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2018, pp. 122–130. URL: https://dl.acm.org/doi/10.5555/3237383.3237395 (visited on 11/25/2024).

[62] Jonathan Letcher et al. “Stable Opponent Shaping in Differentiable Games”. In: *Advances in Neural Information Processing Systems*. Vol. 32. 2019. URL: https://proceedings.neurips.cc/paper_files/paper/2019/file/f4b53633c713eef69431e00fe4341dda-Paper.pdf (visited on 11/25/2024).

[63] Martín Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015. URL: https://www.tensorflow.org/.

[64] Adam Paszke et al. *PyTorch: An Imperative Style, High-Performance Deep Learning Library*. en. Dec. 2019. URL: https://arxiv.org/abs/1912.01703v1 (visited on 11/11/2025).

[65] Robert Tjarko Lange. *gymnax: A JAX-based Reinforcement Learning Environment Library*. Version 0.0.4. 2022. URL: http://github.com/RobertTLange/gymnax.

[66] W. A. Bonner et al. “Fluorescence Activated Cell Sorting”. In: *Review of Scientific Instruments* 43.3 (Mar. 1972). _eprint: https://pubs.aip.org/aip/rsi/article-pdf/43/3/404/19177575/404_1_online.pdf, pp. 404–409. URL: https://doi.org/10.1063/1.1685647.

[67] Ping Wang et al. “Robust Growth of Escherichia coli”. English. In: *Current Biology* 20.12 (June 2010). Publisher: Elsevier, pp. 1099–1103. URL: https://www.cell.com/current-biology/abstract/S0960-9822(10)00524-5 (visited on 01/08/2026).

[68] Andreas Milias-Argeitis et al. “In silico feedback for in vivo regulation of a gene expression circuit”. en. In: *Nature Biotechnology* 29.12 (Dec. 2011). Publisher: Nature Publishing Group, pp. 1114–1116. URL: https://www.nature.com/articles/nbt.2018 (visited on 01/08/2026).

[69] Remy Chait et al. “Shaping bacterial population behavior through computer-interfaced control of individual cells”. en. In: *Nature Communications* 8.1 (Nov. 2017). Publisher: Nature Publishing Group, p. 1535. URL: https://www.nature.com/articles/s41467-017-01683-1 (visited on 01/08/2026).

[70] Jean-Baptiste Lugagne, Haonan Lin, and Mary J. Dunlop. “DeLTA: Automated cell segmentation, tracking, and lineage reconstruction using deep learning”. en. In: *PLOS Computational Biology* 16.4 (Apr. 2020). Publisher: Public Library of Science, e1007673. URL: https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1007673 (visited on 01/08/2026).

[71] Jessica James et al. “Optimisation strategies for directed evolution without sequencing”. In: *PLOS Computational Biology* 20.12 (2024), e1012695. URL: https://doi.org/10.1371/journal.pcbi.1012695.

[72] Sebastian Towers et al. “Sequence-free landscape inference for directed evolution”. In: *PLOS Computational Biology* 22.9 (2026), e1014713. URL: https://doi.org/10.1371/journal.pcbi.1014713.

[73] Sebastian Rene Towers et al. “ADIOS: Antibody Development via Opponent Shaping”. In: *Proceedings of the 42nd International Conference on Machine Learning*. Vol. 267. Proceedings of Machine Learning Research. PMLR, 2025, pp. 59904–59923. URL: https://proceedings.mlr.press/v267/towers25a.html.

[74] Google DeepMind. *AlphaFold 3: Performance*. Official documentation; model-inference benchmarks exclude compilation and the data pipeline. Google DeepMind. URL: https://github.com/google-deepmind/alphafold3/blob/main/docs/performance.md (visited on 09/20/2026).

[75] Isaiah Berlin. *The Hedgehog and the Fox: An Essay on Tolstoy’s View of History*. English. Ed. by Henry Hardy. 2nd. Princeton University Press, June 2013.