
OPPONENT SHAPING FOR ANTIBODY DEVELOPMENT

A PREPRINT

Sebastian Towers*^{1,2}, Aleksandra Kalisz*^{1,2}, Philippe A. Robert³, Alicia Higuero⁴,
Francesca Vianello⁵, Ming-Han Chloe Tsai⁶, Harrison Steel², and Jakob Foerster^{1,2}

¹Foerster Lab for AI Research, University of Oxford

²Department of Engineering Science, University of Oxford

³Department of Biomedicine, University of Basel

⁴Isomorphic Labs, London

⁵Exscientia, Oxford

⁶Epsilon Ltd., London

ABSTRACT

Anti-viral therapies are typically designed or evolved towards the current strains of a virus. In learning terms, this corresponds to a *myopic* best response, i.e., not considering the possible adaptive moves of the opponent. However, therapy-induced selective pressures act on viral antigens to drive the emergence of mutated strains, against which initial therapies have reduced efficacy. To motivate our work, we consider antibody designs that target not only the current viral strains but also the wide range of possible future variants that the virus might evolve into under the evolutionary pressure exerted by said antibodies. Building on a computational model of binding between antibodies and viral antigens (the Absolut! framework), we design and implement a genetic simulation of the viral evolutionary escape. Crucially, this allows our antibody optimisation algorithm to consider and influence the entire escape curve of the virus, i.e. to guide (or “shape”) the viral evolution. This is inspired by opponent shaping which, in general-sum learning, accounts for the *adaptation* of the co-player rather than playing a myopic best response. Hence we call the optimised antibodies *shapers*. Within our simulations, we demonstrate that our shapers target both current and simulated future viral variants, outperforming the antibodies chosen in a myopic way. Furthermore, we show that shapers exert specific evolutionary pressure on the virus compared to myopic antibodies. Altogether, shapers modify the evolutionary trajectories of viral strains and minimise the viral escape compared to their myopic counterparts. While this is a simple model, we hope that our proposed paradigm will enable the discovery of better long-lived vaccines and antibody therapies in the future, enabled by rapid advancements in the capabilities of simulation tools.

Keywords Machine Learning · Game Theory · Antibody Design · Opponent Shaping · Computational Biology

1 Introduction

Designing effective therapies to fight off viral pathogens is crucial. However, traditional approaches have been largely myopic, focusing on the *current* variant of a virus without considering the future evolutionary pressures induced by the therapies themselves. Although this *shortsighted* approach may yield therapies with high initial efficacy, it fails to account for the virus’s adaptive response, leaving treatments vulnerable to viral escape and the emergence of new therapy-resistant strains [1, 2, 3, 4, 5].

The COVID-19 pandemic starkly illustrated this challenge. While the rapid development of vaccines was a remarkable achievement, concerns quickly arose about their long-term efficacy against new emerging COVID variants [6, 7]. For example, the B.1.351 variant demonstrated that the vaccine loses its effectiveness against new strains [8]. This underscores the need for approaches that consider both the *current and future* efficacy of a designed therapy.

Toward this goal, we propose to design therapies that target not only current viral variants but also anticipate and target future ones using a model of evolutionary adaptation. Specifically, viruses evolve in response to the selective

*Equal contribution. Correspondence to [sebastian.towers, aleksandra.kalisz]@eng.ox.ac.uk

pressures imposed by our therapies implying that the therapies we design exert a *degree of influence* over viral evolution. Recognizing this, our goal is to develop therapies which steer the viral evolution away from escape.

This concept aligns closely with the principles of *opponent shaping* [9]. Opponent shaping is a multi-agent reinforcement learning framework that allows agents to anticipate and influence the future strategies of other agents in their environment. This approach, exemplified by methods such as Learning with Opponent-Learning Awareness (LOLA) [9] and Model-Free Opponent Shaping (M-FOS) [10], encourages agents to consider not only their immediate rewards but also the consequences of their actions on their opponents’ future behaviour. In the context of viral therapy, we view the virus as an adversary that adapts to our chosen therapy strategy. By applying opponent shaping techniques to this problem, we suggest optimising therapeutics to be *more robust to future viral mutations* and *shape viral evolution in favourable directions*.

Building on these principles, our work focuses specifically on *antibody design*, framing the interaction between antibodies and surface viral antigens as a *two-player zero-sum game* which we call the *virus-antibody game*. Importantly, we include key functional constraints of the antibody and the virus in the game. Our framework applies opponent shaping techniques to optimise antibodies for long-term effectiveness against evolving viral populations, limiting viral escape.

Viral escape occurs when mutations in the virus allow it to successfully evade the host’s immune system [11]. In the context of our game, the virus reduces or eliminates the antibody’s ability to bind effectively, i.e. it finds an effective best-response strategy to the current antibody. We construct a simulated model of viral escape that captures the adaptive responses to antibody pressure, using the Absolut! framework [12] to approximate the binding strength of protein-protein interactions. Importantly, our simulations reveal that the conventionally “optimal” choice of antibody — one that myopically maximises binding to the current virus — is suboptimal when considering a longer time horizon of viral antigen mutations.

Based on this insight, our algorithm instead optimises antibodies with respect to their expected fitness over the entire viral escape curve. We term these optimised antibodies *shapers*, in contrast to *myopic* antibodies that are selected based solely on their efficacy against the current viral variant. Our results demonstrate that shaper antibodies not only outperform myopic antibodies in long-term efficacy but also exert a shaping influence on viral evolutionary trajectories, effectively constraining the virus’s ability to escape antibody binding.

Our study also explores the trade-offs between the effectiveness of shaper antibodies and the computational resources required for their optimisation, providing insights into the scalability and potential for practical deployment of our method.

Finally, we present an analysis of the key features that distinguish shaper antibodies from their myopic counterparts. This explainability component provides insights into the amino acid sequence characteristics that contribute to long-term antibody efficacy and viral evolution shaping within our simulations.

In summary, this work introduces a novel and general antibody design framework that leverages opponent-shaping principles. By anticipating and influencing expected viral evolution, our method improves the long-term efficacy of antibodies within our Absolut! simulations, and when combined with more accurate future simulators may assist in the design of antiviral therapies that minimise the rate of viral escape.

2 Results

2.1 Computational Implementation of Opponent Shaping for Antibody Design

To implement opponent shaping for antibody design, we develop a computational framework that models the virus-antibody interaction and optimises antibodies for long-term efficacy. Our approach comprises three main components: defining the virus-antibody game with the player’s payoffs, simulating viral escape, and optimising antibody shapers.

2.1.1 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, specifically the third complementarity-determining region (CDRH3), which is the most important part of the antibody for defining its specificity and affinity [15].

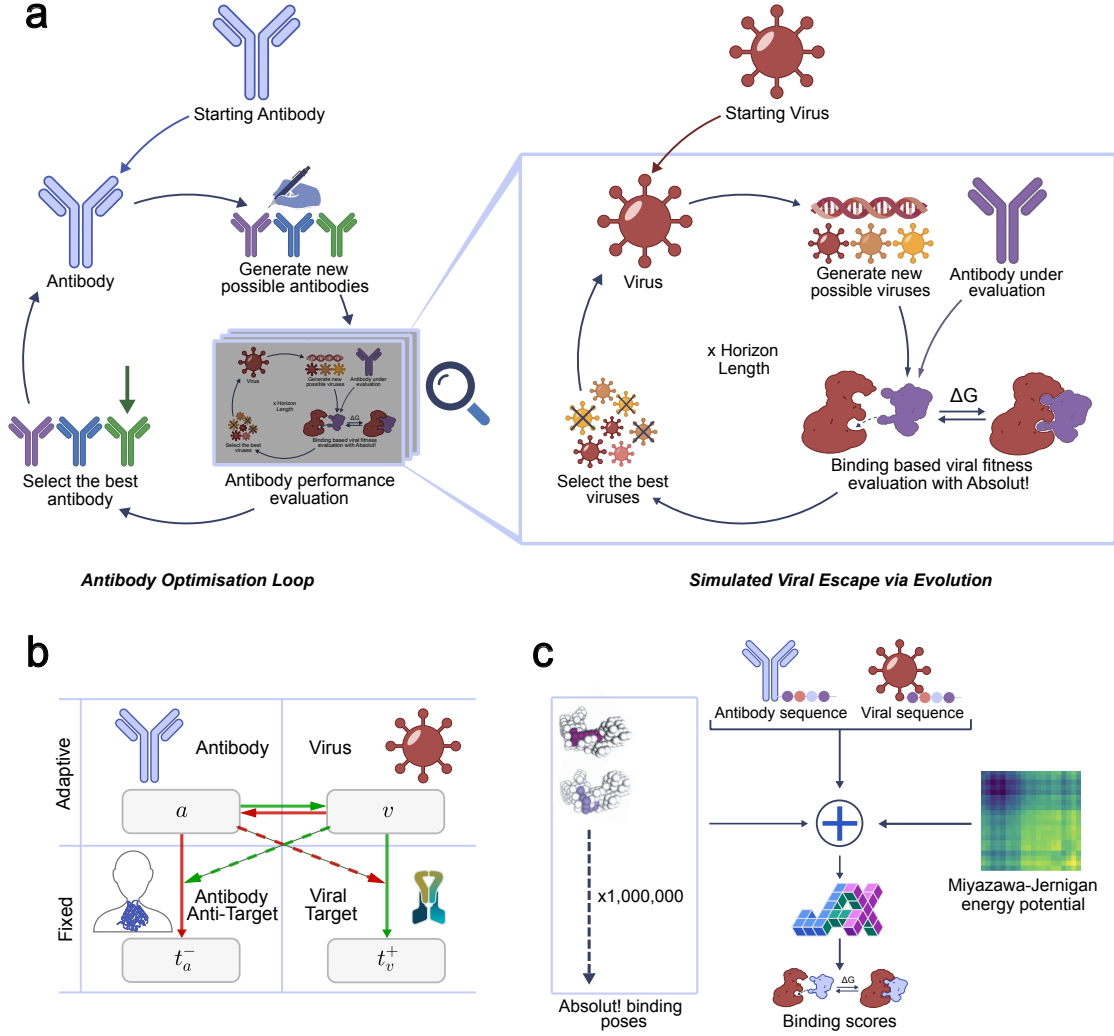


Figure 1: **Computational Implementation of Opponent Shaping for Antibody Design.** **a** An overview of our implementation. In the *Antibody Optimisation Loop* (i.e. the outer loop), we optimise the antibody to perform well against the current and future virus variants and influence the viral evolution. We approximate the future variants through the *Simulated Viral Escape via Evolution* (i.e. the inner loop) where the viral antigens are evolved against 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 minimize, while green arrows represent those they aim to maximize. 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. **c** Efficient JAX [13] implementation of the binding calculation uses binding poses generated by Absolut! [12] and the Miyazawa-Jernigan energy potential matrix [14].

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 host 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:

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

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 1b. This formulation also ensures that neither player can adopt an overly simplistic strategy: the antibody can’t become universally “sticky” due to the anti-target penalty, and the virus can’t become entirely inert without losing its ability to infect host cells. This game structure forms the foundation for our method, see Methods section 4.1 for more details.

To evaluate the binding strength $B(\cdot, \cdot)$ in our setup, we adapt the Absolut! framework [12]. Absolut! is a simplified simulator for antibody-antigen binding. Unlike sequence-based machine learning models [16, 17, 18, 19] that cannot generalise to novel viral mutations, or computationally intensive molecular dynamics simulations [20], Absolut! offers a fast approach that generalises to both viral and antibody mutations, which is necessary for our application.

Specifically, the Absolut! framework operates by discretising the antigen structure, and evaluating the interaction between the antigen and the CDRH3 region of the antibody’s heavy chain - the most variable portion of the antibody [15]. It enumerates all possible binding poses of the CDRH3 to the discretised antigen and computes their binding energy using the Miyazawa–Jernigan energy potential [14]. Absolut! outputs the binding energy E of the lowest energy pose as the binding energy of a query antibody-antigen complex. The binding strength is high when the binding energy is low. Thus we define our binding function as $B(\cdot, \cdot) = -E(\cdot, \cdot)$, see Methods section 4.4 for more details.

To meet the computational demands of our opponent shaping approach, which requires rapid evaluation of numerous antibody-antigen interactions, we reimplement the core binding calculation of Absolut! using JAX [13], a framework that facilitates GPU-accelerated computation, see Figure 1c. 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	2.1×10^{-4}

Table 1: Speed comparison of the binding calculation for a single antibody-antigen query between the original implementation of Absolut! [12] and our reimplement of Absolut! in JAX [13].

2.1.2 Simulated Viral Escape via Evolution

To simulate the space of possible viral escape trajectories against a chosen antibody, we model the virus learning to play the game through an evolutionary algorithm. This approach approximates the process of viral evolution, where mutations that increase fitness (in our case, the virus’s payoff R_v) are more likely to persist, propagate, and become an ancestor of future generations.

We simulate the process of *viral escape* for a specified horizon of H steps, representing multiple generations of viral evolution. This is the target duration of therapy efficacy. For each independent viral trajectory, we use the initial virus strain v to generate a population of P virus variants through random mutations. These are point mutations in the amino acid sequence of the viral antigen and they are introduced according to a predefined mutation rate. Each variant’s fitness is evaluated based on its payoff R_v against the fixed antibody a . Importantly, this process captures the myopic nature of viral evolution: the virus optimises its fitness only with respect to the current antibody, unable to anticipate future human interventions or antibody changes. The next generation of the viruses is then created by probabilistically selecting the variants based on their fitness, with fitter variants more likely to be chosen for replication and subsequent mutation. This process repeats for H generations, creating one trajectory of viral escape, see the *Simulated Viral Escape via Evolution* in Figure 1a and Methods section 4.2 for more details.

We use the antigen protein from the Dengue Virus as the virus for all our experiments, specifically, the structure with PDB code 2R29 [21]. We discretise this structure using the Absolut! framework to generate the docking poses, which we then use to calculate the binding strength B of any antibody-virus pair using our JAX implementation. In the viral escape step, we mutate only the amino acid sequence of the dengue envelope antigen, which is composed of 97 amino acids, and do not consider other components of Dengue Virus. Importantly, while we alter the amino acid sequence, we assume the structure of the antigen does not significantly change for the calculation of binding strength.

2.1.3 Optimisation of Antibody Shapers

Our antibody optimisation process aims to generate antibodies that can prevent viral escape. We start with a random antibody sequence and generate a population of antibodies of size P_a through mutations in the amino acid sequence. Each new antibody in this population is created by introducing a single mutation to the original sequence.

We define the antibody fitness function $F_v^H(a)$ as the antibody’s performance over multiple potential viral escape trajectories of virus v :

$$F_v^H(a) = \mathbb{E}_{\eta \text{ escape trajectories of } v} \left[\mathbb{E}_{H \text{ steps of escape}} (R_a(v', a)) \right] \quad (1)$$

Specifically, for each of the η trajectories, we simulate the viral escape process for H steps and calculate the average antibody payoff R_a over these steps, where v' represents the possible mutations of the original virus v . The final fitness is the average over these η simulations, i.e. η Monte Carlo roll-outs. The fitness function $F_v^H(a)$ captures the average antibody’s long-term effectiveness against evolving viral populations.

The optimisation proceeds iteratively for N steps. At each step, we select the antibody with the highest fitness and generate a new population of P_a antibodies by introducing single mutations to the selected antibody. This process of evaluation, selection, and mutation continues, gradually improving the antibody’s robustness against, and ability to limit, viral escape, see *Antibody Optimisation Loop* in Figure 1a and Methods section 4.3 for more details.

It’s worth noting the computational intensity of this process: a single step of antibody optimisation requires simulating $P_a \times \eta$ independent viral escapes, each for H steps. Consequently, we refer to different optimised antibodies as $H = 5$ *shapers* or $H = 100$ *shapers*, reflecting the horizon length chosen for the viral escape simulations during their optimisation. *Myopic* antibodies are antibodies where horizon length $H = 0$ and thus $F_v^0(a) = R_a(v, a)$. The horizon length significantly impacts the antibody’s long-term effectiveness and the computational resources required for optimisation.

This process can be viewed as meta-optimisation or meta-learning. The *inner loop* of our optimisation is the *viral escape simulation*, where we model the virus learning to evade the current antibody. The *outer loop* or *meta-loop* is the antibody optimisation process itself, which learns to generate antibodies that perform well across many possible viral escape trajectories. We are designing antibodies that anticipate and influence the escape process of the virus, embodying the core principle of opponent shaping in this biological context.

2.2 Antibody Shapers Minimise Viral Escape

2.2.1 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 . The myopic antibodies optimise the fitness function $F_v^0(a) = R_a(v, a)$, which considers only the immediate payoff against the non-mutated virus. For our shaper antibodies, we select a long horizon of $H = 100$ to capture extended viral escape trajectories. Both shapers and myopic antibodies are optimised for $N = 30$ steps. Figure 2a presents the performance distributions of shapers and myopic antibodies under both objective functions.

Our results demonstrate a clear advantage for antibody 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 2a. Notably, none of the myopic antibodies outperforms 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 compared to myopic antibodies in optimising the immediate payoff $R_a(v, a)$ against the initial, non-mutated virus v .

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 2b 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.

2.2.2 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

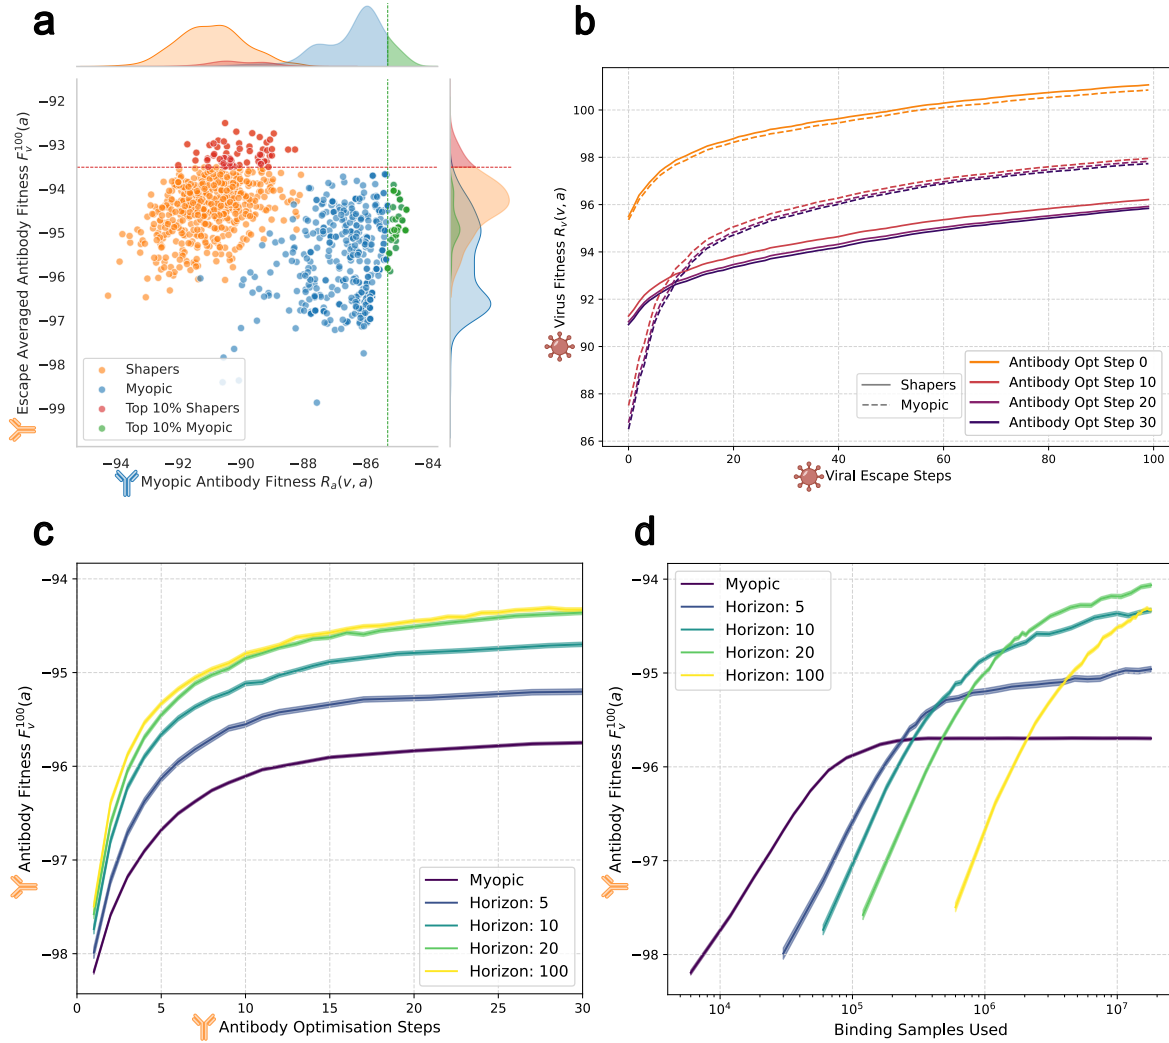


Figure 2: **Optimised antibody shapers outperform myopic antibodies.** **a** Distribution of antibody shapers (in orange) which were optimised with horizon $H = 100$ vs. myopic antibodies distribution (in blue). We highlight the top 10% shapers with respect to $F_v^{100}(a)$ in red and the top 10% myopic 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 over 100 steps of viral escape. Higher values on both axes indicate better performance. **b** Viral escape curves for different steps of the antibody optimisation process 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). **c, d** Antibody optimisation learning curves for a varying horizon length. The x-axis shows the antibody optimisation steps (**c**) and the number of samples from the binding simulator (**d**). The y-axis represents the escape averaged antibody fitness over 100 steps of viral escape $F_v^{100}(a)$. Error bars correspond to the standard error. Higher values indicate better performance.

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 2c 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 antibody objective $F_v^{100}(a)$. Consequently, the optimal horizon varies depending on the available computational budget.

To illustrate this trade-off, we conduct an additional experiment shown in Figure 2d. 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.

We find that the optimisation horizon significantly influences the performance of antibody shapers. While longer horizons generally 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.

2.3 How the Antibody Shapers Prevent Viral Escape

2.3.1 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.

To discern which strategy our antibody shapers adopt, we investigate whether shapers optimised with different horizons exert different types of evolutionary pressure on the virus. We generate antibodies for each horizon (*Myopic*, $H = 5$, $H = 10$, $H = 20$, $H = 100$) and simulate viral escape against each antibody for 100 steps, resulting in viruses v_H . We then evaluate the antibody payoff $R_a(v_H, a_{H'})$ for each horizon antibody against the final viral variants v_H induced by other antibodies $a_{H'}$. Figure 3 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 a_{100} shapers 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 right most column of Figure 3). This trade-off indicates that to exert a stronger shaping influence on viral evolution, $H = 100$ shapers sacrifice 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: some optimised for shaping the virus’s evolutionary trajectory, and others designed for strong immediate binding.




Figure 3: **Effect of myopic and shaper antibodies on viral antigens induced by other antibodies.** We optimise 80 different antibodies for each horizon (*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 escaped 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 against the final escape viral variant v_H induced by other antibodies $a_{H'}$. Darker colours correspond to better antibody payoff.

2.3.2 Amino Acid Distribution in Shapers and Myopic Antibodies

To further understand the performance differences between myopic antibodies and shapers, we analyse the amino acid distributions of the antibodies. We compared myopic antibodies with shapers optimised using horizons $H = 5, 10, 20, 100$ over $N = 30$ antibody optimisation steps. Figure 4 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.




Figure 4: **Distribution of amino acids in myopic antibodies and shapers.** The antibodies are optimised using different viral escape horizons H . Longer horizon shapers push the amino acid distribution closer to a uniform distribution.

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.

2.3.3 Influence of Antibody Shapers on Binding Poses

In the Absolut! framework [12], 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 [14]. 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 Methods Section 4.4 for more details. Importantly, only a small part of the antigen 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 matrix² where one dimension corresponds to the antibody amino acids and the other to the viral amino acids, both contributing to the lowest energy pose. The entries in this matrix represent the number of interactions between specific amino acid pairs.

Figure 5a presents average pose matrices from multiple optimisation runs of both myopic antibodies and long-horizon shapers. We observe two key trends:

1. 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.
2. 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.

²20 is the number of possible amino acids.

To further understand these pose dynamics, we aggregate the pose matrices along the antibody axis (Figure 5b) and the virus axis (Figure 5c). These figures show the change in interaction counts between viral escape steps 0 and 100.

Figure 5b reveals a viral escape strategy: including more of the antibody’s lowest binding amino acids (particularly K) in the pose. Notably, long-horizon shapers, especially $H = 100$ shapers, are most effective at preventing this increase in K interactions. Furthermore, Figure 5c shows another viral escape strategy: removing its high-binding amino acids (I and M) from the pose. Again, $H = 100$ shapers are most successful in counteracting this trend, although they cannot completely prevent it.

Based on these observations, we hypothesise that the shaping ability of $H = 100$ shapers relies on two main mechanisms:

1. Preventing the virus from including the antibody’s lowest binding amino acids in the pose.
2. 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 long-horizon shapers to produce antibodies that more effectively prevent viral escape.

The transferability of these findings to more realistic binding simulators and viral escape models remains an open question. However, these results are encouraging: we have demonstrated that shapers not only perform better and shape viral escape trajectories, but we can also explain the strategies they employ to achieve this shaping effect. This analysis underscores the potential of opponent shaping in antibody design and opens new avenues for developing more effective, long-term strategies against rapidly evolving pathogens.

3 Discussion

This study introduces a novel framework for designing antibodies using the principle of opponent shaping [9] that can effectively combat evolving pathogens. The core of our framework is a two-player game that models the interaction between antibodies and viruses. Applying opponent shaping to the game allows us to optimise for antibody shapers that anticipate and influence viral evolution.

The proposed framework is adaptable. While we demonstrated its efficacy using the Absolut! [12] binding simulator accelerated with JAX [13] for payoff estimation, and our model of the simulated viral escape with an evolutionary process, these components are interchangeable. As more accurate and efficient binding simulators or viral evolution models become available, they can be integrated into our framework, potentially enhancing its biological applicability.

Our results demonstrate a clear advantage of antibody shapers over myopic antibodies when considering long horizons of viral escape. We observed that myopic antibodies, while initially effective, struggle to maintain their efficacy as the virus evolves. In contrast, shaper antibodies, particularly those optimized for longer horizons, show sustained performance over extended periods of viral evolution, see Figure 2.

A notable computational trade-off emerges in the performance of antibody shapers. While longer optimisation horizons generally lead to better long-term performance, they come at an increased computational cost. Figure 2d offers guidance for selecting optimal optimisation horizons based on available computational resources and runtime of a binding simulation which is particularly valuable for future applications of our framework.

Furthermore, our analysis reveals that long-horizon shapers shape the virus into more exploitable variants as opposed to just becoming robust to any virus, see Figure 3. In fact, the short-horizon shapers bind stronger to the viruses included by the long-horizon shapers than the long-horizon shapers themselves. This suggests exploring a mixture of antibodies, some optimised for virus shaping and others for binding. This could leverage the benefits of both shaping the viral evolutionary landscape and maintaining strong binding efficacy. Interestingly, a longitudinal study on rare patients who naturally developed broadly neutralizing antibodies (bnAbs) to HIV showed that early neutralizing antibodies induce HIV escape to those antibodies [22]. The bnAbs were later developed with the cooperation between the B cell lineages. This indicates the constant evolving dynamics between the virus and our immune system in nature. In the future, our approach could enable us to better understand and develop therapies for similar settings by simulating a game with more than one antibody playing against the virus.

To provide insights into the mechanisms underlying the effectiveness of antibody shapers, we conduct an explainability analysis focused on the amino acid distribution in shapers (Figure 4) and on binding poses (Figure 5). While these results offer insights into the strategies employed by our antibody shapers, it’s important to note that they are specific to

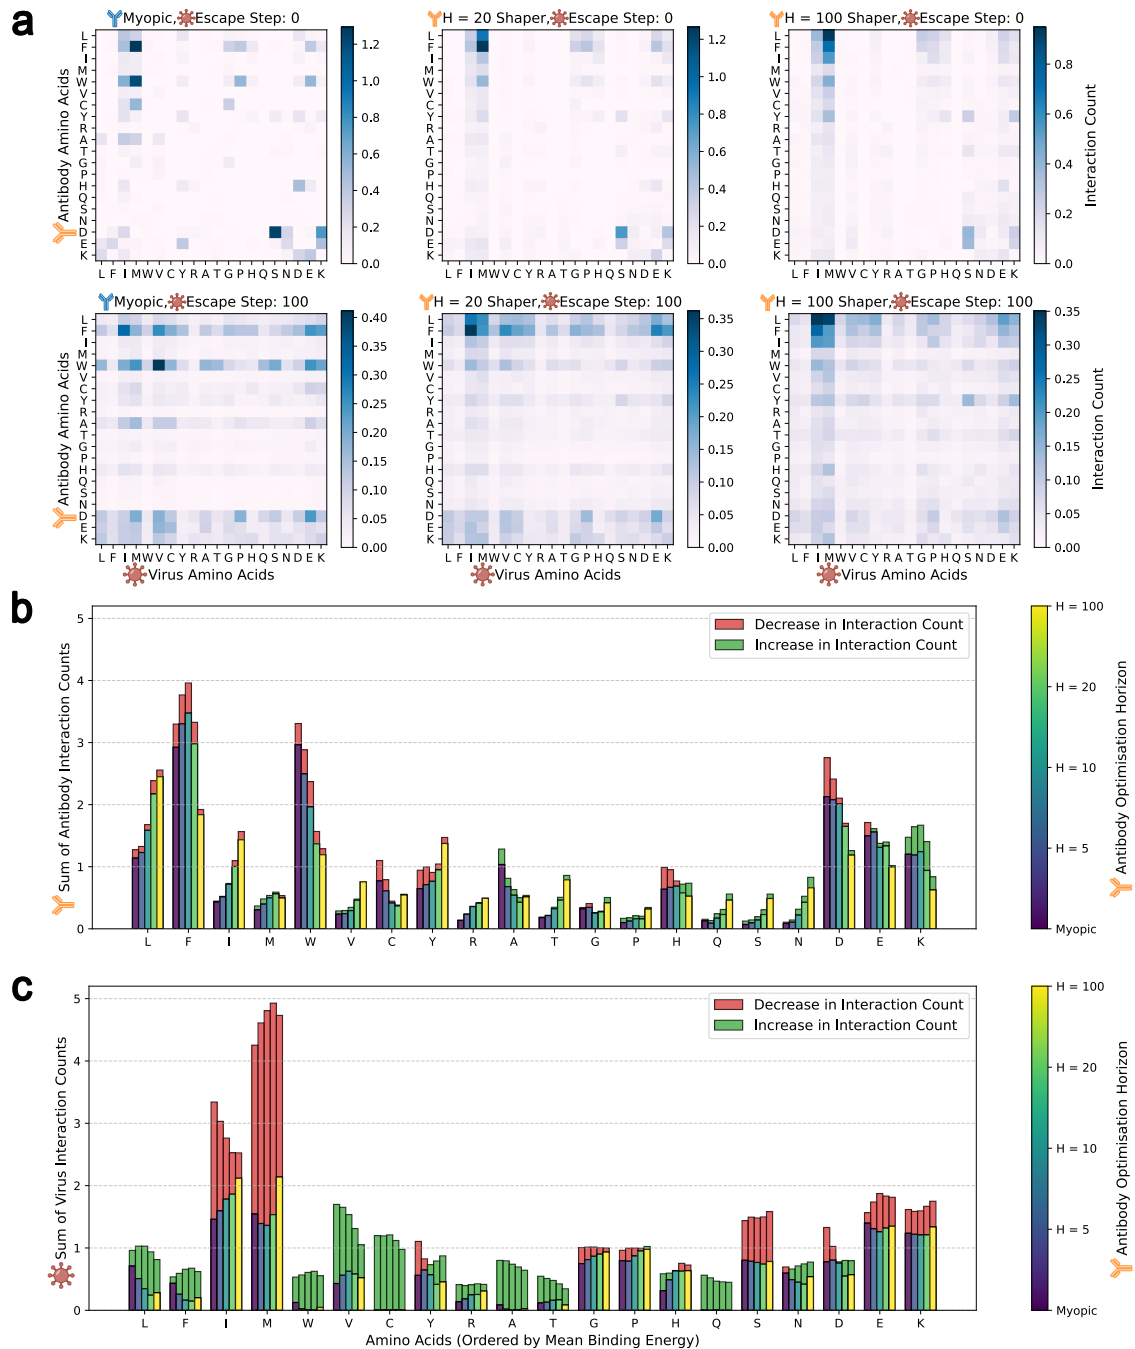


Figure 5: **Influence of antibody shapers on binding poses.** **a** Average pose matrices between antibodies optimised using different horizons and the virus at various stages of its escape. The escape steps increase from left to right and the horizon increases from top to bottom. The full grid of matrices with more antibody horizons and virus escape steps is available in the Supplementary Information, Figure 7. **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.

the Absolut! binding model [12, 14]. Antibodies designed using more accurate binding simulators might be able to employ other strategies in their interactions with viruses. Nevertheless, the fact that we can explain the behaviour of our antibody shapers is encouraging. It suggests that, even as we move towards more complex and realistic models, we may be able to maintain a degree of interpretability in our approach.

While our approach shows promise, it’s important to acknowledge its current limitations, particularly in the modelling of antibody-antigen interactions. Our study relies on the Absolut! framework which is a dramatically simplified model of reality.[12]. Moreover, here we focus only on the CDRH3 binding domain of the antibody and consider a fixed length (11 amino acid) region. Insertions and deletions are not within our scope of the study. Regardless of the CDRH3 region being the most critical for both antibody specificity and affinity [15], the other five binding regions (CDR1-3 on both the heavy and light chains) also partially contribute to the binding profiles, either by making direct contacts with the antigen or putting constraints on the conformational space. In future work, we want to extend the binding simulation to all the CDR regions with both light and heavy chains.

A significant limitation of our current approach is the static structure of the viral antigen. In reality, mutations to the virus induce changes in the structure of the viral antigen. However, in our experiments, we consistently use the original structure of the 2R29 dengue viral antigen [21] and a fixed set of Absolut!-generated docking poses based on that structure. As we introduce mutations to both the virus and the antibody, the identity of the lowest-energy pose changes, but the underlying set of possible poses remains constant. This simplification potentially overlooks important structural dynamics that influence antibody-antigen interactions. Ideally, we could model the structural changes that occur with each introduced mutation to capture a more realistic representation of the evolving virus-antibody interaction landscape. AlphaFold3 has demonstrated the capability to predict structures of antibody-antigen complexes from their sequences with high accuracy [23]. Incorporating methods like AlphaFold3 into our framework would allow for the dynamic updating of antigen structures in response to our simulated mutations. This would enhance the biological relevance of our method.

Another promising avenue for enhancing our framework lies in the integration of more accurate viral escape forecasting methods such as EVEscape [24]. EVEscape successfully anticipated the emergence of COVID-19 variants using only pre-pandemic data. Incorporating such sophisticated viral escape forecasting as the simulated viral escape component of our framework could significantly boost its real-world applicability. This integration can help to bridge the gap between our computational approach and practical therapeutic development.

The implications of our work extend beyond the specific context of antiviral therapies. While our experiments focus on the dengue virus, our framework is adaptable to any evolving pathogen or even the heterogeneity of cancers. This opens up possibilities for applications in various fields where antibody therapies are increasingly important. Several programs aim to develop antibody-based therapies against cancer in both clinical and pre-clinical stages, with a few already being approved and included in the treatments.[25].

As computational models of protein interactions and evolutionary processes continue to improve, the principles of opponent shaping in therapy design are likely to become increasingly relevant. Our work presents a “proof-of-concept” for developing antiviral therapies that not only combat current viral strains but actively shape the evolutionary landscape to our advantage. While significant challenges remain in translating these computational insights into pre-clinical development, our study demonstrates the potential of opponent shaping in designing effective therapies against evolving pathogens.

4 Methods

Our method applies the principles of opponent shaping [9] to a two-player game between the antibody shaper player or *agent* and a naive virus agent (Figure 1a). In this section, we describe our method in detail.

Firstly, we introduce the **virus-antibody game** in Section 4.1 where we define the action spaces and the payoffs of both players. Secondly, we define the **simulated viral evolution** process in Section 4.2 where the virus evolves to escape in response to the *current antibody*. Finally, in Section 4.3 we define the **antibody optimisation** process where we optimise the antibody shapers in a way that *accounts for future viral mutations* and *learns to influence viral evolution away from escape*.

Importantly, implementing the virus-antibody game requires a way to determine the strength of binding of any antibody-antigen pair, or as we refer to it later in this section a **binding function**. In section 4.4 we describe our binding function choice, based on the Absolut! framework [12].

4.1 Virus-Antibody Game

We frame our setting as a *two-player zero-sum game* between the virus (v) and the antibody (a), meaning that the gain of one player is equal to the loss of the other player, and vice-versa. The action space of both agents is their respective amino acid sequence, and the payoffs are defined by Figure 1b and are the negative of each other - i.e. an improvement in antibody payoff results in a corresponding reduction of virus payoff of equal magnitude.

We define the set of 20 amino acids as \mathbb{A} . Let N_v be the sequence length of the viral antigen, and N_a be the sequence length of the antibody. So an action of the virus will be $v \in \mathbb{A}^{N_v}$, and an action of the antibody will be $a \in \mathbb{A}^{N_a}$. In our experiments, we use the dengue envelope antigen [21] as the virus, it is composed of 97 amino acids, thus $N_v = 97$. N_a is the sequence length of the antibody’s CDRH3 region which Absolut! [12] assumes to be 11 amino acids long, so $N_a = 11$.

Let $B : \mathbb{A}^{N_v} \times \mathbb{A}^{N_a} \rightarrow \mathbb{R}$ be our binding function, which measures the strength of the binding between the antibody and the viral antigen with increasing values corresponding to stronger binding³. Intuitively, the main payoff/reward for the antibody R_a should be $R_a = B(v, a)$ since it describes the capacity of the antibody to bind the virus, and the corresponding reward for the virus should be $R_v = -B(v, a)$.

However, such a choice would lead to a globally optimal strategy for the viral antigen to be completely inert, while the antibody would evolve to be maximally “sticky”. Such results fail to capture the biological realities: viruses must retain their ability to bind to host cell receptors to function, and antibodies need to maintain a certain level of specificity to avoid off-target binding [26]. Altogether, to create a biologically relevant and realistic game setting, we must incorporate additional incentives into our objectives for both the antibody and the virus.

To achieve this, we introduce a **binding target** t_v^+ for the virus. In the real world, this corresponds to the cell surface receptor protein, that is used by the virus to enter the host cell. To mimic the requirement that antibodies should not become self-reactive to the host proteins and trigger auto-immunity we also introduce a dummy **anti-target** t_a^- representing a protein for the self, that is a target to which the antibody should not bind.

This establishes the complete antibody payoff that can be used in the framework of our two-player virus-antibody game:

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

With $R_v(v, a) = -R_a(v, a)$. In this framework, the antibody gets a higher payoff not only if it binds the virus with higher affinity, but also if the virus fails to bind its receptor or if the antibody has a lower risk of autoimmunity.

We can now define the simulated viral escape via evolution. In the following section, we leverage the viral payoff $R_v(v, a)$ as a fitness function to model how viruses might evolve to evade antibody binding. This simulates viral trajectories that viruses could take in response to therapeutic pressure.

4.2 Simulated Viral Escape via Evolution

We model the escape of the virus \hat{v} by simulating the viral evolution from a fixed antibody a over a horizon H number of steps. The *simulated viral escape via evolution* (see Figure 1a) is defined as follows. Given a starting virus \hat{v} , the fixed antibody a induces a distribution $\text{Ev}(\hat{v}, a)$ over sequences of viruses $\hat{v}^1, \hat{v}^2, \hat{v}^3, \dots$. We write that $\hat{v} = \hat{v}^0, \hat{v}^1, \hat{v}^2, \dots, \hat{v}^H$, where $\hat{v}^0 = \hat{v}$ and H is the horizon length. Then a viral escape trajectory is given by $\hat{v} \sim \text{Ev}(\hat{v}, a)$. Below we define the process of generating the escape trajectories inductively.

At generation i we have a population of P viruses $v_1^i, v_2^i \dots v_P^i$. In our experiments $P = 15$. When $i = 0$, the population is the starting virus \hat{v} duplicated P times. 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 based on the fitness values, in particular:

$$\mathbb{P}(\hat{v}^i = v_k^i) \propto \exp(\beta R_v(v_k^i, a))$$

With duplicates in the population being considered distinct, 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 $\beta \rightarrow \infty$ reflecting deterministic max-fitness selection.

To generate the next population, \hat{v}^i is then duplicated and randomly mutated P times such that on average there is one amino acid mutation in the viral sequence, so:

$$v_k^{i+1} = \hat{v}^i \oplus \text{Mutation}$$

After horizon H generations, a full escape trajectory $\hat{v} = \hat{v}^0, \hat{v}^1, \hat{v}^2, \dots, \hat{v}^H$ is generated and the simulated viral escape process ends.

³This is opposite to binding **energies**, which are smaller for stronger binding.

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} , antibody fitness is:

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

This is a precise definition of the antibody fitness, Equation 1 in Section 2.1.3 uses a simplified notation. 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 antibody shapers, we employ Monte Carlo simulations to estimate the antibody fitness combined with an evolutionary algorithm. We refer to this process as the *antibody optimisation loop*, see Figure 1a. 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 $\hat{a} = \hat{a}^0, \hat{a}^1, \hat{a}^2, \dots, \hat{a}^N$, where N is the number of antibody optimisation steps (or 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 antibody optimisation step $i > 0$, we have a population of P_a antibodies $a_1^i, a_2^i \dots a_{P_a}^i$. In 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. We then select the best-performing antibody, \hat{a}^i :

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

When $i = 0$, $\hat{a}^0 = \hat{a}$ is the starting random antibody. We generate the next population of antibodies $a_1^{i+1}, a_2^{i+1} \dots a_{P_a}^{i+1}$ by taking both the antibody \hat{a}^i and $P_a - 1$ copies of it, each copy with a single random mutation to the amino acid sequence of \hat{a}^i .

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

4.4 Binding Function

The methodology described so far is independent of the choice of the binding function $B : \mathbb{A}^{N_v} \times \mathbb{A}^{N_a} \rightarrow \mathbb{R}$. In our work, we rely on the Absolut! framework [12] to implement the binding function and 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 [12].

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 residues⁴ 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* $M : \mathbb{A} \times \mathbb{A} \rightarrow \mathbb{R}$ which is determined experimentally by [14].

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

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

⁴A single amino acid position on a protein

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

$$B(v, a) = -E(v, a; S, M) = -\min_{\phi \in S} \hat{E}(v, a; \phi, M) \quad (2)$$

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 [27]. 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, and again 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 [21]. 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 1000^5 times faster for a negligible performance 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 does reveal 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 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 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. The reason is 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 superior, effective antibodies.

References

- [1] Yiska Weisblum, Fabian Schmidt, Fengwen Zhang, Justin DaSilva, Daniel Poston, Julio CC Lorenzi, Frauke Muecksch, Magdalena Rutkowska, Hans-Heinrich Hoffmann, Eleftherios Michailidis, Christian Gaebler, Marianna Agudelo, Alice Cho, Zijun Wang, Anna Gazumyan, Melissa Cipolla, Larry Luchsinger, Christopher D Hillyer, Marina Caskey, Davide F Robbiani, Charles M Rice, Michel C Nussenzweig, Theodora Hatzioannou, and Paul D Bieniasz. Escape from neutralizing antibodies by SARS-CoV-2 spike protein variants. *eLife*, 9:e61312, October 2020. Publisher: eLife Sciences Publications, Ltd.
- [2] Michael B. Doud, Juhye M. Lee, and Jesse D. Bloom. How single mutations affect viral escape from broad and narrow antibodies to H1 influenza hemagglutinin. *Nature Communications*, 9(1):1386, April 2018. Publisher: Nature Publishing Group.

⁵In practice, the difference is closer to 10,000

- [3] Juhye M Lee, Rachel Eguia, Seth J Zost, Saket Choudhary, Patrick C Wilson, Trevor Bedford, Terry Stevens-Ayers, Michael Boeckh, Aeron C Hurt, Seema S Lakdawala, Scott E Hensley, and Jesse D Bloom. Mapping person-to-person variation in viral mutations that escape polyclonal serum targeting influenza hemagglutinin. *eLife*, 8:e49324, August 2019. Publisher: eLife Sciences Publications, Ltd.
- [4] Adam S. Dingens, Dana Arenz, Haidyn Weight, Julie Overbaugh, and Jesse D. Bloom. An Antigenic Atlas of HIV-1 Escape from Broadly Neutralizing Antibodies Distinguishes Functional and Structural Epitopes. *Immunity*, 50(2):520–532.e3, February 2019. Publisher: Elsevier.
- [5] Allison J. Greaney, Tyler N. Starr, Pavlo Gilchuk, Seth J. Zost, Elad Binshtein, Andrea N. Loes, Sarah K. Hilton, John Huddleston, Rachel Eguia, Katharine H. D. Crawford, Adam S. Dingens, Rachel S. Nargi, Rachel E. Sutton, Naveenchandra Suryadevara, Paul W. Rothlauf, Zhuoming Liu, Sean P. J. Whelan, Robert H. Carnahan, James E. Crowe, and Jesse D. Bloom. 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. Publisher: Elsevier.
- [6] Alessandro M. Carabelli, Thomas P. Peacock, Lucy G. Thorne, William T. Harvey, Joseph Hughes, Thushan I. de Silva, Sharon J. Peacock, Wendy S. Barclay, Thushan I. de Silva, Greg J. Towers, and David L. Robertson. SARS-CoV-2 variant biology: immune escape, transmission and fitness. *Nature Reviews Microbiology*, 21(3):162–177, March 2023. Publisher: Nature Publishing Group.
- [7] Jie Hu, Pai Peng, Kai Wang, Liang Fang, Fei-yang Luo, Ai-shun Jin, Bei-zhong Liu, Ni Tang, and Ai-long Huang. Emerging SARS-CoV-2 variants reduce neutralization sensitivity to convalescent sera and monoclonal antibodies. *Cellular and Molecular Immunology*, 18(4):1061–1063, April 2021.
- [8] Shabir A. Madhi, Vicky Baillie, Clare L. Cutland, Merryn Voysey, Anthonet L. Koen, Lee Fairlie, Sherman D. Padayachee, Keertan Dheda, Shaun L. Barnabas, Qasim E. Bhorat, Carmen Briner, Gaurav Kwatra, Khatija Ahmed, Parvinder Aley, Sutika Bhikha, Jinal N. Bhiman, As’ad E. Bhorat, Jeanine du Plessis, Aliasgar Esmail, Marisa Groenewald, Elizea Horne, Shi-Hsia Hwa, Aylin Jose, Teresa Lambe, Matt Laubscher, Mookho Malahleha, Masebole Masenya, Mduduzi Masilela, Shakeel McKenzie, Kgaogelo Molapo, Andrew Moultrie, Suzette Oelofse, Faezah Patel, Sureshnee Pillay, Sarah Rhead, Hylton Rodel, Lindie Rossouw, Carol Taoushanis, Houriiyah Tegally, Asha Thombrayil, Samuel van Eck, Constantinos K. Wibmer, Nicholas M. Durham, Elizabeth J. Kelly, Tonya L. Villafana, Sarah Gilbert, Andrew J. Pollard, Tulio de Oliveira, Penny L. Moore, Alex Sigal, Alane Izu, 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.
- [9] Jakob Foerster, Richard Y. Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. Learning with Opponent-Learning Awareness. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS ’18*, pages 122–130, Richland, SC, July 2018. International Foundation for Autonomous Agents and Multiagent Systems.
- [10] Chris Lu, Timon Willi, Christian Schroeder de Witt, and Jakob Foerster. Model-Free Opponent Shaping, November 2022. arXiv:2205.01447 [cs].
- [11] Michaela Lucas, Urs Karrer, Andrew Lucas, and Paul Klenerman. Viral escape mechanisms – escapology taught by viruses. *International Journal of Experimental Pathology*, 82(5):269–286, 2001.
- [12] Philippe A. Robert, R. Akbar, R. Frank, Milena Pavlović, Michael Widrich, I. Snapkov, Andrei Slabodkin, Maria Chernigovskaya, Lonneke Scheffer, Eva Smorodina, Puneet Rawat, Brij Bhushan Mehta, Mai Ha Vu, Ingvild Frøberg Mathisen, Aurel Prosz, K. Abram, Alexandru Olar, Enkelejda Miho, Dag Trygve Tryslew Haug, F. Lund-Johansen, S. Hochreiter, Ingrid Hobæk Haff, G. Klambauer, G. K. Sandve, and V. Greiff. Unconstrained generation of synthetic antibody–antigen structures to guide machine learning methodology for antibody specificity prediction. *Nature Computational Science*, 2022.
- [13] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, and Skye Wanderman-Milne. JAX: composable transformations of Python+ NumPy programs. 2018.
- [14] Sanzo Miyazawa and Robert L. Jernigan. An empirical energy potential with a reference state for protein fold and sequence recognition. *Proteins: Structure, Function, and Bioinformatics*, 36(3):357–369, 1999.
- [15] L. VanDyk and K. Meek. Assembly of IgH CDR3: mechanism, regulation, and influence on antibody diversity. *International Reviews of Immunology*, 8(2-3):123–133, 1992.
- [16] Derek M. Mason, Simon Friedensohn, Cédric R. Weber, Christian Jordi, Bastian Wagner, Simon M. Meng, Roy A. Ehling, Lucia Bonati, Jan Dahinden, Pablo Gainza, Bruno E. Correia, and Sai T. Reddy. Optimization of therapeutic antibodies by predicting antigen specificity from antibody sequence via deep learning. *Nature Biomedical Engineering*, 5(6):600–612, June 2021.

- [17] Yoong Wearn Lim, Adam S. Adler, and David S. Johnson. Predicting antibody binders and generating synthetic antibodies using deep learning. *mAbs*, 14(1):2069075, 2022.
- [18] Jeffrey A. Ruffolo, Lee-Shin Chu, Sai Pooja Mahajan, and Jeffrey J. Gray. Fast, accurate antibody structure prediction from deep learning on massive set of natural antibodies. *Nature Communications*, 14(1):2389, April 2023.
- [19] 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.
- [20] Scott A. Hollingsworth and Ron O. Dror. Molecular dynamics simulation for all. *Neuron*, 99(6):1129–1143, September 2018.
- [21] Shee-Mei Lok, Victor Kostyuchenko, Grant E. Nybakken, Heather A. Holdaway, Anthony J. Battisti, Soila Sukupolvi-Petty, Dagmar Sedlak, Daved H. Fremont, Paul R. Chipman, John T. Roehrig, Michael S. Diamond, Richard J. Kuhn, and Michael G. Rossmann. Binding of a neutralizing antibody to dengue virus alters the arrangement of surface glycoproteins. *Nature Structural & Molecular Biology*, 15(3):312–317, March 2008.
- [22] Feng Gao, Mattia Bonsignori, Hua-Xin Liao, Amit Kumar, Shi-Mao Xia, Xiaozhi Lu, Fangping Cai, Kwan-Ki Hwang, Hongshuo Song, Tongqing Zhou, Rebecca M. Lynch, S. Munir Alam, M. Anthony Moody, Guido Ferrari, Mark Berrong, Garnett Kelsoe, George M. Shaw, Beatrice H. Hahn, David C. Montefiori, Gift Kamanga, Myron S. Cohen, Peter Hraber, Peter D. Kwong, Bette T. Korber, John R. Mascola, Thomas B. Kepler, and Barton F. Haynes. Cooperation of B Cell Lineages in Induction of HIV-1-Broadly Neutralizing Antibodies. *Cell*, 158(3):481–491, July 2014. Publisher: Elsevier.
- [23] Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J. Ballard, Joshua Bambrick, Sebastian W. Bodenstein, David A. Evans, Chia-Chun Hung, Michael O’Neill, David Reiman, Kathryn Tunyasuvunakool, Zachary Wu, Akvilė Žemgulytė, Eirini Arvaniti, Charles Beattie, Ottavia Bertolli, Alex Bridgland, Alexey Cherepanov, Miles Congreve, Alexander I. Cowen-Rivers, Andrew Cowie, Michael Figurnov, Fabian B. Fuchs, Hannah Gladman, Rishub Jain, Yousuf A. Khan, Caroline M. R. Low, Kuba Perlin, Anna Potapenko, Pascal Savy, Sukhdeep Singh, Adrian Stecula, Ashok Thillaisundaram, Catherine Tong, Sergei Yakneen, Ellen D. Zhong, Michal Zielinski, Augustin Žídek, Victor Bapst, Pushmeet Kohli, Max Jaderberg, Demis Hassabis, and John M. Jumper. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, 630(8016):493–500, June 2024.
- [24] Nicole N. Thadani, Sarah Gurev, Pascal Notin, Noor Youssef, Nathan J. Rollins, Daniel Ritter, Chris Sander, Yarin Gal, and Debora S. Marks. Learning from pre-pandemic data to forecast viral escape. *Nature*, 622(7984):818–825, October 2023. Publisher: Nature Publishing Group.
- [25] David Zahavi and Louis Weiner. Monoclonal Antibodies in Cancer Therapy. *Antibodies*, 9(3):34, July 2020.
- [26] Lilia A. Rabia, Alec A. Desai, Harkamal S. Jhaji, and Peter M. Tessier. Understanding and overcoming trade-offs between antibody affinity, specificity, stability and solubility. *Biochemical engineering journal*, 137:365–374, September 2018.
- [27] wwPDB consortium. Protein Data Bank: the single global archive for 3D macromolecular structure data. *Nucleic Acids Research*, 47(D1):D520–D528, October 2018.

A Viral Mutation Distribution




Figure 6: Distribution of selected viral mutations when evolved against different antibody horizons. Due to more neutral mutations, the distribution is closer to uniform. Longer horizon antibodies cause fewer mutations on average.

B Influence of Antibody Shapers on Binding Poses




Figure 7: 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.