SHAPED, NOT STIRRED: CO-EVOLVING COMBINATION THERAPIES TO PREVENT VIRAL ESCAPE

Anonymous authors
Paper under double-blind review

ABSTRACT

Combination therapies are the gold standard for treating rapidly evolving viruses like HIV or Influenza, where single-drug approaches quickly lose efficacy due to viral adaptation. However, current therapy design is fundamentally *myopic*, and optimises for immediate efficacy against present viral strains without anticipating how therapeutic pressure will influence future viral evolution. We introduce **Co-ADIOS** (**Co**operative **A**ntibody **D**esign vi**a** **O**pponent **S**haping), a framework that designs antibody combination therapies, or *cocktails*, that steer viral evolution toward weaker variants. We formulate the interaction between antibody cocktails and the virus as a *Stackelberg Evolutionary Game*, where antibodies act as *leaders* that shape the fitness landscape to steer the adaptive trajectory of the virus (the *evolving follower*). In the game, viral evolutionary dynamics are governed by a *biophysically grounded lattice simulator* that accounts for binding energies and evolutionary plausibility. We evaluate Co-ADIOS on the Influenza H5 virus in two clinically relevant scenarios: (1) designing de novo antibody cocktails, and (2) co-designing a partner antibody to complement an established therapeutic. We demonstrate that Co-ADIOS antibodies outperform myopic baselines by using cooperative strategies which achieve stronger shaping over viral evolution, at a modest cost in efficacy against the current strain. Finally, our analysis shows that the designed cocktails leave the virus far fewer single mutations worth taking, and that their two antibodies bind a shared site while covering different escape mutations, offering physically explainable insights for the design of effective real-world viral therapies.

1 INTRODUCTION

Viruses pose a fundamental challenge for therapeutic design: they continuously evolve in response to environmental pressure. Traditional approaches are *myopic*, i.e. they design therapies to target only the current viral strains, but this leaves treatments vulnerable as viruses inevitably mutate and become therapy-resistant. The COVID-19 pandemic exemplified this challenge, where vaccines lost efficacy against emerging variants like B.1.351 (Madhi et al., 2021). Instead, the ADIOS framework proposed a different approach: designing antibodies that actively shape viral evolution toward less dangerous variants (Towers et al., 2025). ADIOS simulates viral adaptation and selects antibody *shapers* that retain long-term efficacy by steering evolution towards weaker variants.

However, ADIOS models a single-drug selective pressure, whereas pathogens in clinical settings are rarely subject to only one. The gold standard for treating evolving diseases are **combination therapies**, which have been shown to be more effective in treating diseases like HIV (Hammer et al., 1997), Hepatitis (Afdhal et al., 2014), and COVID-19 (Weinreich et al., 2021), inherently creating multi-pressure environments as multiple drugs target the same pathogen (Shyr et al., 2021). Yet, the design of these *cocktails* remains fundamentally myopic: individual drugs are optimised in isolation and then combined, creating an uncoordinated multi-pressure environment. This strategy applies multi-pressure but not *shaped and coordinated* evolutionary pressure, an opportunity neither traditional cocktail design nor single-antibody ADIOS addresses.

This motivates coordinated multi-antibody ADIOS, or **Co-ADIOS**, in which we design antibody cocktails to shape viral evolution together. We formulate the interaction as a *Stackelberg evolutionary game* (Stein et al., 2023): the cocktail leads by selecting its action first, and the virus follows by

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

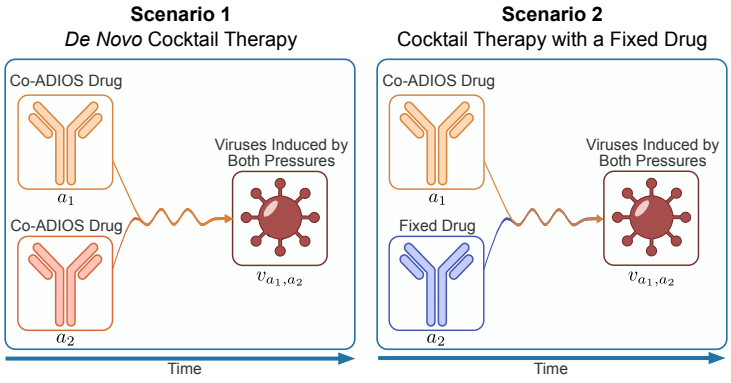

Figure 1: **Two clinically relevant antibody combination therapy design scenarios.** Co-ADIOS optimises combination therapies for performance across entire viral evolutionary trajectories rather than just current variants. *Scenario 1:* Co-ADIOS jointly optimises both antibodies a_1 and a_2 to cooperatively shape viral evolution. *Scenario 2:* Co-ADIOS designs antibody a_1 to work synergistically with a fixed therapeutic antibody a_2.

evolving against the combined pressure the cocktail exerts, mutating under BLOSUM62 substitution propensities (Henikoff & Henikoff, 1992) so that the escape routes it explores are ones seen in nature. Co-ADIOS optimises the cocktail over the viral trajectories it induces using a cooperative co-evolutionary algorithm.

We demonstrate the effectiveness of Co-ADIOS on two clinically relevant scenarios shown in Figure 1. **Scenario 1** considers the design of a fully *de novo* **antibody cocktail**, where both antibodies are co-designed for cooperative shaping. **Scenario 2** considers the design of a partner antibody for an existing approved therapeutic, which offers the benefits of combination therapy while minimising clinical development costs and regulatory hurdles.

Taken together, our contributions are:

- **Cocktail-vs-virus as a Stackelberg evolutionary game**: A formulation in which the cocktail leads, and the virus evolves in response, with payoffs that score both antibodies jointly through the fractional-activity model (Einav & Bloom, 2020) rather than one antibody at a time.
- **Co-ADIOS**: An algorithm that solves this game by using an inner viral-escape simulation inside an outer cooperative co-evolutionary loop with one sub-population per antibody, extending opponent shaping to coordinated cocktail design.
- **Empirical evaluation on Influenza H5**: Across two clinically-relevant scenarios in the lattice-discretised Absolut! simulator (Robert et al., 2022), Co-ADIOS cocktails outperform myopic baselines and independent ADIOS shapers, achieving stronger long-term shaping, which the myopic baselines trade for efficacy against the current strain.
- **Emergent coordination strategies**: Explainability analysis shows that Co-ADIOS cocktails leave the virus fewer and weaker single-mutation escape routes, with the two antibodies sharing a binding site and dividing the escape mutations between them.

Readers unfamiliar with biological terminology may find Section A a useful reference throughout.

2 Related Work

Computational Antibody Design. Recent generative models design antibodies that bind a specified antigen target (Joubbi et al., 2024). Diffusion-based methods (Luo et al., 2022) and equivariant graph methods (Kong et al., 2023) generate complementarity-determining regions or full-atom antibody structures conditioned on the antigen. Protein language models mature antibodies against current viral variants (Hie et al., 2024), and structure-informed extensions further improve neutralisation against escape variants of concern (Shanker et al., 2024). Closest to our setting, sequence-level optimisers

search subsequences of the antibody heavy chain against the same Absolut! scoring function we use: AntBO applies combinatorial Bayesian optimisation to this problem (Khan et al., 2023), and Vogt et al. (2023) train reinforcement-learning policies for it. Yet all of these methods remain myopic: they optimise against a fixed antigen and offer no mechanism to anticipate the variants that will emerge once the therapy is deployed.

Antibody Design Against Multiple Variants. One antibody can be designed to neutralise many viral variants at once. Multistate protein design scores a single sequence across several target structures and combines the state-specific scores into one objective (Leaver-Fay et al., 2011). Sevy et al. (2018) redesign a broadly neutralising antibody over 180 HIV strains with integer linear programming, and Shan et al. (2022) optimise a human antibody for breadth across SARS-CoV-2 variants. The variant set is chosen before the optimisation starts, and so it does not account for the therapeutic pressure induced on the viruses.

Modelling Viral Evolution and Escape. Predictive evolutionary models forecast which viral variants will emerge: antigenic-phenotype models trace influenza drift through the viral evolutionary tree (Neher et al., 2016), beth-1 (Lou et al., 2024) forecasts influenza evolution by modelling site-wise mutation fitness, and CoVFit (Ito et al., 2025) scores SARS-CoV-2 variant fitness directly from spike sequence. EVEscape (Thadani et al., 2023) generalises to viruses without prior surveillance, using only pre-pandemic sequences. At the population level, phylodynamic methods couple this tree to epidemic transmission (Kühnert et al., 2011), and PyR_0 (Obermeyer et al., 2022) infers per-mutation fitness from millions of genomes. None of these models account directly for the effect of therapeutic pressure on viral evolution; using them to inform therapy design risks distribution shift once the deployed therapy reshapes the fitness landscape.

Designing Combination Therapies. Given the members of a combination, synergy-prediction frameworks rank pairs by their immediate combined potency (Chen et al., 2025), and pharmacological models optimise the dose of each member over the course of treatment (Brill et al., 2018). Choosing the members themselves is usually structural: cocktail design pairs antibodies that bind different sites, or *epitopes*, on the virus (Su et al., 2021), on the assumption that a mutation disrupting one binding site leaves the other antibody effective (Baum et al., 2020). Different binding sites need not imply different escape routes: Starr et al. (2021) report a single mutation that escapes both antibodies of an approved two-antibody cocktail. However, the choice of cocktail members is not optimised against the evolution the combination will induce.

3 Background

Opponent Shaping. In multi-agent reinforcement learning (MARL), opponent shaping methods equip designated shaper agents with an update rule that anticipates how their policies influence the learning of other (potentially myopic) agents. Learning with Opponent-Learning Awareness (LOLA) (Foerster et al., 2018) adds a higher-order gradient term that differentiates through the opponent's anticipated update, with consistent variants such as COLA (Willi et al., 2022) addressing fixed-point inconsistencies in the original formulation. Model-Free Opponent Shaping (M-FOS) (Lu et al., 2022) instead frames the interaction as a meta-game in which a meta-policy selects the inner policy played each episode, learning to shape via standard model-free reinforcement learning. ADIOS (Towers et al., 2025) applies this paradigm to antibody design in the Absolut! simulator, optimising single antibody shapers that steer viral evolution toward weaker variants. We extend ADIOS from a single shaper to a coordinated antibody cocktail.

Stackelberg Evolutionary Games. Stackelberg evolutionary game (SEG) theory (Stein et al., 2023) studies the optimal strategy for a rational leader interacting with a follower population whose composition changes via evolutionary dynamics, combining classical Stackelberg leadership with evolutionary game theory. Static aggressive policies are typically suboptimal in this setting: the Stackelberg-equilibrium leader anticipates and steers the eco-evolutionary trajectory of the follower, a principle that adaptive cancer therapy has applied to delay or avoid resistance (Staňková et al., 2019). We instantiate a SEG (Figure 2) between an antibody cocktail (the leader) and a viral population (the evolving follower), with the cocktail's fitness defined over the induced viral evolutionary trajectory.

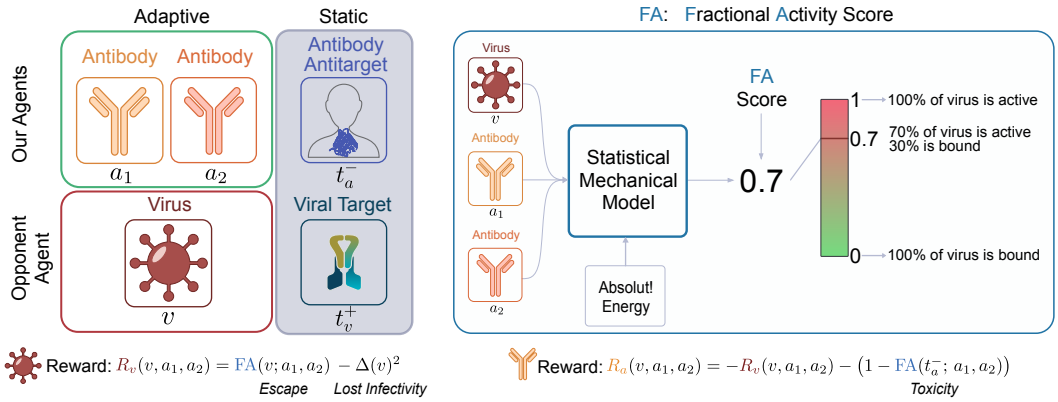

Figure 2: **The Stackelberg evolutionary game between an antibody cocktail (leader) and an evolving virus (follower).** Antibody cocktail (a_1, a_2) acts on virus v subject to two static constraints: a viral target t_v^+ that the virus must bind to remain infectious, and an antibody antitarget t_a^- that the antibodies must avoid. The cocktail's joint binding to the virus is scored by a statistical-mechanical fractional-activity model (FA) over Absolut! binding energies: FA $= 1$ means the viral population is fully active, FA $= 0$ means it is fully bound. The bottom row gives the per-agent rewards, where $\Delta(v) = \max\big(\mathrm{FA}(v; t_v^+) - \mathrm{FA}(v_0; t_v^+),\, 0\big)$ is the target binding the virus has given up relative to the wildtype v_0.

Binding and Fractional Activity. The payoffs of the game are computed from binding energies. We score the binding of an antibody a to a virus v with Absolut! (Robert et al., 2022), which discretises the viral antigen onto a 3D lattice and exhaustively enumerates the poses the antibody can adopt against it. Each pose is scored by summing Miyazawa–Jernigan contact energies over neighbouring residue pairs, and the binding energy $E(a, v)$ is that of the lowest-energy pose (see Section B for details). When multiple antibodies target the same virus simultaneously, their combined effect is not additive. Einav & Bloom (2020) introduced a statistical-mechanical model of *fractional activity* (FA) that maps binding energies to the fraction of the virus left unbound, and so still active: FA $= 1$ when none of it is bound and FA $= 0$ when all of it is (see FA score in Figure 2). We write $\mathrm{FA}(v; a)$ for a single antibody and $\mathrm{FA}(v; a_1, a_2)$ for a pair, and we apply the same model to the viral target, where $\mathrm{FA}(v; t_v^+)$ is the fraction of virus not bound to its receptor. Whereas ADIOS (Towers et al., 2025) relies on single-antibody binding scores, FA reduces a cocktail to the single value we optimise (see Section C for details).

Cooperative Co-Evolution. Designing a cocktail means optimising two antibody sequences to maximise their joint payoff. Cooperative co-evolutionary algorithms (CCEAs), first proposed by Potter & De Jong (1994) and surveyed by Ma et al. (2019), decompose the optimisation problem into multiple sub-populations searched in parallel. Individuals in each sub-population are evaluated collaboratively by pairing them with a champion from the other sub-population, so that fitness rewards compatibility over isolated quality and the sub-populations specialise into complementary roles. We adopt this structure for the Co-ADIOS outer loop, with one sub-population per antibody in the cocktail, jointly searching the space of two-antibody combinations under a shared shaping objective.

4 Method

Co-ADIOS extends the ADIOS opponent-shaping framework (Towers et al., 2025) from a single *shaper antibody* to a coordinated antibody *cocktail* acting on an evolving viral population. We first define the cocktail-vs-virus interaction as a **Stackelberg evolutionary game** (Section 4.1), with the cocktail as leader and the virus as evolving follower. An **inner-loop viral escape simulator** then models the viral evolution (Section 4.2), and a **cooperative co-evolutionary outer loop** optimises the cocktail against the whole simulated viral trajectory (Section 4.3).

4.1 COCKTAIL-VS-VIRUS AS A STACKELBERG EVOLUTIONARY GAME

We frame the cocktail-vs-virus interaction as a Stackelberg evolutionary game (Figure 2) with tuple

$$\mathcal{G} = \langle \mathcal{A}_a, \mathcal{A}_v, T, R_v, R_a, H, t_v^+, t_a^- \rangle.$$

Antibodies and viruses are both sequences over a 20-letter amino acid alphabet $\mathbb{A}$. We optimise an antibody only at its CDRH3 loop[1] (length N_a). The cocktail is a pair of antibodies, giving action spaces $\mathcal{A}_a = \mathbb{A}^{N_a} \times \mathbb{A}^{N_a}$ and $\mathcal{A}_v = \mathbb{A}^{N_v}$, where N_v is the length of the antigen sequence. Given the leader's action $(a_1, a_2) \in \mathcal{A}_a$, the follower's state evolves over H steps as a Markov chain $v_{t+1} \sim T(\cdot \mid v_t, a_1, a_2)$ from initial wildtype virus $v_0 \in \mathcal{A}_v$. The leader maximises an escape-averaged return $F_H(a_1, a_2)$ defined in Section 4.3; we specify the transition kernel T in Section 4.2.

The game also features two static targets, both fixed amino-acid sequences over $\mathbb{A}$. The viral target t_v^+ is a structure the virus must keep binding to remain *infectious* (e.g., a host-cell receptor site). The antibody antitarget t_a^- is a structure the cocktail must not bind (e.g., a human protein), since binding it would cause clinical *toxicity*. Specific instantiations are introduced in Section 5.

Viral reward R_v combines escape from the cocktail (high $\text{FA}(v; a_1, a_2)$) with a penalty for the target binding the virus gives up as it mutates, written

$$\Delta(v) = \max\big(\text{FA}(v; t_v^+) - \text{FA}(v_0; t_v^+),\, 0\big),$$

the loss in viral-target binding relative to the wildtype v_0, weighted by $\lambda_{\text{vt}} > 0$. The floor at zero makes the penalty one-sided: a virus that binds its target better than the wildtype doesn't gain extra payoff because the wildtype would already bind strongly to its target. Cocktail reward R_a directly opposes R_v and adds a toxicity penalty for binding the antitarget, weighted by $\lambda_{\text{tox}} > 0$.

$$R_v(v, a_1, a_2) = \text{FA}(v; a_1, a_2) - \lambda_{\text{vt}}\, \Delta(v)^2$$
$$R_a(v, a_1, a_2) = -R_v(v, a_1, a_2) - \lambda_{\text{tox}}\big(1 - \text{FA}(t_a^-; a_1, a_2)\big)$$

Importantly, the cocktail does not maximise R_a per-step; the outer loop maximises its escape-averaged version F_H (Section 4.3).

4.2 INNER LOOP: VIRAL ESCAPE UNDER COCKTAIL PRESSURE

During the inner loop, the virus v evolves for horizon H steps, while the antibody cocktail remains unchanged. Each inner-loop step (Figure 3, right) performs two operations: mutate the current virus to generate a cohort of offspring, then probabilistically select one as the next generation seed. Given the current virus v_t, we sample P_v mutant offspring,

$$v^{(k)} \sim \text{Mutate}(v_t), \quad k = 1, \ldots, P_v,$$

where Mutate substitutes a position of the viral sequence with a different amino acid with probability $\mu_v/|N_v|$, so μ_v is the expected number of substitutions per sequence. To draw the replacement, we sample the amino acid distribution proportionally to the BLOSUM62 score (Henikoff & Henikoff, 1992) which favours biologically plausible substitutions. In our setting, we constrain the plausible viral mutation positions to a subset $\mathcal{S}_v \subseteq \{1, \ldots, N_v\}$ where escape is more likely to occur. The remaining $N_v - |\mathcal{S}_v|$ positions are held fixed. The next-generation virus is then sampled by softmax over the offspring cohort $\mathcal{O} = \{v^{(1)}, \ldots, v^{(P_v)}\}$:

$$v_{t+1} \sim \text{softmax}_{v' \in \mathcal{O}}\big(R_v(v', a_1, a_2)/\beta\big),$$

where temperature β controls stochasticity. Composing these two operations realises the stochastic transition kernel $T(\cdot \mid v_t, a_1, a_2)$ from Section 4.1.

4.3 OUTER LOOP: COOPERATIVE CO-EVOLUTIONARY COCKTAIL DESIGN

In the outer loop we optimise the cocktail (a_1, a_2) over M generations, each of which runs the full inner loop for every candidate. While the per-step reward R_a scores the cocktail's instantaneous

[1] CDRH3 is the third complementarity-determining region or loop on the antibody that dominates binding specificity (VanDyk & Meek, 1992).

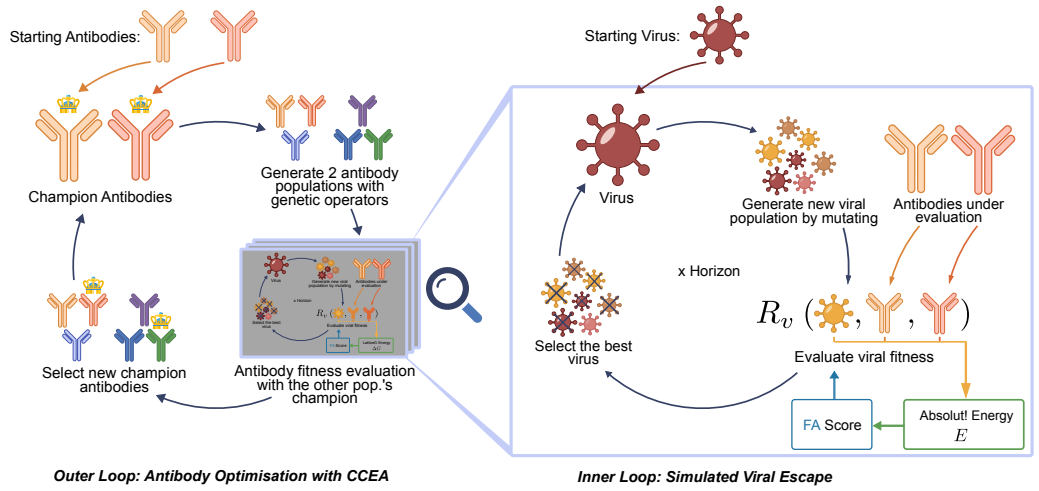

Figure 3: **Co-ADIOS nests an outer cooperative co-evolutionary loop over an inner viral-escape loop.** *Inner loop* (right): the current virus mutates into a cohort, which is then sampled by softmax on the cocktail-conditioned viral reward R_v; binding is scored by Absolut! (E) and aggregated into the fractional activity (FA). *Outer loop* (left): two antibody sub-populations evolve in parallel via CCEA; each candidate's fitness is the inner-loop return F_H computed by pairing it with the other sub-population's current champion.

neutralisation of a single virus, optimising for shaping requires evaluating across the whole induced trajectory of viral escape attempts. Let $\tau = (v_\tau^0, \ldots, v_\tau^H)$ denote such a trajectory, with $v_\tau^0 = v_0$ and $v_\tau^{t+1} \sim T(\cdot \mid v_\tau^t, a_1, a_2)$. The cocktail's *escape-averaged return* is

$$F_H(a_1, a_2) = \mathbb{E}_\tau \left[\frac{1}{H+1} \sum_{t=0}^{H} R_a(v_\tau^t, a_1, a_2) \right].$$

We approximate F_H by averaging over K Monte-Carlo rollouts of the inner loop. The horizon H counts the escape steps we simulate, so a trajectory spans $H+1$ scored viral states: at $H = 0$, F_H collapses to $R_a(v_0, a_1, a_2)$, the one-shot scoring against the wildtype we use as the myopic baseline, and at $H > 0$ the cocktail is rewarded for steering evolution toward weaker variants rather than for binding the wildtype alone.

To optimise F_H in $\mathcal{G}$, we use a CCEA (Section 3), shown in Figure 3. At each generation m, the algorithm maintains two sub-populations $\mathcal{P}_1^{(m)}, \mathcal{P}_2^{(m)} \subset \mathbb{A}^{N_a}$ of size P_a, plus a current champion $\hat{a}_k^{(m)}$ for each, $k \in \{1, 2\}$. Fitness is computed by champion-pairing: each ith candidate $a_{k,i}^{(m)}$ scores $F_H(a_{k,i}^{(m)}, \hat{a}_{-k}^{(m)})$ against the partner's champion.

We initialise $\mathcal{P}_1^{(0)}, \mathcal{P}_2^{(0)}$ uniformly from $\mathbb{A}^{N_a}$ and seed the champions by an arbitrary cross-pairing. A generation then runs one sub-round per sub-population. In each, only the active sub-population breeds, using standard genetic operators (tournament selection, uniform crossover at rate p_{xover}, BLOSUM62-weighted mutation at rate μ_a, and elitism (Eiben & Smith, 2015)), and its candidates are scored against the partner's current champion, which is frozen for that sub-round. The champions therefore update sequentially:

$$\hat{a}_1^{(m+1)} = \arg\max_i F_H\left(a_{1,i}^{(m+1)}, \hat{a}_2^{(m)}\right),$$
$$\hat{a}_2^{(m+1)} = \arg\max_j F_H\left(\hat{a}_1^{(m+1)}, a_{2,j}^{(m+1)}\right).$$

We run for M generations and report the final champion pair $(\hat{a}_1^{(M)}, \hat{a}_2^{(M)})$. By scoring every $a_{1,i}$ against the same $\hat{a}_2^{(m)}$, fitness reflects each candidate's individual contribution to joint cocktail performance rather than its partner's quality, preventing one sub-population from *hitchhiking* on the other's strength (Potter & De Jong, 1994). Champion-pairing therefore rewards an antibody for complementing its partner rather than for its individual strength.

5 EXPERIMENTAL SETUP

5.1 BINDING SIMULATOR

The game is played inside a binding simulator, the fixed map from a cocktail (a_1, a_2) and a virus v to the payoffs of Section 4.1. We instantiate it with Absolut! (Robert et al., 2022), which enumerates the *poses*, the discrete positions an antibody can take against the antigen, and returns the binding energy $E(a, v)$ of the best pose of antibody a against virus v. Each cocktail antibody is scored separately, yielding two energies, $E(a_1, v)$ and $E(a_2, v)$. The statistical-mechanical model of Einav & Bloom (2020) turns those two energies and the corresponding antibody doses into the single fractional activity $\mathrm{FA}(v; a_1, a_2)$, and scores the viral target and the antitarget the same way (see Section C for more details). We calibrate the model on the energy values the viral trajectories reach, so that a single viral substitution changes FA. We optimise cocktails at a coarse structural resolution of the simulator and re-score each champion at a finer one, which creates a small transfer gap between the training and evaluation environments; Section D gives both and the calibration.

5.2 INFLUENZA H5 CASE STUDY

Influenza is a natural candidate for shaping: its adapts between seasons, so a therapy designed for one strain loses efficacy against the next (Neher et al., 2016). In our experiments the virus v is influenza H5 haemagglutinin (PDB 5JW4 (Kallewaard et al., 2016)) and the antitarget t_a^- is a human protein (PDB 1ADQ (Corper et al., 1997)) that the antibody must avoid binding to. We keep only the poses on the head of the haemagglutinin, so that the lowest-energy pose falls closer to the receptor-binding sites where escape mutations are known to arise (Yen et al., 2009). Each antibody is eleven residues long ($N_a = 11$), and the virus may mutate the residues any retained pose touches: $|\mathcal{S}_v| = 42$ at the coarse resolution the search runs on, and 44 at the fine one it is evaluated on. It must also keep binding its target receptor t_v^+. We select CH65 as t_v^+ (Whittle et al., 2011), which is the closest thing to sugar (our desired target) that Absolut! can simulate.

5.3 SCENARIOS AND BASELINES

Our two main experiments are the scenarios of Figure 1. In Scenario 1 both antibodies are designed *de novo*, against a *single shaper*, which omits the second antibody; a *non-cooperative cocktail*, which concatenates the pair into one sequence and optimises with a single population instead of coevolving two; and *myopic cocktail* and *myopic single* baselines optimised at $H = 0$. Scenario 2 fixes a deployed antibody and Co-ADIOS designs a complement beside it. We use two baselines: the partner administered alone, and a *myopic complement*. We repeat this for seven deployed partners, chosen to span a wide range of potency against the antigen. Throughout, a cocktail splits a fixed total dose evenly, so a pair and a single antibody are compared at equal dose, and every design is optimised on 20 seeds and reported on 20 held-out ones. Beyond designing antibodies, we also score every pair in a library of existing therapeutics under the same objective, which asks whether a combination can be assembled from antibodies already in clinical development rather than designed; Section G reports that experiment.

6 RESULTS

We evaluate Co-ADIOS on Scenario 1 (Section 6.1) and Scenario 2 (Section 6.2), then examine what differentiates the Co-ADIOS cocktails from the baselines (Section 6.3). Unless stated otherwise, designs are compared by held-out F_{100}, paired over 20 seeds.

6.1 SCENARIO 1: DE NOVO COCKTAIL DESIGN

A designed cocktail outperforms a single designed antibody at equal total dose. Co-ADIOS improves the held-out return over a single shaper on 18 of 20 seeds (Figure 4d). However, optimising against the viral trajectory matters more than the second antibody. Shaping is worth close to four times the improvement from adding an antibody, and the held-out return rises monotonically with the shaping horizon length (Figure 4c). Without it the cocktail advantage disappears: at $H = 0$ a pair is indistinguishable from a single antibody. The trajectories demonstrate why (Figure 4a): myopic

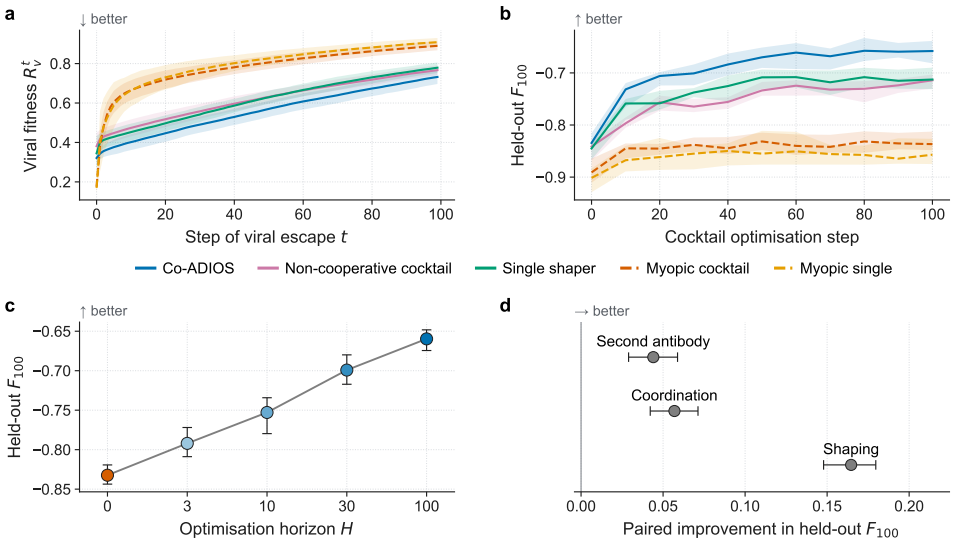

Figure 4: **Scenario 1: Designing a cocktail *de novo*. a** Viral fitness along held-out rollouts, averaged over 20 seeds, with a ± 1 s.d. band across seeds. **b** Each optimisation step's champion re-scored on held-out rollouts at horizon 100, median over seeds with the interquartile band. **c** Held-out F_{100} against the optimisation horizon, interquartile mean with 95% bootstrap intervals. **d** Paired improvement of Co-ADIOS over each ablation, mean with 95% confidence interval over the 20 shared seeds.

designs suppress the wildtype v_0 best of all, but the virus then escapes them fastest. Section E separates this into the terms of the objective, where the same pattern holds for neutralisation.

The second antibody is only beneficial when the two are designed together. A non-cooperative cocktail is no better than a single designed antibody, while Co-ADIOS beats it on 19 of 20 seeds (Figure 4d). Coordination is therefore worth more than the extra antibody itself, and the gap is visible throughout the learning curve of the cocktail (Figure 4b).

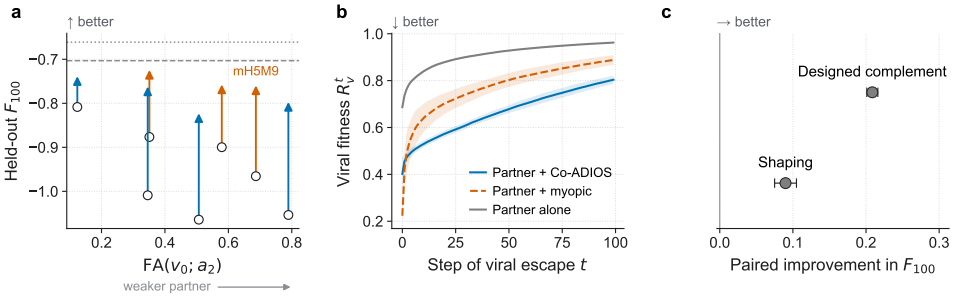

Figure 5: **Scenario 2: Designing a complement to an existing partner. a** Each arrow runs from a partner administered alone at full dose to the same partner paired with a designed complement, positioned by viral activity FA(v_0; a_2) against the partner a_2. Blue arrows are the four designed partners and orange the three published ones; the dotted and dashed grey lines mark the Scenario 1 Co-ADIOS cocktail and single shaper, for reference. **b** Viral fitness along held-out rollouts for the mH5M9 partner, averaged over 20 seeds with a ± 1 s.d. band. **c** Paired improvement at that partner, mean with 95% confidence interval over the 20 shared seeds.

6.2 Scenario 2: Cocktail with a Fixed Partner

A designed complement improves every partner on every seed, recovering between half and three quarters of the distance to a freely designed cocktail (Figure 5a). This improvement is largely independent of the partner: the final cocktail always reaches a similar performance regardless of the starting score of the fixed partner FA(v_0; a_2). mH5M9 (Zhu et al., 2013), the only partner which is experimentally verified to bind our antigen, behaves like designed partners of comparable potency.

Shaping remains beneficial here but is less significant than in Scenario 1. A myopic complement optimised against the wildtype alone recovers more than half of what a Co-ADIOS complement achieves(Figure 5c). The trajectories separate as in Scenario 1 (Figure 5b): the myopic complement is stronger against the wildtype and weaker after a few steps of viral escape.

6.3 Explainability

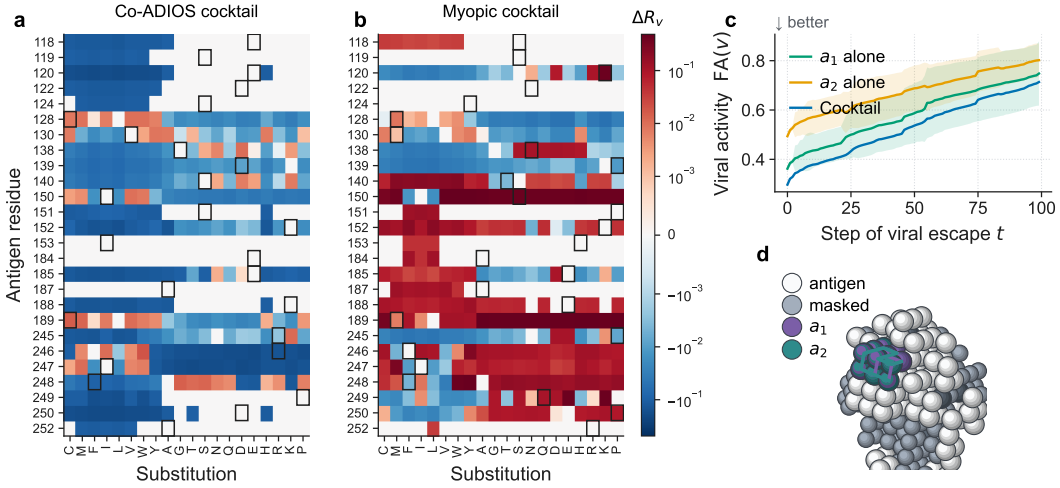

Figure 6: **Viral escape landscapes with and without Co-ADIOS. a,b** Change in viral reward for every single amino acid substitution of the viral sequence, against a Co-ADIOS cocktail and a myopic one, for one seed. Rows are mutable antigen residues, columns the substituted amino acid; red favours the virus, blue blocks viral escape, and black outlines show the wildtype sequence v_0. **c** Viral activity along held-out rollouts of that cocktail, for the pair and for each member, averaged over 20 rollouts with a ± 1 s.d. band. **d** The poses the cocktail's two members occupy, each with its lattice walk; the masked residues are grey.

Co-ADIOS blocks most of the escape routes still available against a myopic cocktail. Scoring every single substitution of the escape set against each design (Figure 6a,b), 56 of 880 raise viral reward under Co-ADIOS against 188 under the myopic cocktail on average, and the best move available is an order of magnitude smaller. The two antibodies bind the same site. Their poses always overlap, sharing at least five of eleven lattice positions and ten contacted residues, and in over half the replicates they are identical (Figure 6d), differing in which of their own positions face which residue. They divide the escape mutations instead: of the 122 substitutions that raise viral activity against a_1 alone, 38 are closed by a_2, and 34 run the other way. For all cocktail pairs one member is the stronger binder, leaving 0.62 of the virus active on average against 0.74 for its partner and accounting for 63% of the bound antigen. Viral activity under the pair still sits below it at every step (Figure 6c), so the weaker member continues to contribute. Section F gives the per-member landscapes.

7 Conclusion and Future Work

Co-ADIOS designs antibody cocktails as the leader of a Stackelberg evolutionary game, optimising the pair over the viral trajectories it induces rather than against the current variant only. On influenza H5 the resulting cocktails keep more of the virus neutralised across a hundred steps of escape than myopic cocktails, single shapers and uncoordinated pairs, and a second antibody helps only when the two are designed together. They are weaker against the wildtype v_0 in exchange, which is the trade the shaping horizon makes. What the cocktails change is the escape landscape itself: they leave the virus far fewer single mutations worth taking, and divide the ones that remain between two antibodies that share a binding site. These results come from a lattice simulator, one antigen and cocktail pairs, and whether the same coordination emerges for larger combinations and under a higher-fidelity binding model is an open question for future work.

AI USE STATEMENT

In this work, we used generative AI tools to assist with writing and refactoring code, and to draft and edit prose. All AI-assisted code was read, run and tested by the authors: the binding implementation is checked against the reference C++ implementation of Absolut! (Section B), and the analysis code is covered by unit tests in the released repository. Every number reported in this paper comes from running that code, not from a language model. AI tools were also used to search for related work, and every resulting reference was read by the authors before citing it. All citations were added manually by the authors. We did not use generative AI to produce research ideas, experimental designs, proofs or results. We have reviewed all AI-assisted work and take responsibility for the final content of this paper, including its text, claims and artifacts.

ETHICS STATEMENT

This work is computational throughout. It involves no human subjects, no personal data and no laboratory work, and the antigens are public PDB structures. Antibodies and viruses are fixed-length symbol strings scored by a published lattice model, not designs intended for synthesis: the simulator resolves binding on a 5.25 Å lattice with a contact potential, which is far too coarse for real-world deployment. The application we study is adjacent to pandemic preparedness, in that the objective rewards therapies that keep a simulated virus neutralised as it adapts. We note nonetheless that methods which anticipate viral escape can in principle be inverted to search for escape, and that this risk is inherent to any escape-aware design method rather than specific to ours; the coarseness of the simulator significantly limits what either direction can produce here.

REPRODUCIBILITY STATEMENT

An anonymised copy of the code is available at `https://anonymous.4open.science/r/co-adios-D1CF/`. It contains the binding simulator, the inner and outer loops, the configuration file of every arm reported here, and the scripts that produce each figure and table, together with unit tests. The game and its payoffs are defined in Sections 4.1 to 4.3, with the activity calibration and the reward weights in Section D. Results are reported over 20 training seeds with held-out evaluation seeds. The discretised antigens are regenerated from the public Absolut! distribution following `docs/data/README.md`.

REFERENCES

Nezam Afdhal, Stefan Zeuzem, Paul Kwo, Mario Chojkier, Norman Gitlin, Massimo Puoti, Manuel Romero-Gomez, Jean-Pierre Zarski, Kosh Agarwal, Peter Buggisch, Graham R. Foster, Norbert Bräu, Maria Buti, Ira M. Jacobson, G. Mani Subramanian, Xiao Ding, Hongmei Mo, Jenny C. Yang, Phillip S. Pang, William T. Symonds, John G. McHutchison, Andrew J. Muir, Alessandra Mangia, Patrick Marcellin, and ION-1 Investigators. Ledipasvir and sofosbuvir for untreated HCV genotype 1 infection. *The New England Journal of Medicine*, 370(20):1889–1898, May 2014. ISSN 1533-4406. doi: 10.1056/NEJMoa1402454.

Alina Baum, Benjamin O. Fulton, Elzbieta Wloga, Richard Copin, Kristen E. Pascal, Vincenzo Russo, Stephanie Giordano, Kathryn Lanza, Nicole Negron, Min Ni, Yi Wei, Gurinder S. Atwal, Andrew J. Murphy, Neil Stahl, George D. Yancopoulos, and Christos A. Kyratsous. Antibody cocktail to SARS-CoV-2 spike protein prevents rapid mutational escape seen with individual antibodies. *Science*, 369(6506):1014–1018, August 2020. ISSN 1095-9203. doi: 10.1126/science.abd0831.

M. J. E. Brill, A. N. Kristoffersson, C. Zhao, E. I. Nielsen, and L. E. Friberg. Semi-mechanistic pharmacokinetic–pharmacodynamic modelling of antibiotic drug combinations. *Clinical Microbiology and Infection*, 24(7):697–706, July 2018. ISSN 1198-743X. doi: 10.1016/j.cmi.2017.11.023. URL `https://www.sciencedirect.com/science/article/pii/S1198743X17306705`.

Jinghong Chen, Anqi Lin, Aimin Jiang, Chang Qi, Zaoqu Liu, Quan Cheng, Shuofeng Yuan, and Peng Luo. Computational frameworks transform antagonism to synergy in optimizing combination therapies. *npj Digital Medicine*, 8(1):44, January 2025. ISSN 2398-

6352. doi: 10.1038/s41746-025-01435-2. URL https://www.nature.com/articles/s41746-025-01435-2.

Adam L. Corper, Maninder K. Sohi, Vincent R. Bonagura, Michael Steinitz, Royston Jefferis, Arnold Feinstein, Dennis Beale, Michael J. Taussig, and Brian J. Sutton. Structure of human IgM rheumatoid factor Fab bound to its autoantigen IgG Fc reveals a novel topology of antibody—antigen interaction. *Nature Structural Biology*, 4(5):374–381, May 1997. ISSN 1545-9985. doi: 10.1038/nsb0597-374. URL https://www.nature.com/articles/nsb0597-374.

A.E. Eiben and J.E. Smith. *Introduction to Evolutionary Computing*. Natural Computing Series. Springer, Berlin, Heidelberg, 2015. ISBN 978-3-662-44873-1 978-3-662-44874-8. doi: 10.1007/978-3-662-44874-8. URL https://link.springer.com/10.1007/978-3-662-44874-8.

Tal Einav and Jesse D. Bloom. When two are better than one: Modeling the mechanisms of antibody mixtures. *PLOS Computational Biology*, 16(5):e1007830, May 2020. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1007830. URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1007830.

Jakob N. Foerster, Richard Y. Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. Learning with Opponent-Learning Awareness, September 2018. URL http://arxiv.org/abs/1709.04326. arXiv:1709.04326 [cs].

Scott M. Hammer, Kathleen E. Squires, Michael D. Hughes, Janet M. Grimes, Lisa M. Demeter, Judith S. Currier, Joseph J. Eron, Judith E. Feinberg, Henry H. Balfour, Lawrence R. Deyton, Jeffrey A. Chodakewitz, Margaret A. Fischl, John P. Phair, Louise Pedneault, Bach-Yen Nguyen, and Jon C. Cook. A Controlled Trial of Two Nucleoside Analogues plus Indinavir in Persons with Human Immunodeficiency Virus Infection and CD4 Cell Counts of 200 per Cubic Millimeter or Less. *New England Journal of Medicine*, 337(11):725–733, September 1997. ISSN 0028-4793. doi: 10.1056/NEJM199709113371101. URL https://www.nejm.org/doi/full/10.1056/NEJM199709113371101.

S Henikoff and J G Henikoff. Amino acid substitution matrices from protein blocks. *Proceedings of the National Academy of Sciences*, 89(22):10915–10919, November 1992. doi: 10.1073/pnas.89.22.10915. URL https://www.pnas.org/doi/abs/10.1073/pnas.89.22.10915.

Brian L. Hie, Varun R. Shanker, Duo Xu, Theodora U. J. Bruun, Payton A. Weidenbacher, Shaogeng Tang, Wesley Wu, John E. Pak, and Peter S. Kim. Efficient evolution of human antibodies from general protein language models. *Nature Biotechnology*, 42(2):275–283, February 2024. ISSN 1546-1696. doi: 10.1038/s41587-023-01763-2. URL https://www.nature.com/articles/s41587-023-01763-2.

Jumpei Ito, Adam Strange, Wei Liu, Gustav Joas, Spyros Lytras, and Kei Sato. A protein language model for exploring viral fitness landscapes. *Nature Communications*, 16(1):4236, May 2025. ISSN 2041-1723. doi: 10.1038/s41467-025-59422-w. URL https://www.nature.com/articles/s41467-025-59422-w.

Sara Joubbi, Alessio Micheli, Paolo Milazzo, Giuseppe Maccari, Giorgio Ciano, Dario Cardamone, and Duccio Medini. Antibody design using deep learning: from sequence and structure design to affinity maturation. *Briefings in Bioinformatics*, 25(4):bbae307, July 2024. ISSN 1477-4054. doi: 10.1093/bib/bbae307. URL https://doi.org/10.1093/bib/bbae307.

Nicole L. Kallewaard, Davide Corti, Patrick J. Collins, Ursula Neu, Josephine M. McAuliffe, Ebony Benjamin, Leslie Wachter-Rosati, Frances J. Palmer-Hill, Andy Q. Yuan, Philip A. Walker, Matthias K. Vorlaender, Siro Bianchi, Barbara Guarino, Anna De Marco, Fabrizia Vanzetta, Gloria Agatic, Mathilde Foglierini, Debora Pinna, Blanca Fernandez-Rodriguez, Alexander Fruehwirth, Chiara Silacci, Roksana W. Ogrodowicz, Stephen R. Martin, Federica Sallusto, JoAnn A. Suzich, Antonio Lanzavecchia, Qing Zhu, Steven J. Gamblin, and John J. Skehel. Structure and Function Analysis of an Antibody Recognizing All Influenza A Subtypes. *Cell*, 166(3):596–608, July 2016. ISSN 0092-8674, 1097-4172. doi: 10.1016/j.cell.2016.05.073. URL https://www.cell.com/cell/abstract/S0092-8674(16)30720-6.

Asif Khan, Alexander I. Cowen-Rivers, Antoine Grosnit, Derrick-Goh-Xin Deik, Philippe A. Robert, Victor Greiff, Eva Smorodina, Puneet Rawat, Rahmad Akbar, Kamil Dreczkowski, Rasul Tutunov, Dany Bou-Ammar, Jun Wang, Amos Storkey, and Haitham Bou-Ammar. Toward real-world automated antibody design with combinatorial Bayesian optimization. *Cell Reports Methods*, 3(1): 100374, January 2023. ISSN 2667-2375. doi: 10.1016/j.crmeth.2022.100374.

Xiangzhe Kong, Wenbing Huang, and Yang Liu. End-to-End Full-Atom Antibody Design, February 2023. URL https://arxiv.org/abs/2302.00203v4.

Denise Kühnert, Chieh-Hsi Wu, and Alexei J. Drummond. Phylogenetic and epidemic modeling of rapidly evolving infectious diseases. *Infection, Genetics and Evolution*, 11(8):1825–1841, December 2011. ISSN 1567-1348. doi: 10.1016/j.meegid.2011.08.005. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC7106223/.

Andrew Leaver-Fay, Ron Jacak, P. Benjamin Stranges, and Brian Kuhlman. A Generic Program for Multistate Protein Design. *PLOS ONE*, 6(7):e20937, July 2011. ISSN 1932-6203. doi: 10.1371/journal.pone.0020937. URL https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0020937.

Jingzhi Lou, Weiwen Liang, Lirong Cao, Inchi Hu, Shi Zhao, Zigui Chen, Renee Wan Yi Chan, Peter Pak Hang Cheung, Hong Zheng, Caiqi Liu, Qi Li, Marc Ka Chun Chong, Yexian Zhang, Eng-kiong Yeoh, Paul Kay-Sheung Chan, Benny Chung Ying Zee, Chris Ka Pun Mok, and Maggie Haitian Wang. Predictive evolutionary modelling for influenza virus by site-based dynamics of mutations. *Nature Communications*, 15(1):2546, March 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-46918-0. URL https://www.nature.com/articles/s41467-024-46918-0.

Chris Lu, Timon Willi, Christian Schroeder de Witt, and Jakob Foerster. Model-Free Opponent Shaping, November 2022. URL http://arxiv.org/abs/2205.01447. arXiv:2205.01447 [cs].

Shitong Luo, Yufeng Su, Xingang Peng, Sheng Wang, Jian Peng, and Jianzhu Ma. Antigen-Specific Antibody Design and Optimization with Diffusion-Based Generative Models for Protein Structures. October 2022. URL https://openreview.net/forum?id=jSorGn2Tjg.

Xiaoliang Ma, Xiaodong Li, Qingfu Zhang, Ke Tang, Zhengping Liang, Weixin Xie, and Zexuan Zhu. A Survey on Cooperative Co-Evolutionary Algorithms. *IEEE Transactions on Evolutionary Computation*, 23(3):421–441, June 2019. ISSN 1941-0026. doi: 10.1109/TEVC.2018.2868770. URL https://ieeexplore.ieee.org/document/8454482.

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, Faeezah 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. ISSN 1533-4406. doi: 10.1056/NEJMoa2102214.

Richard A. Neher, Trevor Bedford, Rodney S. Daniels, Colin A. Russell, and Boris I. Shraiman. Prediction, dynamics, and visualization of antigenic phenotypes of seasonal influenza viruses. *Proceedings of the National Academy of Sciences*, 113(12):E1701–E1709, March 2016. doi: 10.1073/pnas.1525578113. URL https://www.pnas.org/doi/10.1073/pnas.1525578113.

Fritz Obermeyer, Martin Jankowiak, Nikolaos Barkas, Stephen F. Schaffner, Jesse D. Pyle, Leonid Yurkovetskiy, Matteo Bosso, Daniel J. Park, Mehrtash Babadi, Bronwyn L. MacInnis, Jeremy Luban, Pardis C. Sabeti, and Jacob E. Lemieux. Analysis of 6.4 million SARS-CoV-2 genomes identifies mutations associated with fitness. *Science*, 376(6599):1327–1332, June 2022. doi: 10.

1126/science.abm1208. URL https://www.science.org/doi/10.1126/science.abm1208.

Mitchell A. Potter and Kenneth A. De Jong. A cooperative coevolutionary approach to function optimization. In Yuval Davidor, Hans-Paul Schwefel, and Reinhard Männer (eds.), *Parallel Problem Solving from Nature — PPSN III*, pp. 249–257, Berlin, Heidelberg, 1994. Springer. ISBN 978-3-540-49001-2. doi: 10.1007/3-540-58484-6_269.

Matthew I J Raybould, Claire Marks, Alan P Lewis, Jiye Shi, Alexander Bujotzek, Bruck Taddese, and Charlotte M Deane. Thera-SAbDab: the Therapeutic Structural Antibody Database. *Nucleic Acids Research*, 48(D1):D383–D388, January 2020. ISSN 0305-1048. doi: 10.1093/nar/gkz827. URL https://doi.org/10.1093/nar/gkz827.

Philippe A. Robert, Rahmad Akbar, Robert Frank, Milena Pavlović, Michael Widrich, Igor Snapkov, Andrei Slabodkin, Maria Chernigovskaya, Lonneke Scheffer, Eva Smorodina, Puneet Rawat, Brij Bhushan Mehta, Mai Ha Vu, Ingvild Frøberg Mathisen, Aurél Prósz, Krzysztof Abram, Alex Olar, Enkelejda Miho, Dag Trygve Tryslew Haug, Fridtjof Lund-Johansen, Sepp Hochreiter, Ingrid Hobæk Haff, Günter Klambauer, Geir Kjetil Sandve, and Victor Greiff. 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.

Alexander M. Sevy, Swetasudha Panda, James E. Crowe Jr, Jens Meiler, and Yevgeniy Vorobeychik. Integrating linear optimization with structural modeling to increase HIV neutralization breadth. *PLOS Computational Biology*, 14(2):e1005999, February 2018. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1005999. URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1005999.

Sisi Shan, Shitong Luo, Ziqing Yang, Junxian Hong, Yufeng Su, Fan Ding, Lili Fu, Chenyu Li, Peng Chen, Jianzhu Ma, Xuanling Shi, Qi Zhang, Bonnie Berger, Linqi Zhang, and Jian Peng. Deep learning guided optimization of human antibody against SARS-CoV-2 variants with broad neutralization. *Proceedings of the National Academy of Sciences of the United States of America*, 119(11):e2122954119, March 2022. ISSN 1091-6490. doi: 10.1073/pnas.2122954119.

Varun R. Shanker, Theodora U. J. Bruun, Brian L. Hie, and Peter S. Kim. 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/doi/10.1126/science.adk8946.

Zeenat A. Shyr, Yu-Shan Cheng, Donald C. Lo, and Wei Zheng. Drug combination therapy for emerging viral diseases. *Drug Discovery Today*, 26(10):2367–2376, October 2021. ISSN 1359-6446. doi: 10.1016/j.drudis.2021.05.008. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8139175/.

Tyler N. Starr, Allison J. Greaney, Amin Addetia, William W. Hannon, Manish C. Choudhary, Adam S. Dingens, Jonathan Z. Li, and Jesse D. Bloom. Prospective mapping of viral mutations that escape antibodies used to treat COVID-19. *Science*, 371(6531):850–854, February 2021. ISSN 1095-9203. doi: 10.1126/science.abf9302.

Kateřina Staňková, Joel S. Brown, William S. Dalton, and Robert A. Gatenby. Optimizing Cancer Treatment Using Game Theory. *JAMA oncology*, 5(1):96–103, January 2019. ISSN 2374-2437. doi: 10.1001/jamaoncol.2018.3395. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC6947530/.

Alexander Stein, Monica Salvioli, Hasti Garjani, Johan Dubbeldam, Yannick Viossat, Joel S. Brown, and Kateřina Staňková. Stackelberg evolutionary game theory: how to manage evolving systems. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378(1876):20210495, March 2023. ISSN 0962-8436. doi: 10.1098/rstb.2021.0495. URL https://doi.org/10.1098/rstb.2021.0495.

Shih-Chieh Su, Tzu-Jing Yang, Pei-Yu Yu, Kang-Hao Liang, Wan-Yu Chen, Chun-Wei Yang, Hsiu-Ting Lin, Mei-Jung Wang, Ruei-Min Lu, Hsien-Cheng Tso, Meng-Jhe Chung, Tzung-Yang Hsieh, Yu-Ling Chang, Shin-Chang Lin, Fang-Yu Hsu, Feng-Yi Ke, Yi-Hsuan Wu, Yu-Chyi Hwang, I.-Ju Liu, Jian-Jong Liang, Chun-Che Liao, Hui-Ying Ko, Cheng-Pu Sun, Ping-Yi Wu, Jia-Tsrong Jan, Yuan-Chih Chang, Yi-Ling Lin, Mi-Hua Tao, Shang-Te Danny Hsu, and Han-Chung Wu. Structure-guided antibody cocktail for prevention and treatment of COVID-19. *PLoS pathogens*, 17(10):e1009704, October 2021. ISSN 1553-7374. doi: 10.1371/journal.ppat.1009704.

Nicole N. Thadani, Sarah Gurev, Pascal Notin, Noor Youssef, Nathan J. Rollins, Daniel Ritter, Chris Sander, Yarin Gal, and Debora S. Marks. Learning from prepandemic data to forecast viral escape. *Nature*, 622(7984):818–825, 2023. ISSN 0028-0836. doi: 10.1038/s41586-023-06617-0. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10599991/.

Sebastian Towers, Aleksandra Kalisz, Philippe A. Robert, Alicia Higueruelo, Francesca Vianello, Ming-Han Chloe Tsai, Harrison Steel, and Jakob N. Foerster. ADIOS: Antibody Development via Opponent Shaping, June 2025. URL http://arxiv.org/abs/2409.10588. arXiv:2409.10588 [q-bio].

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. ISSN 0883-0185. doi: 10.3109/08830189209055568.

Yannick Vogt, Mehdi Naouar, Maria Kalweit, Christoph Cornelius Miething, Justus Duyster, Roland Mertelsmann, Gabriel Kalweit, and Joschka Boedecker. Stable Online and Offline Reinforcement Learning for Antibody CDRH3 Design, November 2023. URL http://arxiv.org/abs/2401.05341. arXiv:2401.05341 [q-bio.BM].

David M. Weinreich, Sumathi Sivapalasingam, Thomas Norton, Shazia Ali, Haitao Gao, Rafia Bhore, Jing Xiao, Andrea T. Hooper, Jennifer D. Hamilton, Bret J. Musser, Diana Rofail, Mohamed Hussein, Joseph Im, Dominique Y. Atmodjo, Christina Perry, Cynthia Pan, Adnan Mahmood, Romana Hosain, John D. Davis, Kenneth C. Turner, Alina Baum, Christos A. Kyratsous, Yunji Kim, Amanda Cook, Wendy Kampman, Lilia Roque-Guerrero, Gerard Acloque, Hessam Aazami, Kevin Cannon, J. Abraham Simón-Campos, Joseph A. Bocchini, Bari Kowal, A. Thomas DiCioccio, Yuhwen Soo, Gregory P. Geba, Neil Stahl, Leah Lipsich, Ned Braunstein, Gary Herman, George D. Yancopoulos, and Trial Investigators. REGEN-COV Antibody Combination and Outcomes in Outpatients with Covid-19. *The New England Journal of Medicine*, 385(23):e81, December 2021. ISSN 1533-4406. doi: 10.1056/NEJMoa2108163.

James R. R. Whittle, Ruijun Zhang, Surender Khurana, Lisa R. King, Jody Manischewitz, Hana Golding, Philip R. Dormitzer, Barton F. Haynes, Emmanuel B. Walter, M. Anthony Moody, Thomas B. Kepler, Hua-Xin Liao, and Stephen C. Harrison. Broadly neutralizing human antibody that recognizes the receptor-binding pocket of influenza virus hemagglutinin. *Proceedings of the National Academy of Sciences*, 108(34):14216–14221, August 2011. doi: 10.1073/pnas.1111497108. URL https://www.pnas.org/doi/full/10.1073/pnas.1111497108.

Timon Willi, Alistair Letcher, Johannes Treutlein, and Jakob Foerster. COLA: Consistent Learning with Opponent-Learning Awareness, June 2022. URL http://arxiv.org/abs/2203.04098. arXiv:2203.04098 [cs].

Hui-Ling Yen, Jerry R. Aldridge, Adrianus C. M. Boon, Natalia A. Ilyushina, Rachelle Salomon, Diane J. Hulse-Post, Henju Marjuki, John Franks, David A. Boltz, Dorothy Bush, Aleksandr S. Lipatov, Richard J. Webby, Jerold E. Rehg, and Robert G. Webster. Changes in H5N1 influenza virus hemagglutinin receptor binding domain affect systemic spread. *Proceedings of the National Academy of Sciences of the United States of America*, 106(1):286–291, January 2009. ISSN 1091-6490. doi: 10.1073/pnas.0811052106.

Xueyong Zhu, Yong-Hui Guo, Tao Jiang, Ya-Di Wang, Kwok-Hung Chan, Xiao-Feng Li, Wenli Yu, Ryan McBride, James C. Paulson, Kwok-Yung Yuen, Cheng-Feng Qin, Xiao-Yan Che, and Ian A. Wilson. A Unique and Conserved Neutralization Epitope in H5N1 Influenza Viruses Identified by an Antibody against the A/Goose/Guangdong/1/96 Hemagglutinin. *Journal of Virology*, 87(23):12619–12635, December 2013. ISSN 0022-538X. doi: 10.1128/JVI.01577-13. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC3838140/.

A Biological Glossary

We collect definitions of the biological terms used throughout the paper here for the reader's convenience; readers from a machine-learning background may find this a useful reference.

Antibody. A Y-shaped protein produced by the immune system that binds specific antigens via its hypervariable loops to neutralise pathogens. We represent antibodies by their CDRH3 sequence (see below).

Antigen. A molecule (typically a protein) that an antibody recognises and binds. For viruses, the antigen is usually a surface protein (e.g., influenza haemagglutinin).

Residue. A single amino-acid unit at a fixed position within a protein sequence. "Viral residues" are the positions in the antigen sequence the virus can mutate.

CDRH3. The third complementarity-determining region of an antibody's heavy chain; the most variable loop and the dominant determinant of antigen-binding specificity. We optimise the cocktail entirely at this region.

Pose. A specific spatial configuration of an antibody on an antigen. In lattice models, an enumerable discrete docking arrangement specifying which antibody residues contact which antigen residues.

Lattice. A 3D grid onto which Absolut!-style models discretise the antigen chain, allowing fast enumeration and scoring of antibody poses.

Miyazawa–Jernigan contact energies. A 20×20 lookup table of empirical pairwise energies between amino acids, derived from statistics over known protein structures. The energy of a pose is the sum of these over the residue pairs it brings into contact.

Haemagglutinin. The surface protein influenza uses to attach to and enter host cells, and the antigen in our experiments. Its *head* carries the receptor-binding site and most antibody escape mutations, while its *stalk* is more conserved. We keep only poses on the head.

Receptor and receptor-binding site. The molecule on the host cell the virus must attach to in order to infect it, and the part of the antigen that does the attaching. A virus that mutates away from an antibody but loses receptor binding is not viable, which is what the viral-target term of the reward represents.

Neutralisation. Binding by an antibody that blocks a virus from infecting cells. We measure it as the fraction of virus bound, one minus the fractional activity.

Binding affinity / binding energy. Two related measures of how tightly an antibody binds its antigen. Affinity is "high" when binding is tight, and binding energy is "negative" when binding is favourable. We use $E(a, v)$ for binding energy throughout.

Epitope. The site on an antigen that a given antibody binds. Two antibodies with different epitopes contact different parts of the same virus.

Glycan. A sugar chain attached to a protein surface. Glycans shield parts of the viral surface and take part in receptor binding, and lattice models cannot represent them.

Antitarget. A protein native to the host that the antibody must avoid binding to remain clinically safe, since antibodies that bind host proteins can cause autoimmunity and toxicity. Ours is a human Fc fragment, denoted t_a^-.

Viral escape. A set of viral mutations that reduce or 'escape' antibody recognition while preserving infectivity.

Wildtype. The unmutated reference strain of the virus, denoted v_0, that serves as the starting point of the inner-loop viral evolution.

B Lattice Discretisation and Pose Enumeration

The Absolut! (Robert et al., 2022) simulator pre-enumerates antibody-antigen binding poses on a 3D lattice surrounding the antigen, enabling fast scoring of arbitrary CDRH3 sequences (Figure B.1). For each (antigen sequence, antibody CDRH3) pair, the binding score is the minimum energy across pre-computed poses, where each pose's energy sums pairwise contact energies over its contact graph.

Taking the minimum assumes a steep energy funnel, in which the free energy of the complex is dominated by its best pose. Summing Boltzmann weights over low-energy poses would additionally capture the entropic width of the binding basin, but we use the global minimum throughout, as Absolut! and ADIOS do.

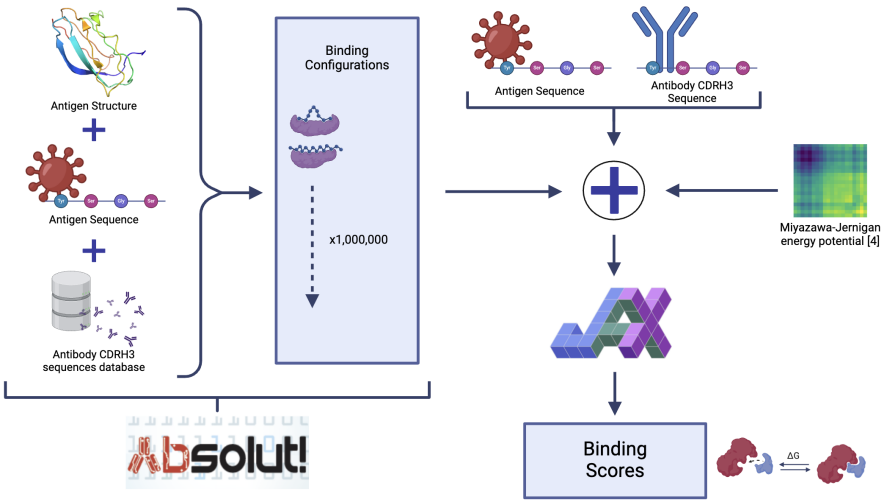

Figure B.1: **The JAX-Absolut pipeline.** Left: Absolut! pre-enumerates binding poses for a fixed antigen using its 3D structure and lattice discretisation. Right: any (antigen sequence, antibody CDRH3) pair is scored in JAX by summing Miyazawa–Jernigan contact energies over each pose's contact graph. The lowest pose energy gives the binding score $E(a, v)$.

Figure B.2 shows the lattice discretisations for the two antigens used in our experiments: the Influenza H5 antigen (PDB 5JW4) and the antibody antitarget (PDB 1ADQ Fc fragment). Each amino-acid residue is mapped to a coloured sphere on the 3D lattice surface; antibody poses are enumerated as discrete configurations against this surface.

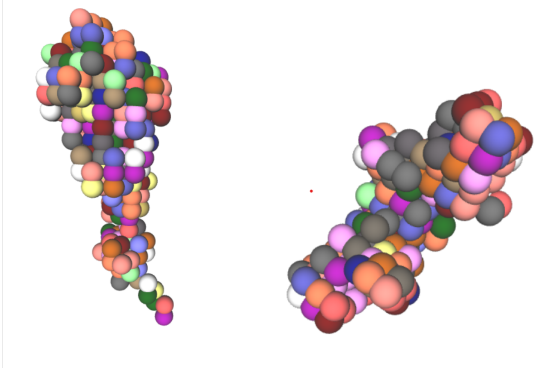

Figure B.2: **Lattice discretisations of the antigens used in our experiments.** Left: Influenza H5 antigen (PDB 5JW4). Right: Antibody antitarget (PDB 1ADQ Fc fragment). Each sphere represents a single amino-acid residue at its lattice position; colours indicate residue identity.

C Fractional Activity from Binding Energies

This section derives the fractional activity used throughout the paper from the binding energies the simulator returns, and the effective single energy a cocktail is equivalent to.

One binder. A virus is either free or bound, and we assume a bound virus is neutralised, so the fractional activity is the probability of the free state. With antibody concentration c and dissociation constant K_D, the two Boltzmann weights give the partition function $Z = 1 + c/K_D$ and

$$\text{FA} = \frac{1}{Z} = \frac{1}{1 + c/K_D}. \tag{C.1}$$

The binding free energy the simulator returns fixes K_D through $E(a, v) = RT \ln K_D$, so that $K_D = e^{E(a,v)/RT}$ and $c/K_D = c\, e^{-E(a,v)/RT}$, where RT sets the thermal energy scale. A more negative energy is a tighter binder, a larger Boltzmann weight and a smaller free fraction, as expected.

A cocktail. The poses of both antibodies are enumerated against the same antigen surface, so we treat the members of a cocktail as competing for one site rather than occupying independent ones. A virus is then free, bound by a_1, or bound by a_2, giving

$$\text{FA}(v; a_1, a_2) = \frac{1}{1 + c_1 e^{-E(a_1,v)/RT} + c_2 e^{-E(a_2,v)/RT}}, \tag{C.2}$$

which reduces to the single-binder expression when either concentration is zero. The same expression scores the antitarget, and the viral target under its own curve.

Calibration. In practice we do not work with absolute concentrations. Writing d_i for the fraction of the total dose given to a_i, we replace each Boltzmann weight by $r_i = d_i \exp\big(-(E(a_i, v) - e_{50})/t_{\text{eff}}\big)$ and compute $\text{FA} = 1/(1 + \sum_i r_i)$. Here t_{eff} plays the role of RT and e_{50} is the energy at which a unit dose leaves half the population free, both measured per antigen and pose set as described in Section D. For a single binder this is a logistic in the binding energy with midpoint e_{50} and width t_{eff}.

The effective single energy. ADIOS optimises a single binding energy, so it is useful to ask which energy a cocktail is equivalent to. Equating the free fraction of the pair with that of one binder carrying the same total dose gives

$$E_{\text{net}}(a_1, a_2, v) = -t_{\text{eff}} \ln \left(d_1 e^{-E(a_1,v)/t_{\text{eff}}} + d_2 e^{-E(a_2,v)/t_{\text{eff}}} \right), \tag{C.3}$$

a soft minimum of the two energies weighted by dose. As $t_{\text{eff}} \to 0$ the stronger binder dominates and $E_{\text{net}} \to \min_i E(a_i, v)$, while as t_{eff} grows it approaches the dose-weighted average $d_1 E(a_1, v) + d_2 E(a_2, v)$. A cocktail is therefore never worse than its weaker member and never better than its stronger one at equal total dose, which is why a second antibody pays off through the trajectory it induces rather than through the instantaneous score.

D Scoring, Calibration and Dosing

This section records the pose sets, activity curves and dosing model behind the numbers reported in Section 5.2.

Pose sets. Absolut! enumerates every lattice pose an eleven-residue antibody can adopt against a discretised antigen. For 5JW4_A we prune that enumeration to the globular head and use two nested subsets, the coarse and fine resolutions of Section 5.1: 933 poses covering 42 antigen residues, on which we optimise, and 2763 poses covering 44, on which every reported score is recomputed. The antitarget 1ADQ_A is scored on its own nested pair, 899 poses over 133 residues at the coarse resolution and 2273 at the fine one. The fine set contains the coarse one, so the two differ by the poses the optimisation never scored, which is the transfer gap the reported numbers cross.

Activity curves. The model of Section C maps a binding energy to a fractional activity through two measured parameters, the midpoint e_{50} and the width t_{eff}. They are specific to the antigen, the pose set and the reaction scored, so our experiments use six curves: antibody-antigen, antibody-antitarget and virus-target, each at both resolutions (Table D.1). Both ends of each curve are measured by

probing the simulator rather than assumed. For the antibody curves the strong end is a hill-climbed optimum and the weak end is the energy a well-optimised antibody is driven to by a virus escaping it, so that the trajectory stays inside the calibrated band instead of running off its tail. The virus-target curve is instead anchored at the wildtype, which starts the run at half activity with the slope at its steepest, because the virus begins at one sequence and moves locally; its width is set so that one substitution is resolvable.

Table D.1: The six activity curves used in our experiments. Energies are in Absolut! units.

reaction	pose set	e_{50}	t_{eff}
antibody-antigen (5JW4_A)	coarse	-74.62	11.75
antibody-antigen (5JW4_A)	fine	-77.90	10.25
antibody-antitarget (1ADQ_A)	coarse	-94.22	8.33
antibody-antitarget (1ADQ_A)	fine	-94.02	8.25
virus-target (5JW4_A)	coarse	-75.58	2.36
virus-target (5JW4_A)	fine	-75.58	2.34

Dose and therapy lifetime. A therapy carries a total dose of 1.0, divided between its members by a fixed ratio, so that a cocktail and a monotherapy are compared at equal total dose. The dose then decays over a trajectory, as a real therapy is cleared from the body: at viral step t each member is present at $d_i e^{-t/\tau}$, with the therapy's lifetime $\tau = 86$ steps. Since dose and energy enter the activity model together, this is equivalent to weakening each binding energy by $(t_{\text{eff}}/\tau)\, t$, so a therapy loses about one curve width over the hundred steps we simulate and cannot suppress the virus indefinitely by binding it once.

E Scenario 1 Reward Components

The held-out return is one scalar per design, so it cannot say which term moved. Figure E.1 and Figure E.2 take it apart along the rollout: how much virus the therapy neutralises, how much antitarget it leaves alone, what the viral-target penalty costs the virus, and where the virus's own receptor binding ends up. The first varies the optimisation horizon, the second the design method at the full horizon.

Two patterns recur. Myopic designs neutralise the wildtype hardest and lose that advantage within roughly ten viral steps, which is the trajectory-level version of the endpoint result in Section 6.1. Longer horizons also accept more antitarget binding, so part of what the horizon buys is spent on specificity rather than won outright.

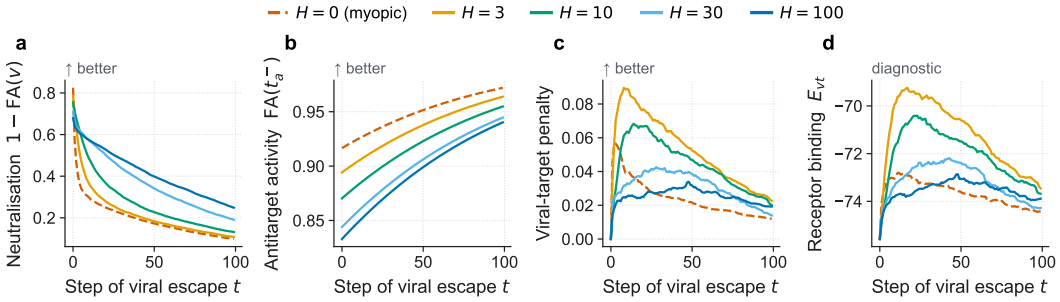

Figure E.1: **Reward components across the optimisation horizon.** Each panel averages over 20 seeds and their held-out rollouts. **a** Fraction of virus neutralised. **b** Antitarget activity, where higher means less antitarget bound. **c** The viral-target penalty the virus pays for losing receptor binding. **d** The virus's own receptor binding energy, a diagnostic rather than a term of the objective.

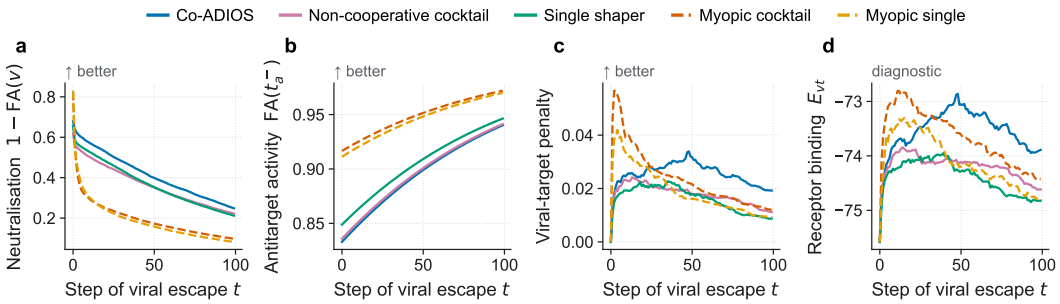

Figure E.2: **Reward components by design method.** Panels as in Figure E.1, with each arm of Section 6.1 at its own horizon.

F Per-Member Escape Landscapes

Viral reward is a property of the therapy the virus faces, so it exists for a pair and not for one member of it. A member is therefore read on the activity axis instead: the fraction of virus it would leave active at the dose it is given inside the cocktail, which is the term the reward is built from. Binding energies need no dose at all, so both members come off the same scan rather than from a second pass that re-doses each one as a monotherapy.

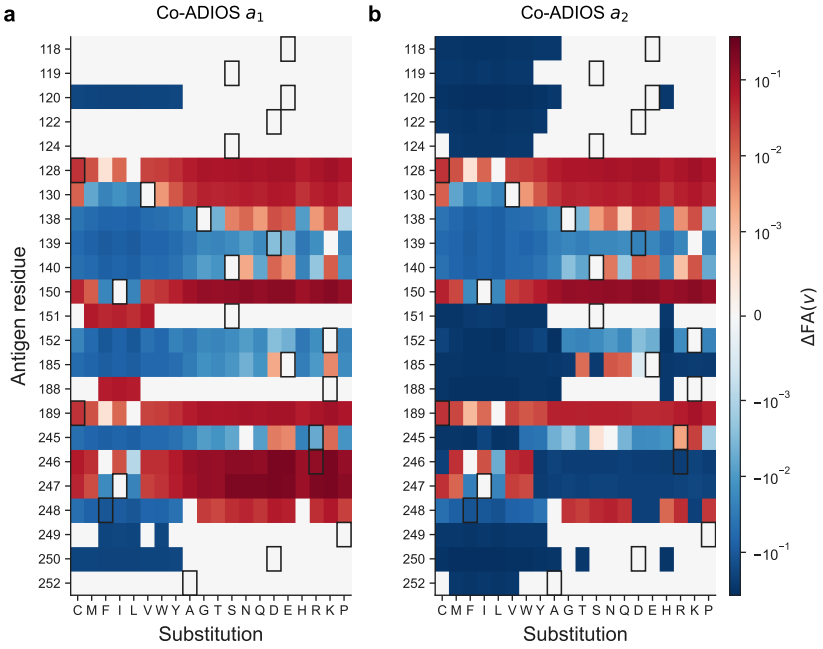

Figure F.1: **Viral Escape Landscapes for two Co-ADIOS cocktail members.** Change in viral activity for every single substitution of the escape set, against each member of the cocktail of Figure 6 at its own dose. Rows, columns and outlines as in Figure 6; red favours the virus. The two share a site but not a set of open substitutions.

G Assembling a Combination from Existing Antibodies

A combination can be designed, or it can be assembled from antibodies that already exist. This section asks whether the objective of Section 4.3 ranks combinations usefully inside a real therapeutic library, and how far the best such pair sits from a designed one. No optimiser enters the claim: the pair space is small enough to enumerate, which separates ranking combinations well from searching well.

Setup. The library is the Thera-SAbDab therapeutic set (Raybould et al., 2020) docked against 5JW4_A, 877 antibodies whose CDRH3s are reduced to their best eleven-residue window. Every unordered pair, including self-pairs, gives 385,003 combinations, and a self-pair at equal doses is one antibody at full dose, so the monotherapy control lives in the same enumeration. A maximum over 385,003 noisy estimates would select the pair whose seed noise was luckiest, so we select on the 20 training seeds and verify the top 200 on 20 disjoint held-out seeds. The 877 self-pairs were verified in full rather than as a shortlist, because selection and held-out ranking disagree at the top of that set.

Results. The objective ranks combinations usefully inside a real library: the best pair beats the best single antibody in it by 0.0305 (Table G.1), a gain of the same sign and roughly the same size as the one a designed cocktail has over a designed single shaper. The library does not, however, reach what design reaches. The best assembled pair sits 0.0926 below the designed cocktail, and a designed *single* antibody already beats the best pair drawn from 385,003 combinations of existing ones.

Table G.1: The best combinations and single antibodies in the therapeutic library, verified on 20 held-out seeds, against the designs of Section 5.3. Higher is better.

	escape-averaged return
best pairs in the library	
Petosemtamab + Monalizumab	-0.7537
Petosemtamab + Tibulizumab	-0.7621
Petosemtamab + Tabalumab	-0.7621
best single antibodies in the library	
Monalizumab	-0.7843
Gremubamab	-0.8059
Petosemtamab	-0.8083
designed	
Co-ADIOS cocktail	$\mathbf{-0.6611}$
single shaper	-0.7032

Caveats. This is a maximum over an enumeration verified on held-out seeds, not a paired test, and it carries no p value: the statement is that the best pair beats the best single by this much, not that pairs beat singles on average. The winning pair is also carried by a promiscuous binder, since Monalizumab leaves 0.447 of the antitarget free against a library median of 0.858, and across this pool antigen potency and antitarget stickiness are close to one axis, so a selectivity threshold would change the answer. Finally, the library is a therapeutic set docked against this antigen rather than a set of antibodies raised against it, and their real targets are unrelated to influenza, so a pair found here is a statement about the objective and not a deployable therapy.