

FUTURE-PROOF VACCINE DESIGN WITH A GENERATIVE MODEL OF ANTIBODY CROSS-REACTIVITY

Sarah Gurev^{*†}

EECS
Massachusetts Institute of Technology

Noor Youssef^{*†}

Systems Biology
Harvard Medical School

Hannah Pierce-Hoffman

Biomedical Informatics
Harvard Medical School

Debora S. Marks[†]

Systems Biology
Harvard Medical School

ABSTRACT

Mosaic nanoparticle vaccines incorporating naturally diverse sarbecovirus receptor binding domains (RBDs) represent a promising approach for pan-coronavirus vaccine design. Mosaic nanoparticles elicit broad, cross-reactive immune responses, likely because antibodies utilize avidity effects to preferentially bind to conserved regions where they can cross-link across neighboring RBDs. However, the diversity in natural RBDs is limited, leading to ‘off-target’ antibodies that do not bind to low-mutability regions. We therefore develop a novel future-proof vaccine design method, building upon a probabilistic generative model of antibody escape, to computationally design RBDs with further diversity. This approach aims to focus antibody responses to regions that are (1) neutralizing, (2) accessible and (3) unlikely to mutate during future viral evolution. The designs will be assessed by immunizing mice and testing the breadth of neutralizability of the sera compared to a nanoparticle composed of naturally diverse strains.

1 INTRODUCTION

The COVID-19 pandemic, particularly the waning immunity from SARS-CoV-2 vaccines, reinforced the necessity for universal vaccines that are broadly protective against diverse members of a viral family. A pan-coronavirus vaccine would not only protect against presently circulating SARS-CoV-2 variants but also against future variants and other species within the family. This includes existing endemic human coronaviruses and potential spill-over events that could seed a new pandemic. This urgent need is underscored by the billions of dollars earmarked for the development of pan-coronavirus vaccines across a range of approaches including mRNA-based, nanoparticle-based, and protein subunit vaccines (Dolgin, 2022; Musunuri et al., 2024; Caradonna & Schmidt, 2021).

Mosaic nanoparticles represent one such promising approach for pan-coronavirus vaccines, as well as other pan-genus or pan-family vaccines for viruses such as influenza. In a nanoparticle, antigenic proteins are attached to a central protein scaffold (Brune & Howarth, 2018). This vaccine modality inherently lends itself to broad protection in two crucial ways. Firstly, in the case of a mosaic nanoparticle, different antigens are presented, from diverse viral species (as opposed to a homotypic nanoparticle where each antigen is the same). Secondly, when adjacent antigens differ, antibodies capable of cross-linking between conserved regions in neighboring proteins, using avidity binding, are elicited (Fig 1A). These conserved epitopes, shared across viral species, are less prone to future mutation and antibodies binding them are more likely to protect against the entire viral family, even beyond strains present on the nanoparticle (Cohen et al., 2021a; 2022). Nanoparticles containing 8 naturally diverse sarbecovirus RBDs elicit broader cross-reactive responses compared to homotypic nanoparticles or convalescent plasma. This broader response offers protection both against other sarbecoviruses and SARS-CoV-2 variants not present on the nanoparticle (Cohen et al., 2021a; 2022; 2024).

^{*}Equal contribution

[†]Correspondence: sgurev@mit.edu, noor_youssef@hms.harvard.edu, debbie@hms.harvard.edu

Despite its potential, the current mosaic nanoparticle presents significant challenges. There is evidence of waning immunity to newer SARS-CoV-2 variants (Cohen et al., 2024). Additionally, the nanoparticle elicits off-target antibodies, many of which are inaccessible in the context of the full Spike protein encountered during an infection, failing to provide effective protection. Epitopes may also be conserved between a pair of viruses on the nanoparticle despite a lack of conservation throughout the sarbecovirus sub-genus, highlighting potential susceptibility to future mutations (Cohen et al., 2021a; 2022).

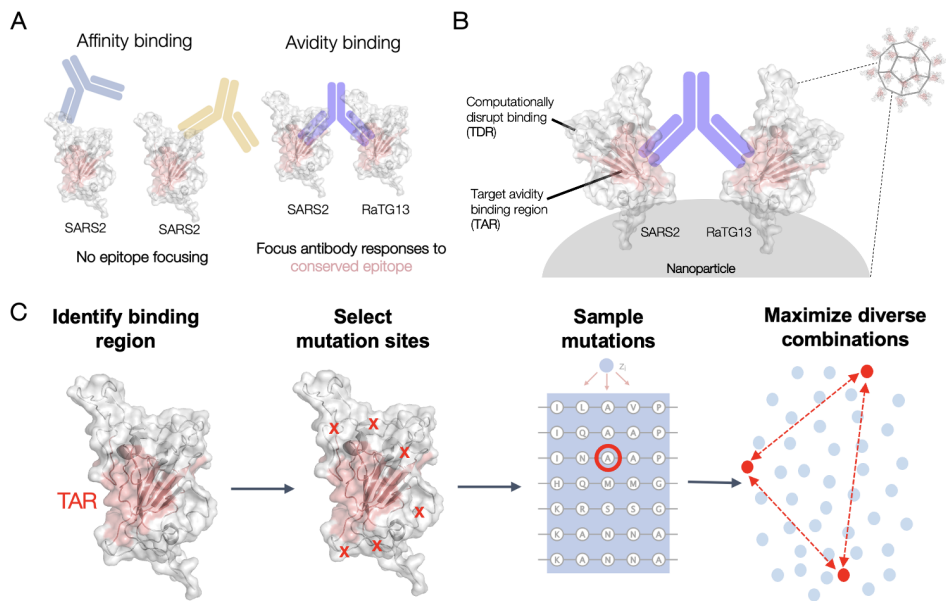


Figure 1: **A.** A single RBD elicits antibodies that bind with affinity to many regions. Conversely, two adjacent, different RBDs focus antibody responses to conserved regions across the RBDs through avidity effects. **B.** RBD antigens were designed to have a Target Avidity Region (TAR) that is conserved, neutralizing, and less likely to mutate during viral evolution. To optimize antibody binding to the TAR we computationally designed mutations elsewhere (in the Target Disruption Region or TDR) to prevent off-target binding. **C.** Steps in nanoparticle design process.

We propose to address these challenges by designing mosaic RBD nanoparticles that focus the antibody response. Our approach involves generating RBDs with greater diversity than natural strains outside a selected conserved, neutralizing region. This is aimed at refining antibody responses to focus on regions that are accessible and less likely to mutate, focusing antibodies beyond what is naturally conserved between any pair of natural strains on the original nanoparticle. Prior approaches to masking undesirable epitopes (Caradonna & Schmidt, 2021), including hyperglycosylation (Thornlow et al., 2021) and random mutation with relatively inert amino acids (Frei et al., 2018), have achieved moderate success for other vaccine modalities.

Given the substantial number of mutations to optimize in a mosaic nanoparticle, we build upon our prior work creating a deep generative model of antibody escape, EVEscape (Thadani et al., 2023; Youssef et al., 2023), to tackle this novel vaccine design approach computationally (Fig 1B). EVEscape, trained on pre-pandemic data, has been shown to be as accurate as high-throughput experimental scans using pandemic antibodies at anticipating SARS-CoV-2 mutations escaping prior immunity (Thadani et al., 2023). EVEscape has also demonstrated remarkable success in designing antigenically diverse SARS-CoV-2 sequences for vaccine testing which remain infectious at distances of up to 46 mutations from the ancestral strain (Youssef et al., 2023). While EVEscape was framed previously as a model of antibody escape from prior infection or vaccination, it can more generally be considered to be a measure of cross-reactivity between antigens—whether antibodies directed against a first antigen can successfully bind a different antigen. In the case of modeling antibody escape in a pandemic, the first strain is the ancestral strain and the second is a future variant. However, it can be applied to cross-reactivity between two antigens on a single nanoparticle. Therefore, EVEscape provides an ideal foundation for an approach focused on generating functional

heterotypic antigens with targeted mutations that eliminate cross-reactivity at any potential off-target epitope.

2 DESIGN METHODOLOGY

Here we address a novel sequence design problem: given a set of naturally diverse antigens in a viral family, identify a region to target antibody binding and generate mutations elsewhere, while retaining antigen functionality. The mutations should minimize the pairwise antibody cross-reactivity (where antibodies can bind with avidity) at any possible off-target epitope. The goal is a group of designs that, when immunizing as a mosaic nanoparticle, generates sera that is more broadly neutralizing against diverse existing and potential future species in a viral family, compared to a mosaic nanoparticle composed of the original, natural strains. We base our designs on the natural mosaic sarbecoviruses used previously: SARS-CoV-2 Beta, RaTG13, Rs4081, SHC014, Pang17, RmYN02, Rf1, and WIV1 (Cohen et al., 2021a; 2022; 2024). Critically, the model is unsupervised and does not rely on prior knowledge of specific epitopes or experimental data.

We develop a multi-step design process (Fig 1C)—rather than a joint optimization that directly results in a final group of designed antigens—due to the complexity of the design objectives. Moreover, with diverse antigens as the key goal, it is important to down-sample from many potential designs to avoid prematurely constraining design choices based on initially selected sequences. This design process builds upon EVEscape, which models antibody escape as the product of three components: (1) mutation maintains fitness, (2) mutation is accessible to antibodies, (3) mutation is sufficiently dissimilar to disrupt antibody binding (Thadani et al., 2023). By virtue of EVEscape’s flexible, modular framework, we repurpose each of its components, separately and in combination, in the relevant steps of our approach.

Identify target avidity binding region We first identify an optimal antibody binding region, a contiguous, accessible region with (1) physicochemical conservation across sarbecoviruses and (2) neutralization potential. Sequence conservation is often measured via Shannon entropy, calculated from the relative frequency of each amino acid in a multiple sequence alignment position. However this conservation ignores physicochemical similarities between some amino acids. We therefore develop a physicochemical-based conservation using the dissimilarity component from EVEscape (Thadani et al., 2023) to calculate a von Neumann Entropy (VNE) (Caffrey et al., 2004). This considers not just the relative frequencies f of each amino acid i , but also a physicochemical amino acid similarity matrix P , such that

$$VNE = \sum_i \lambda_i \log_{20} \lambda_i, \text{ where } \lambda_i \text{ are the eigenvalues of } \text{diag}(f)P.$$

Second, since most neutralizing antibodies for the RBD physically block the host receptor ACE2 from binding, we use distance to the ACE2 binding site as a proxy for neutralization potential. Amongst accessible, non-Spike contacting positions we create a contiguous region by selecting patches of residues to include with a high rank product of neutralization and conservation scores amongst all residues within a radius of 8Å. This resulted in a target avidity binding region (TAR) of 66 residues, approximately one third of the RBD. This selected TAR has several advantages over the regions in each pair of natural sequences where antibodies bind with avidity: the TAR is composed of more non-Spike contacting positions, contains residues that are not just similar across a given pair of natural sequences on the nanoparticle but have physicochemical conservation across the entire sarbecovirus family (which has down to 68% sequence identity), and is selected to ignore decoy epitopes from the natural strains that may be conserved but are of low neutralizing potential.

Select mutation sites In order to minimize antibody cross-reactivity to pairs of antigens outside the TAR, we select combinations of positions to optimally eliminate off-target binding. We begin by identifying a region that due to great evolutionary diversity, is already sufficiently diverse to mitigate off-target binding, minimizing the number of mutations we must make. We call this the evolutionary target disruption region (TDR) or eTDR. We then select positions to mutate in the remaining region that is more likely to have off-target binding, which we call the designed TDR or dTDR.

First, we identify the already sufficiently dissimilar evolutionary TDR through the same patch-based approach described above, but from physicochemical-conservation of the natural strains on the

nanoparticle alone. This creates an eTDR of 66 residues, with 87 residues remaining in neither the eTDR or TAR. Second, we identify positions to mutate within the designed TDR. In order to disrupt each potential epitope, we wish to select spatially dispersed sites to mutate. To do so, we use k-means clustering on the Cartesian coordinates of the alpha carbon of each residue – such that each cluster represents a potential off-target epitope to disrupt. We find that using 8 centroids generates clusters in which the maximum distance between residues is less than $\sim 25\text{\AA}$, the typical diameter of conformational B cell epitopes (Sun et al., 2011). Within each of the 8 spatially distributed clusters, we select one position to mutate by centroid proximity. We rerun the k-means algorithm with 5000 random initializations to generate many sets of mutation sites to use to create designs.

Sample mutations After selecting sites to mutate, we create designed sequences by making substitutions with a probabilistic generative model, EVE (Frazer et al., 2021). EVE is a deep variational autoencoder trained on a multiple sequence alignment of evolutionarily related protein sequences, in this case sarbecovirus RBD sequences aligned to each of the base natural sequence antigens, which learns the constraints that underpin structure and function. We explore the impact of different fitness models in Fig S4. We use three sampling strategies (fittest, sampled, and random) for generating designs to maximize diversity and full-sequence fitness. The first strategy involves creating designs where we select the fittest possible single mutation at each mutational site ($\mathbb{S}_{\text{fittest}}$). The second strategy involves selecting a random mutation at each mutational site without considering fitness ($\mathbb{S}_{\text{random}}$). Finally, for the third strategy, we construct a probability vector normalizing the fitness scores at a site for all mutations with scores above the model’s median single mutation score, which we then select from via multinomial sampling ($\mathbb{S}_{\text{sampled}}$). We remove any designs where no mutation in a cluster is fit. Finally, we down-sample our designs generated with the three strategies to the set with full-sequence EVE score (moving beyond single mutations to scores of 8 mutations in combination) higher than the median score of $\mathbb{S}_{\text{fittest}}$ and maximum score of $\mathbb{S}_{\text{random}}$.

Optimize combinations of designs The final step is selecting groups of diverse sequences to place on the nanoparticle together from the large set of sampled designs. We wish to choose pairs of design sequences $\mathbf{s}_j, \mathbf{s}_k$ that maximize $d[\mathbf{s}_j, \mathbf{s}_k]$, a joint accessibility and physicochemical dissimilarity product across all off-target epitopes \mathbb{E} :

$$d[\mathbf{s}_j, \mathbf{s}_k] = \sum_{e \in \mathbb{E}} \text{acc}[e] \cdot \text{dist}[e, \mathbf{s}_j, \mathbf{s}_k]$$

$$\text{acc}[e] = \frac{1}{|e|} \sum_{r \in e} \mathbf{a}[r]$$

$$\text{dist}[e, \mathbf{s}_j, \mathbf{s}_k] = \frac{1}{|e|} \sum_{r \in e} \mathbf{a}[r] \cdot \mathbf{P}'[\mathbf{s}_j[r], \mathbf{s}_k[r]]$$

where $\text{acc}[e]$ is the accessibility of epitope e , $\text{dist}[e, \mathbf{s}_j, \mathbf{s}_k]$ is the epitope difference between \mathbf{s}_j and \mathbf{s}_k , $\mathbf{a}[r]$ is the accessibility of residue r , and $\mathbf{P}'[\mathbf{s}_j[r], \mathbf{s}_k[r]]$ is the physicochemical dissimilarity between amino acids in position r . Each potential off-target epitope corresponds to a patch around each residue in the target disruption region—where a patch includes all residues within $\sim 10\text{\AA}$. Each epitope difference is weighted by the average accessibility of that epitope, as a measure of the relative importance of each epitope. We use the accessibility and physicochemical dissimilarity components from EVEscape (Thadani et al., 2023), having already selected putatively functional sequences using EVEscape’s generative fitness model in the prior step. To select a mosaic group of designs, we therefore create an accessibility-dissimilarity distance matrix between all pairs of designs. We then use a greedy algorithm to find a solution to the maximum dispersion problem, which seeks to maximize the sum of pairwise distances between selected designs. We start by choosing a random sequence \mathbf{s} to add to the mosaic nanoparticle \mathbb{N} . We continue to add to \mathbb{N} by iteratively selecting the next design which maximizes the minimum distance to all designs already in \mathbb{N} (that uses a different target virus from those already chosen). We repeated this process with different starting sequences and ultimately selected 10 groups of designs \mathbb{N} for experimental characterization for their neutralization potential. Fig S1 showcases a representative group of nanoparticle designs (diverse, natural sarbecoviruses each with 8 different mutations) mapped onto a single RBD structure to see the different mutations sites across viruses.

3 RESULTS

While the ultimate evaluation will be our success at achieving a broadly neutralizing response from immunizing mice with our designed mosaic nanoparticle, there are several criteria we can evaluate computationally. We require each antigen to be fit enough to express and fold into a relatively proper conformation (note that the RBDs do not need to be functional enough to be infective or bind ACE2). Additionally, by spatially dispersing the mutations and mutating an accessible surface (rather than in the core of a protein), we make it more likely for the protein to fold, as clusters of nearby mutations are more likely to experience negative epistasis and buried mutations are under more structural constraints. We therefore demonstrate the ability of our mutation fitness model (EVE) to predict expression from deep mutational scans (Starr et al., 2020) (Table 1), as shown previously (Frazer et al., 2021; Thadani et al., 2023; Youssef et al., 2023). Fig S2 then shows that selected mutations in designs are above a 50% threshold of all single mutations using the EVE fitness model for each target virus. The fitness score of the full RBD for each design is also higher than all generated antigens with random mutations at selected sites. (Figs 2, S3). We also explore the impact of different fitness models on designs (Fig S4). Moreover, we analyze the spatial separation of the mutations and calculated joint accessibility and physicochemical distances, in order to maximize the impact on each potential off-target epitope (Figs S1, S5).




Figure 2: **A.** For each of the 8 natural RBDs in the mosaic-8b nanoparticle vaccine, we generate 15k sequences (5k by random amino acids, 5k by fittest amino acid, 5k by sampling proportional to fitness). Full-RBD fitness scores for 10 selected designs (dashed lines) on the RaTG13 sequence, relative to distributions S_{random} , S_{fittest} , and S_{sampled} . Note that some lines are overlapping. Distributions for remaining RBDs in Fig S3.

4 DISCUSSION

We present a novel future-proof vaccine design method using a deep generative model of antibody escape, EVEscape. Generative models are often divorced from real world applications. By wrapping our generative model into a larger, multi-step optimization process, which we then use to design a mosaic nanoparticle vaccine for experimental testing, we bridge this gap. Given EVEscape’s modular components, we are able to build our design process on its foundation.

We will perform experimental testing by immunizing mice with our designed mosaic nanoparticle, evaluating our vaccine based on its neutralization breadth, against both diverse current and future strains. While we focus on the sarbecovirus sub-genus here, we will later explore our ability to achieve broader protection—for instance of the betacoronavirus genus or indeed of the entire coronavirus family. We will evaluate the breadth of protection relative to a nanoparticle composed of natural strains. To facilitate this comparison, we began our optimization from the same set of natural strains as in the vaccine developed by Cohen et al. (2021a; 2022). Once the success of our approach has been affirmed, an additional step in the design process could be incorporated to select an initial set of diverse viral species as the base of the designs. Moreover, given our initial experimental results, we plan to iterate on our entire design process.

Thus far, our focus has been only on protection by neutralizing antibodies (both in our design objective and evaluation criteria), rather than protection from other sources, such as non-neutralizing

antibody (e.g., ADCC) or T cell responses. Current vaccine practices take this same narrow focus—likely to their detriment. This may be beginning to change, especially for the pursuit of universal vaccines. By virtue of our multi-step process, it is simple for us to likewise incorporate additional design objectives, which we would then match with corresponding evaluation criteria. For example, our current design protocol does not assess the likelihood of triggering an autoimmune response against human proteins. A deimmunization assessment can be easily incorporated into our design protocol to ensure that none of the designed RBDs resemble human proteins.

Remarkably, the field of vaccine design has not yet utilized machine learning efforts to model antibody escape and cross-reactivity. Our approach is readily generalizable – both to other viral families (e.g., influenza nanoparticles already exist and could be similarity optimized (Cohen et al., 2021b)) and to other approaches for choosing nanoparticle antigens (including the full Spike, sub-dominant regions, or even other antigenic proteins like nucleocapsid). Furthermore, this process of optimizing nanoparticles or comparable heteromultimeric vaccination strategies are just the beginning of methods in which EVEscape or other generative models could benefit vaccine design. This is not even considering the new paradigm in generative models for vaccine testing efforts as outlined by (Youssef et al., 2023). Indeed, we will also use EVEscape to design future strains to test protection, beyond evaluating against existing coronavirus diversity, as was previously the only option. The exciting possibilities presented by generative models, including EVEscape, hold immense potential for transforming vaccine design practices, promising far-reaching impacts across applications.

ACKNOWLEDGMENTS

The authors thank Nicole Thadani, Alex Cohen, Luis Caldera, and Pamela Bjorkman, and members of the Marks lab. This work was supported by the Coalition for Epidemic Preparedness Innovations (CEPI).

REFERENCES

- Muriel Aguilar-Bretones, Ron AM Fouchier, Marion PG Koopmans, Gijsbert P van Nierop, et al. Impact of antigenic evolution and original antigenic sin on sars-cov-2 immunity. *The Journal of clinical investigation*, 133(1), 2023.
- Karl D. Brune and Mark Howarth. New routes and opportunities for modular construction of particulate vaccines: Stick, click, and glue. *Frontiers in Immunology*, 2018.
- Daniel R Caffrey, Shyamal Somaroo, Jason D Hughes, Julian Mintseris, and Enoch S Huang. Are protein–protein interfaces more conserved in sequence than the rest of the protein surface? *Protein Science*, 13(1):190–202, 2004.
- Timothy M Caradonna and Aaron G Schmidt. Protein engineering strategies for rational immunogen design. *npj Vaccines*, 6(1):154, 2021.
- Alexander A. Cohen, Priyanthi N. P. Gnanapragasam, Yu E. Lee, Pauline R. Hoffman, Susan Ou, Leesa M. Kakutani, Jennifer R. Keefe, et al. Mosaic nanoparticles elicit cross-reactive immune responses to zoonotic coronaviruses in mice. *Science*, 2021a.
- Alexander A Cohen, Zhi Yang, Priyanthi NP Gnanapragasam, Susan Ou, Kim-Marie A Dam, Haoqing Wang, and Pamela J Bjorkman. Construction, characterization, and immunization of nanoparticles that display a diverse array of influenza ha trimers. *PLoS One*, 16(3):e0247963, 2021b.
- Alexander A Cohen, Neeltje van Doremalen, Allison J Greaney, Hanne Andersen, Ankur Sharma, Tyler N Starr, Jennifer R Keefe, Chengcheng Fan, Jonathan E Schulz, Priyanthi NP Gnanapragasam, et al. Mosaic rbd nanoparticles protect against challenge by diverse sarbecoviruses in animal models. *Science*, 377(6606):eabq0839, 2022.
- Alexander A. Cohen, Jennifer R. Keefe, Ariën Schiepers, Sandra E. Dross, Allison J. Greaney, Annie V. Rorick, Han Gao, Priyanthi N.P. Gnanapragasam, Chengcheng Fan, Jr. Anthony P. West, Arlene I. Ramsingh, Jesse H. Erasmus, Janice D. Pata, Hiromi Muramatsu, Norbert Pardi,

- Paulo J.C. Lin, Scott Baxter, Rita Cruz, Martina Quintanar-Audelo, Ellis Robb, Cristina Serrano-Amatriain, Leonardo Magneschi, Ian G. Fotheringham, Deborah H. Fuller, Gabriel D. Victora, and Pamela J. Bjorkman. Mosaic sarbecovirus vaccination elicits cross-reactive responses in pre-immunized animals. *bioRxiv*, 2024. doi: 10.1101/2024.02.08.576722.
- Elie Dolgin. Pan-coronavirus vaccine pipeline takes form. *Nature Reviews Drug Discovery*, 21(5): 324–26, 2022.
- Robert C. Edgar. Muscle v5 Enables Improved Estimates of Phylogenetic Tree Confidence by Ensemble Bootstrapping. *BioRxiv*, 2021.
- David Eisenberg, Robert M Weiss, and Thomas C Terwilliger. The hydrophobic moment detects periodicity in protein hydrophobicity. *Proceedings of the National Academy of Sciences*, 81: 140–144, 1984.
- Jonathan Frazer, Pascal Notin, Mafalda Dias, Aidan Gomez, Joseph K. Min, Kelly Brock, Yarin Gal, and Debora S. Marks. Disease variant prediction with deep generative models of evolutionary data. *Nature*, 2021.
- Julia C Frei, Ariel S Wirchnianski, Jennifer Govero, Olivia Vergnolle, Kimberly A Dowd, Theodore C Pierson, Margaret Kielian, Mark E Girvin, Michael S Diamond, and Jonathan R Lai. Engineered dengue virus domain iii proteins elicit cross-neutralizing antibody responses in mice. *Journal of virology*, 92(18):10–1128, 2018.
- Allison J Greaney, Tyler N Starr, and Jesse D Bloom. An antibody-escape calculator for mutations to the sars-cov-2 receptor-binding domain. *BioRxiv*, 2021.
- Greg Hamerly and Charles Elkan. “alternatives to the k-means algorithm that find better clusterings. In *Proceedings of the Eleventh International Conference on Information and Knowledge Management*, 2002.
- Steven Henikoff and Jorja G Henikoff. Amino acid substitution matrices from protein blocks. *Proceedings of the National Academy of Sciences*, 89(22):10915–10919, 1992.
- Thomas A Hopf, John B Ingraham, Frank J Poelwijk, Charlotta PI Schärfe, Michael Springer, Chris Sander, and Debora S Marks. Mutation effects predicted from sequence co-variation. *Nature biotechnology*, 35(2):128–135, 2017.
- L. Steven Johnson, Sean R. Eddy, and Elon Portugaly. Hidden markov model speed heuristic and iterative HMM search procedure. *BMC Bioinformatics*, 2010.
- Joshua Meier, Roshan Rao, Robert Verkuil, Jason Liu, Tom Sercu, and Alex Rives. Language models enable zero-shot prediction of the effects of mutations on protein function. *Advances in Neural Information Processing Systems*, 34:29287–29303, 2021.
- Sriharshita Musunuri, Payton AB Weidenbacher, and Peter S Kim. Bringing immunofocusing into focus. *npj Vaccines*, 9(1):11, 2024.
- Erik Nijkamp, Jeffrey A Ruffolo, Eli N Weinstein, Nikhil Naik, and Ali Madani. Progen2: exploring the boundaries of protein language models. *Cell Systems*, 14(11):968–978, 2023.
- Pascal Notin, Lood Van Niekerk, Aaron W. Kollasch, Daniel Ritter, Yarin Gal, and Debora S. Marks. Trancepve: Combining family-specific and family-agnostic models of protein sequences for improved fitness prediction. 2022.
- Yuelong Shu and John McCauley. Gisaid: Global initiative on sharing all influenza data—from vision to reality. *Eurosurveillance*, 22(13):30494, 2017.
- Tyler N. Starr, Allison J. Greaney, Sarah K. Hilton, Daniel Ellis, Katharine H. D. Crawford, Adam S. Dingens, Mary Jane Navarro, et al. Deep mutational scanning of SARS-cov-2 receptor binding domain reveals constraints on folding and ACE2 binding. *Cell*, 2020.
- Jing Sun, Tianlei Xu, Shuning Wang, Guoqing Li, and Zhiwei Cao. Does difference exist between epitope and non-epitope residues? analysis of the physicochemical and structural properties on conformational epitopes from b-cell protein antigens. *Immunome Research*, 2011.

Baris E. Suzek, Hongzhan Huang, Peter McGarvey, Raja Mazumder, and Cathy H. Wu. Uniref: Comprehensive and non-redundant UniProt reference clusters. *Bioinformatics*, 2007.

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*, 2023.

Dana N Thornlow, Andrew N Macintyre, Thomas H Oguin 3rd, Amelia B Karlsson, Erica L Stover, Heather E Lynch, Gregory D Sempowski, and Aaron G Schmidt. Altering the immunogenicity of hemagglutinin immunogens by hyperglycosylation and disulfide stabilization. *Frontiers in Immunology*, 12:737973, 2021.


Noor Youssef, Sarah Gurev, Fadi Ghantous, Kelly Brock, Javier Jaimes, Nicole N Thadani, Ann Dauphin, Amy C Sherman, Leonid Yurkovetskiy, Daria Soto, et al. Protein design for evaluating vaccines against future viral variation. *bioRxiv*, 2023.

A APPENDIX

A.1 GLOSSARY

ACE2	Angiotensin-converting enzyme 2	Human host receptor for SARS-CoV-2 entry
dTDR	Design TDR	Region to add mutations to disrupt off-target antibody avidity binding
eTDR	Evolutionary TDR	Region with sufficient evolutionary diversity to already disrupt off-target antibody avidity binding
RBD	Receptor Binding Domain	Domain of SARS-CoV-2 and other sarbecovirus Spike proteins used for nanoparticle vaccine - an immunodominant, neutralizing region
$\mathbb{S}_{\text{fittest}}$	Fittest sequences	Sequence set with fittest mutation at chosen sites
$\mathbb{S}_{\text{random}}$	Random sequences	Sequence set with random mutation at chosen sites
$\mathbb{S}_{\text{sampled}}$	Sampled sequences	Sequence set with mutations sampled by fitness scores at chosen sites
TAR	Target Avidity Region	Region to target antibody avidity binding towards, where antibodies are likely to be protective across current and future sarbecoviruses
TDR	Target Disruption Region	Region to disrupt off-target antibody avidity binding - where high mutability, inaccessibility, or lack of neutralization potential would lead to less protective antibodies in the future
VNE	von Neumann Entropy	Physicochemical-based amino acid conservation

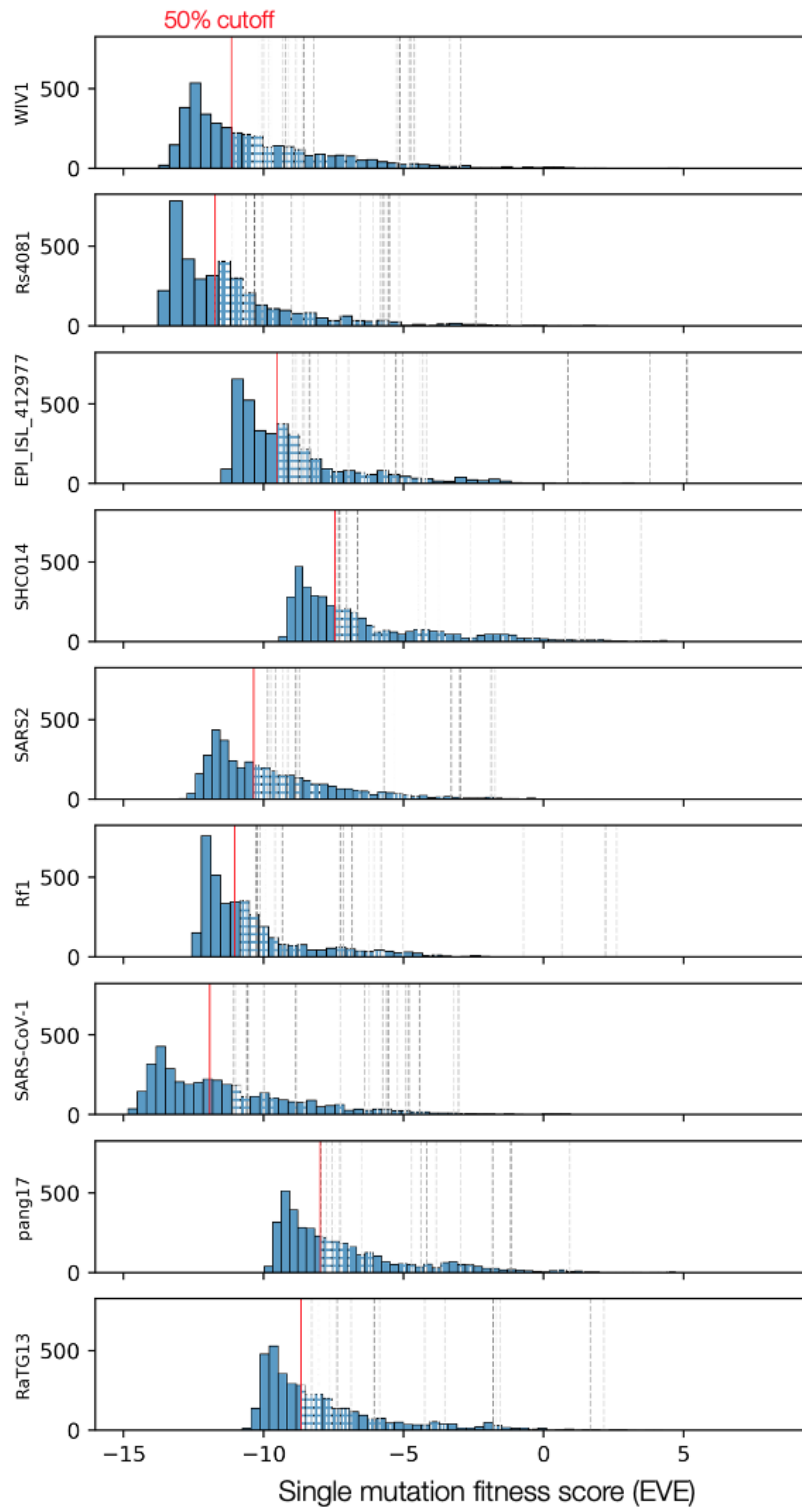
A.2 SUPPLEMENTARY FIGURES



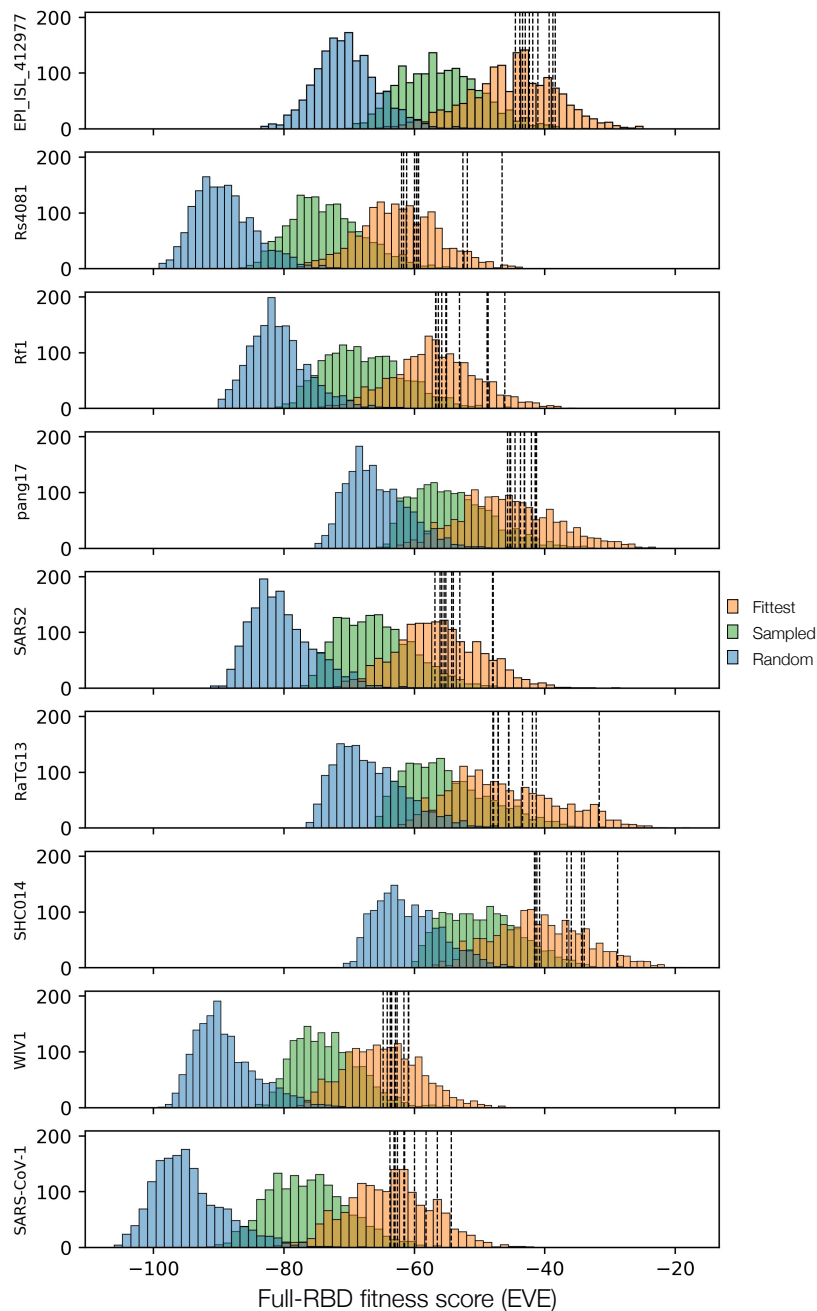
Supplementary Figure 1: Representative group of nanoparticle designs (diverse, natural sarbecoviruses each with 8 different mutations) mapped onto a single RBD structure to see the different mutations sites across viruses. Within target viruses mutations are spatially distributed, and different sets of mutation sites are chosen for different target viruses. Note, a single color is shown even for overlapping positions.

Table 1: Spearman correlations between mutation effect predictions using different fitness models and experimental values from an RBD expression deep mutational scanning experiment (Starr et al., 2020).

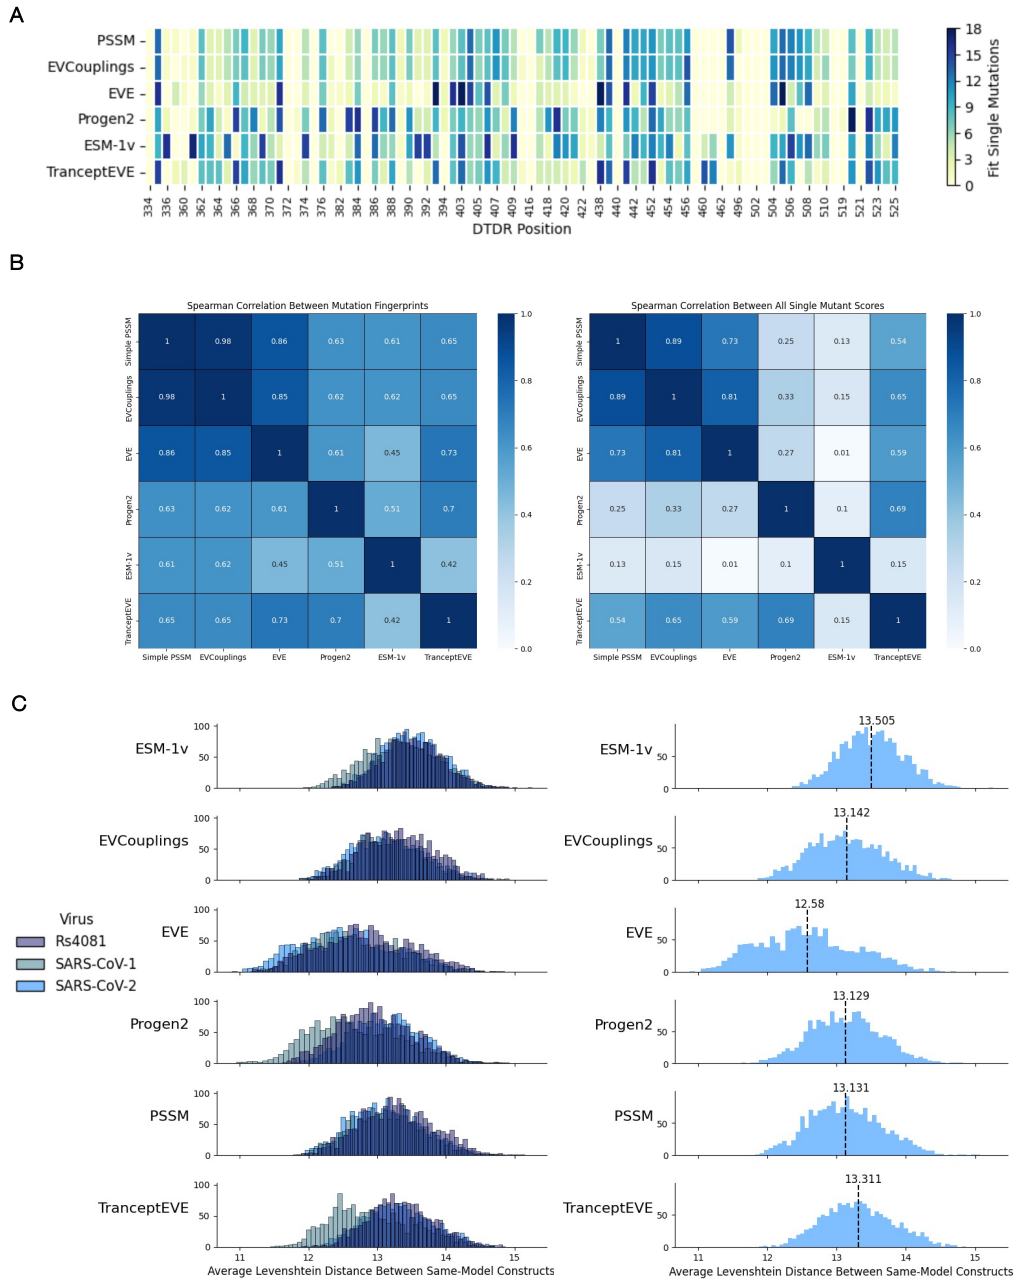
Model	Expression correlation (ρ)
ESM-1v	-0.01
Progen2	0.36
PSSM	0.37
EVCouplings	0.45
EVE	0.51
TranceptEVE	0.52




Supplementary Figure 2: Distribution of single mutation fitness scores with EVE model (Red line indicates median score). Dashed lines show scores of selected mutations across 10 selected nanoparticle groups across the different target viruses (darker lines indicate a mutation that is in more than one selected nanoparticle group).



Supplementary Figure 3: For each of the natural RBDs in the mosaic-8b nanoparticle vaccine, we generate 15k sequences (5k by randomly sampling amino acids, 5k by assigning fittest amino acid, 5k by sampling proportional to fitness). Full-RBD fitness scores for 10 selected designs (dashed lines) on each of the nine target sequences, relative to distributions of designs $\mathbb{S}_{\text{fittest}}$, $\mathbb{S}_{\text{sampled}}$, and $\mathbb{S}_{\text{random}}$. Note that some lines are overlapping.



Supplementary Figure 4: **A**. EVE favors specific positions more than other alignment-based sequence fitness models. This heatmap visualizes the number of possible mutations which were observed at each position within the SARS-CoV-2 design set for a given model. The theoretical maximum number of mutations that could be observed at a specific position is 19, since we do not count the wild-type amino acid as a mutation. Position-specific scoring matrix (PSSM), EVCouplings (Hopf et al., 2017), and TranceptEVE (Notin et al., 2022) show relatively weak preference for specific mutational sites, while EVE (Frazer et al., 2021) shows strong preferences for a few mutational sites. Progen2 (Nijkamp et al., 2023) and Esm-1v (Meier et al., 2021) also show moderate preferences for some dTDR sites. **B**. Alignment-based models generate more similar mutational fingerprints than language models. Spearman correlations between mutational fingerprint in set of SARS-CoV-2 designs generated by each model (left) and SARS-CoV-2 single mutation scores (right). **C**. EVE generates less diverse designs than other models. Distributions of average pairwise Levenshtein distances for three target viruses (left) and SARS-CoV-2 only (right). The dotted line shows the mean of all average pairwise Levenshtein distances between designs.



Supplementary Figure 5: Joint accessibility and physicochemical dissimilarity score distances between all pairs of designs in representative nanoparticle (dashed lines) and across all generated designs (blue distribution).

A.3 EXTENDED METHODS

A.3.1 DATA ACQUISITION

Viral Background Sequences We chose to design two sets of vaccine designs: one set designing mutations on the 8 sarbecoviruses used by Cohen et al. (2021b) in their original mosaic-8b nanoparticle (SARS-CoV-2 Beta, RaTG13, Rs4081, SHC014, Pang17, RmYN02, Rf1, WIV1), and another set containing designed sequences on SARS-CoV-1 instead of SARS-CoV-2. Altogether, we call these the 9 target viruses. We decided to create groups of designs containing SARS-CoV-1 to reduce the chance of generating recall antibodies in individuals with previous SARS-CoV-2 infection or mRNA vaccination, which may diminish preferred antibody responses (Cohen et al., 2024; Aguilar-Bretones et al., 2023). Note, for the purposes of clarity of the main text we focus there only on the original 8 sequences (containing SARS-CoV-2).

Protein Structures We selected a structure from the Protein Data Bank (7C01), which shows the SARS-CoV-2 Spike RBD in complex with a neutralizing antibody (a relevant orientation for vaccine design), and provides full coverage of the desired RBD residues. This structure is used for all distance and accessibility calculations within the RBD. For calculating the distance of each residue to the ACE2 receptor we selected 6M0J, where chain A is the ACE2 receptor and chain B is the RBD.

Multiple Sequence Alignments We use JackHMMER and super5 to construct multiple sequence alignments (MSAs) for each of our 9 target viruses (Johnson et al., 2010; Edgar, 2021). For each virus, we create a alignment using 5 JackHMMER iterations to search against the Uniref100 dataset (Suzek et al., 2007) with bitscore 0.1 normalized to the full sequence length. To avoid bias due to large numbers of SARS-CoV-2 pandemic sequences in UniRef, we remove all SARS-CoV-2 sequences. We concatenate the 9 alignments to create a combined alignment, so all sequences identified by JackHMMer for each target virus will be considered for all target viruses. We use super5 to realign the combined alignment to each target virus, creating 9 final individual sarbecovirus alignments. Lastly, we filter the individual alignments by removing any sequences with less than 50% match to the target virus and removing any residue positions with more than 70% gaps.

A.3.2 EVESCAPE MODEL COMPONENTS

Sequence Fitness To estimate the fitness effect of substitution mutations, and score full-RBD sequences, we train an EVE model Frazer et al. (2021) on the final sarbecovirus MSA for each target virus. EVE is a Bayesian variational autoencoder (VAE) capable of capturing higher-order interactions across sites that learns constraints underpinning structure and function from evolutionarily-related sequences. The fitness of a protein sequence is quantified using the log likelihood ratio of the mutated sequence over that of the wildtype sequence. Thadani et al. (2023) demonstrated that EVE (and EVCouplings) is highly successful at learning constraints important to viral fitness Starr et al. (2020) and antibody escape (Greaney et al., 2021), especially frequent pandemic mutations (Shu & McCauley, 2017). Moreover, Youssef et al. (2023) showed that EVEscape (and EVE) can be used to design infectious pseudovirus Spike proteins with novel combinations of up to 10 mutations from existing Variant of Concern strains, and up to 46 mutations from the ancestral Wuhan strain.

Notably, even when training on an alignment of natural RBD sequences (not the Spike), the EVE model learns constraints underlying contacts with the full Spike. These constraints are not present in deep mutational scanning experiments of RBD expression alone (Starr et al., 2020). This results in some discrepancies between model and experiment, as seen in Table 1, where the models alone capture structural constraints that are meaningful in real life infections with a full Spike protein (Thadani et al., 2023). However, these are in fact unnecessary constraints in the case of RBD nanoparticle design—especially as we have a goal of mitigating off-target antibody binding in full Spike contacts with the RBD.

To run EVE, we use a theta re-weighting value of 0.01 to cluster similar sequences. We use these 9 EVE models to score all possible single mutations for the RBD of each target virus. Based on these fitness scores, we label each single mutant as above or below a 50th percentile fitness score cutoff. We use a 50% cutoff as a strict cutoff based on previous deep mutational scans of the SARS-CoV-2 RBD expression, which indicate that $\sim 30\%$ of single mutations result in a non-functional RBD

(Starr et al., 2020). We use the fitness scores to select mutation sites with at least one fit mutation (section A.3.4), and then to sample fit mutations in our design process that result in functional RBDs with 8 mutations from each target virus (section A.3.5). We also score full-RBDs with the EVCouplings fitness model (Hopf et al., 2017) to further validate the fitness of our choices of designed sequences in section A.3.5.

In the Appendix A.2 we explore the use of alternate fitness models. For alignment-based models, we consider both EVCouplings (Hopf et al., 2017), which considers pairwise relationships between residues, as well as a simple position-specific scoring matrix (PSSM) which considers each column of the MSA in isolation. We also consider Progen2 (Nijkamp et al., 2023) and ESM-1v (Meier et al., 2021), two state-of-the-art transformer-based protein language models which use natural language processing techniques to predict the likelihood of observing a given protein sequence given the lexicon of existing protein sequences. However, viral sequences are relatively underrepresented in the protein sequence universe, especially as only sequences from Uniref90 (which doesn't consider separately sequences that are more than 90% identical, as is true for much of viral variation) are included in training. Finally, we include TranceptEVE (Notin et al., 2022), a hybrid model which combines language modeling with alignment-based modeling. By repeating our design strategy with multiple diverse models, we gain insights into the quality of the designs generated by each model, and the impact of using each model on designed mutations.

Antibody Accessibility The antibody accessibility of each position is computed from its negative weighted residue-contact number from an RBD structure, capturing both protrusion from the core structure and conformational flexibility. Thadani et al. (2023) demonstrated that weighted contact number identifies potential antibody binding sites without prior knowledge of B cell epitopes, and is highly predictive of antibody escape Greaney et al. (2021) and pandemic mutations (Shu & McCauley, 2017). We use the accessibility score in section A.3.6 to select combinations of designs with cross-reactivity at each off-target epitope optimally blocked—with residues and epitopes weighted by their antibody accessibility.

Physicochemical Dissimilarity The dissimilarity between two amino acids is computed using differences in hydrophobicity (Eisenberg et al., 1984) and charge, chosen as physicochemical properties impacting protein-protein interactions. Thadani et al. (2023) demonstrated that this physicochemical score predicts which amino acids within a known escape site actually result in antibody escape (Greaney et al., 2021). We use this physicochemical amino acid similarity matrix to identify a physicochemically conserved region in section A.3.3 and select optimal combinations of designs that eliminate cross-reactivity (as determined by their physicochemical dissimilarity) at off-target epitopes in section A.3.6.

A.3.3 IDENTIFY TARGET AVIDITY BINDING REGION

We define the Spike receptor binding domain (RBD) as containing residues 334-525 of the SARS-CoV-2 Spike protein, following Cohen et al. (2021a). We divide the RBD into two main regions: the Target Avidity Region (TAR) and the Target Disruption Region (TDR). The TAR contains residues which are likely to elicit neutralizing antibodies, but which are conserved throughout sarbecoviruses (and so are unlikely to mutate rapidly) and are accessible in a full Spike protein during an infection (not in an RBD position that is in contact with the rest of the Spike protein). The TDR contains all residues outside the TAR. Our goal is to induce point mutations throughout the TDR, disrupting antibody binding and forcing antibodies to bind to the TAR. We anticipate that the immune system will preferentially elicit antibodies that can cross-link two TARs on adjacent, different RBDs on the nanoparticle due to avidity effects. This will result in increased production of broadly neutralizing sarbecovirus antibodies.

To identify TAR residues, we first score all RBD residues based on two metrics: ACE2 binding proximity and physicochemical-based conservation. We take a rank product of ACE2 binding proximity and physicochemical conservation to generate an overall neutralization-conservation (NC) score for each residue, which prioritizes residues that are both highly conserved and proximal to the ACE2 binding site (e.g., binding likely neutralizing antibodies). To generate a continuous surface for antibodies to bind within the TAR, we use a patch-based approach to select residues to include in the TAR. First, we define a patch around each non-Spike contacting RBD residue containing all residues with side chain centers of mass within 8\AA . We calculate the average and standard deviation of NC

scores for each patch and select all patches within the top 30% of average NC scores and the bottom 30% of standard deviation NC scores across patches. This group of patches constitutes the initial TAR. To improve continuity of potential antibody binding sites, we expand the TAR to include any residues within 5Å of at least two positions in the initial TAR. As a final step to remove gaps, we include any one or two linearly adjacent positions which are linearly surrounded by TAR residues (in the primary sequence). This results in a final TAR of 66 residues, which is about one-third of the entire RBD.

ACE2 binding proximity As a proxy for the neutralizing potential of antibodies binding to a given residue, we calculate the distance to ACE2 for each residue, as most RBD neutralizing antibodies inhibit viral binding to ACE2, thus blocking viral entry. For ACE2 binding proximity, we calculate the Euclidean distance of each residue’s alpha carbon atom to the center of the ACE2 binding site. We define the ACE2 binding site as the center of mass of all residues which have at least one atom within 3.5Å of at least one atom on the ACE2 receptor.

Physicochemical-based conservation We calculate a conservation score for each residue, to identify a conserved region that is therefore less likely to mutate in the future as well as that will permit binding with avidity across target viruses. We found the physicochemical-based conservation score for each residue in the RBD using Von Neumann Entropy on our MSA for each target virus.

Conservation is often measured from multiple sequence alignments via Shannon Entropy

$$SE = \sum_i f_i \log_{20} f_i$$

using the relative frequency f of each of the 20 amino acids i at a given site in the alignment. This uses logarithm base 20 to bound values between zero and one. The frequencies from the alignment are first weighted (by the inverse of the number of sequences in a sequence cluster) if using a large alignment where there may be sampling bias, as is the case for the sarbecovirus alignments. The frequencies are calculated without including gaps.

However, Shannon Entropy ignores that distinct amino acids can have different levels of similarity, and should not be considered to be completely orthogonal. In order to consider antibody binding similarity we use EVEscape’s charge-hydrophobicity score (Thadani et al., 2023) that incorporates physicochemical properties known for impacting protein-protein interactions. The charge-hydrophobicity metric is normalized from 0 to 0.8, with 0 being less similar and 0.8 being more similar, and wildtype residues are given a score of 1.

For Von Neumann Entropy (VNE), the Shannon Entropy equation can be modified by matrix multiplying the diagonal matrix of the frequencies from the alignment with a similarity matrix and doing the entropy calculation instead on the eigenvalues of this matrix, which allows for an importance weighting based on physicochemical similarity (Caffrey et al., 2004). While Caffrey et al. (2004) uses BLOSUM50 (Henikoff & Henikoff, 1992) for a similarity matrix, Thadani et al. (2023) found that the charge-hydrophobicity scores is more predictive of antibody cross-reactivity (Greaney et al., 2021). The final physicochemical-based conservation score at a given site is $1 - VNE$, which is then multiplied by the fraction of non-gapped residues at that position.

A.3.4 SELECT MUTATION SITES

To identify Target Disruption Region (TDR) residues wherein mutations will be made, we first consider all RBD residues not in the TAR. From this group of residues, we subdivide the TDR into the Evolutionary Target Disruption Region (eTDR) and the Designed Target Disruption Region (dTDR). The eTDR contains regions of the RBD which already exhibit sufficient evolutionary variation to prevent binding of bivalent antibodies to adjacent RBDs from different viral species. We use the same patch-based approach described above for the TAR to identify the eTDR, but when calculating physicochemical conservation, we consider only variation within the 9 target viruses, rather than broader variation across the sarbecoviruses. Also, we do not consider ACE2 binding proximity when designating the eTDR. We consider all patches in the bottom 10% of physicochemical-based conservation scores, as well as residues within 5 Å of at least two positions in these patches. This results in an eTDR of 66 residues. The remaining 87 RBD residues constitute the dTDR, the region where we will introduce point mutations to deflect off-target antibody binding. While the region

definitions are created using the SARS-CoV-2 RBD numbering, we map each region across to the appropriate residue numbering system of each target virus by using our multiple sequence alignment of target viruses.

After identifying the dTDR residues, our next objective is to select combinations of these residues to mutate. For any given designed RBD, we would like to introduce enough mutations to disrupt antibody binding throughout each potential off-target epitope within the dTDR. To identify the optimum number of mutations to make, we consider the alpha carbon atoms of residues in the dTDR as a set of coordinates in three-dimensional space. We use the SARS-CoV-2 RBD structure as a proxy structure for all target viruses, since there is a high degree of structural similarity between target RBDs across target sarbecoviruses. We perform k-means clustering on the set of dTDR coordinates for different numbers of centroids between 4 and 16, using Elkan’s algorithm (Hamerly & Elkan, 2002). We find that using 8 centroids generates clusters in which the maximum distance between residues is no larger than $\sim 25\text{\AA}$, which is the average diameter of conformational B cell epitopes (Sun et al., 2011). Therefore, we infer that 8 mutations is the smallest number of mutations which can effectively disrupt antibody binding across each potential off-target epitope in the dTDR.

In order to optimize diversity, we wish to create many potential RBD designs for each target virus, each with 8 mutational sites per RBD, by using k-means clustering on the set of dTDR alpha carbon Euclidean coordinates with 8 centroids. We run the k-means algorithm with 5000 different random initializations for each target virus to generate sufficient diversity. Within each set of 8 clusters, we select one residue position per cluster as a mutation site. To select a cluster-wise mutation site, we iterate through the cluster’s residue positions in order of centroid proximity, selecting the first position where at least one single mutant at this position scores above the 50% EVE fitness threshold for the target virus in question. If no residue positions in a given cluster satisfy this constraint, we do not select any residue position for that cluster. After repeating this process for all centroid initializations, we end up with 5000 designed mutation site sets for each of our 9 target viruses. However, the majority of these sets do not contain 8 distinct mutation sites, since many sets contain at least one low-fitness cluster where no mutation site was selected. After removing all mutation sets which contain fewer than 8 distinct mutation sites, the results in approximately 1700 mutation site sets per target virus.

A.3.5 SAMPLE MUTATIONS

We are now able to generate designed sequences by selecting a single amino acid mutation at each mutational site. We use the EVE single-mutant scores described above in section A.3.2 to guide this process. We generate three versions of each of the ~ 1700 designed mutation site sets per target virus. First, we select the fittest possible single mutation at each of the 8 mutation sites. We call the set of all designed sequences for a target virus generated with this method $\mathbb{S}_{\text{fittest}}$. Next, we select a random mutation at each mutation site without considering fitness. We call the set of designs generated with this method $\mathbb{S}_{\text{random}}$. Finally, we construct a vector for each mutational site, consisting of fitness scores for all possible single mutations at that site which are above the 50th percentile cutoff for single-mutant scores across the RBD. We normalize fitness scores between 0 and 1 within each mutational site. We consider the set of normalized scores as a probability vector, which we use to select a mutation via multinomial sampling. Mutations with a higher fitness score are more likely to be selected, but we allow multiple mutation possibilities at each site in order to introduce diversity across the set of generated designs, particularly in cases where multiple mutations have similarly high fitness. We call the set of designs generated with this method $\mathbb{S}_{\text{sampled}}$. Note that we generate $\mathbb{S}_{\text{fittest}}$, $\mathbb{S}_{\text{random}}$, and $\mathbb{S}_{\text{sampled}}$ separately for each target virus, using our 9 individual EVE models trained on the MSA for each target.

We use the $\mathbb{S}_{\text{fittest}}$ and $\mathbb{S}_{\text{random}}$ distributions in order to select designed sequences from $\mathbb{S}_{\text{sampled}}$ that are diverse yet sufficiently fit. First, we determine the full-RBD EVE fitness score for all designs (each with 8 mutations) in $\mathbb{S}_{\text{fittest}}$ for each target virus. We use the 50th percentile of full-RBD $\mathbb{S}_{\text{fittest}}$ scores as a minimum fitness threshold, and we remove all designs in $\mathbb{S}_{\text{sampled}}$ with EVE fitness scores below this threshold. Second, we remove any remaining designs in $\mathbb{S}_{\text{sampled}}$ with full-RBD EVE fitness scores which are less than the highest full-RBD score in $\mathbb{S}_{\text{random}}$. We cross-check this downselection process by scoring each design also with EVcouplings (Hopf et al., 2017). We remove any design which do not meet these fitness benchmarks (fitter than the 50th percentile of $\mathbb{S}_{\text{fittest}}$, and fitter than the highest-scoring sequence in $\mathbb{S}_{\text{random}}$) when scored with both EVE and EVcouplings.

A.3.6 OPTIMIZE COMBINATIONS OF DESIGNS

The final step of designing a nanoparticle vaccine is selecting a group of 8 sequences, one per target virus, for construction of a mosaic-8 nanoparticle. Each mosaic group may contain either SARS-CoV-2 or SARS-CoV-1, but not both, since we would like to perform separate experiments with SARS-CoV-1 nanoparticles to investigate their effect on the production of recall antibodies. When selecting our final mosaic group, we optimize our design choice in a pair-wise manner, as our goal is to minimize off-target cross-reactivity of epitopes in TDR between pairs of sequences on the nanoparticle. This will minimize antibodies that can bind with avidity to the TDR, and so preferentially elicit antibodies to the more conserved TAR.

Pairwise Distance In order to calculate pairwise distances between all designed sequences, we first remap designs from all target viruses onto the SARS-CoV-2 RBD sequence so they can be easily compared. We use a patch-wise approach to define the pairwise distance between sequences, as each patch represents a potential off-target cross-reactive epitope. We define a potential epitope patch \mathbb{E} around each TDR residue R as the set of all neighboring sites r within 10.58\AA of R , by calculating distances between sidechain centers of mass. 10.58\AA was selected as the maximum distance between the sidechain centers of any two linearly adjacent residues in the SARS-CoV-2 RBD. After defining potential epitope patches, we can compare the same epitope patch on any pair of designs using an accessibility-dissimilarity distance metric. We use EVEscape’s physicochemical dissimilarity score P' , which considers differences in charge and hydrophobicity between amino acids, and accessibility score \mathbf{a} , which uses the weighted contact number of residues to identify surface exposed and protruding positions where antibodies are likely to bind (Thadani et al., 2023). The joint accessibility-dissimilarity distance is calculated as the average of the products of dissimilarity and accessibility across all residues in an epitope:

$$\text{dist}[e, \mathbf{s}_j, \mathbf{s}_k] = \frac{1}{|e|} \sum_{r \in e} \mathbf{a}[r] \cdot P'[\mathbf{s}_j[r], \mathbf{s}_k[r]]$$

By combining these two scores, we are able to consider the physicochemical dissimilarity between amino acids in the two sequences at a given position, and then weight by the antibody accessibility of that position. We finally aggregate distance scores to get a distance per pair of designs as the weighted average over all epitopes, where the joint accessibility-dissimilarity score for each epitope is weighted by the average accessibility score of that epitope.

$$\text{acc}[e] = \frac{1}{|e|} \sum_{r \in e} \mathbf{a}[r]$$

Mosaic Nanoparticles To select a mosaic group of 9 sequences, we use the joint accessibility-dissimilarity pairwise distance metric described above to create a distance matrix between all pairs of designs. We then follow a greedy solution to the maximum dispersion problem, which seeks to find a set of 9 sequences from the complete set of designs such that the sum of pairwise distances between the selected designs are maximized. Note that so pairs of designs from the same target virus will never be chosen, we consider these pairs of designs to have a negative distance value, and hence all 9 sequences will be from a different target virus. For our greedy approach, we start in turn with adding each possible sequence \mathbf{s} and adding it to the mosaic nanoparticle \mathbb{N} . We then add to \mathbb{N} by iteratively selecting the next design which maximizes the minimum distance to all designs already in \mathbb{N} . When selecting SARS-CoV-2 sequences, we do not consider distance to SARS-CoV-1 sequences, and vice versa, since no mosaic group will contain both viruses.

Experimental Validation The above process generates a potential mosaic nanoparticle \mathbb{N} for each possible starting sequence \mathbf{s} . We repeated this process with each starting sequence and ultimately selected 10 nanoparticle designs for experimental validation, which have the highest rank sum of distances among all chosen nanoparticles. Designs will be characterized first at the individual RBD level for expression and non-aggregation, as well as binding to a class of known broadly neutralizing antibodies. After passing these initial checks, designs will be constructed as a mosaic RBD nanoparticle vaccine, and used to immunize mice. The resulting sera will be characterized for its breadth of neutralization against existing and potential future sarbecoviruses.