Atomically accurate de novo design of single-domain antibodies

Nathaniel R. Bennett[‡1,2,3], Joseph L. Watson*[‡1,2], Robert J. Ragotte[‡1,2], Andrew J. Borst[‡1,2], Déjenaé L. See[#1,2,4], Connor Weidle[#1,2], Riti Biswas[1,2,3], Ellen L. Shrock[1,2], Philip J. Y. Leung[1,2,3], Buwei Huang[1,2,4], Inna Goreshnik[1,2,5], Russell Ault[6,7], Kenneth D. Carr[2], Benedikt Singer[1,2], Cameron Criswell[1,2], Dionne Vafeados[2], Mariana Garcia Sanchez[2], Ho Min Kim[8,9], Susana Vázquez Torres[1,2,10], Sidney Chan[2], David Baker*[1,2,5]

‡ Contributed Equally
# Co-second author
* To whom correspondence should be addressed

1. Department of Biochemistry, University of Washington, Seattle, WA 98105, USA
2. Institute for Protein Design, University of Washington, Seattle, WA 98105, USA
3. Graduate Program in Molecular Engineering, University of Washington, Seattle, WA 98105, USA
4. Department of Bioengineering, University of Washington, Seattle, WA, USA
5. Howard Hughes Medical Institute, University of Washington, Seattle, WA, USA
6. Department of Pediatrics, Children's Hospital of Philadelphia, Philadelphia, PA 19104, USA
7. Perelman School of Medicine at the University of Pennsylvania, Philadelphia, PA 19104, USA
8. Center for Biomolecular and Cellular Structure, Institute for Basic Science (IBS), Daejeon, 34126, Republic of Korea
9. Department of Biological Sciences, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, 34141, Republic of Korea
10. Graduate Program in Biological Physics, Structure and Design, University of Washington, Seattle, WA, USA

Abstract

Despite the central role that antibodies play in modern medicine, there is currently no way to rationally design novel antibodies to bind a specific epitope on a target. Instead, antibody discovery currently involves time-consuming immunization of an animal or library screening approaches. Here we demonstrate that a fine-tuned RFdiffusion network is capable of designing de novo antibody variable heavy chains (VHH's) that bind user-specified epitopes. We experimentally confirm binders to four disease-relevant epitopes, and the cryo-EM structure of a designed VHH bound to influenza hemagglutinin is nearly identical to the design model both in the configuration of the CDR loops and the overall binding pose.

Introduction

Antibodies are the dominant class of protein therapeutics with over 160 antibody therapeutics currently licensed globally and a market value expected to reach $445 billion in the next five years[1]. Despite immense pharmaceutical interest, therapeutic antibody development still relies on animal immunization or screening of antibody libraries to identify candidate molecules that bind to a desired target. These methods are laborious, time-consuming, and can fail to produce antibodies that interact with the therapeutically relevant epitope[2]. Efforts at computational design of antibodies have grafted residues into existing antibody structures, sampled alternative native CDR loops to improve affinities[3,4] and used Rosetta[5] sequence design to improve the interacting regions. More recently, structure-based and sequence-based deep learning networks have been trained to design novel antibody sequences[6–8], but de novo (no homology to an existing antibody targeting that epitope) design of structurally accurate antibodies has remained elusive. There has been recent progress in designing binding proteins (not antibodies) using RFdiffusion[9,10] which, unlike previous methods, does not require pre-specification of the protein binder backbone, permitting the design of very diverse binders with inherent shape complementarity to the user-specified epitope[9,10]. However, as with other methods for de novo interface design[11,12], these binders almost exclusively rely on regular secondary structure (helical or strand) based interactions with the target epitope, and RFdiffusion is therefore unable to design antibodies de novo (Extended Data Fig. 1).

An ideal method for designing de novo antibodies would enable 1) targeting of any specified epitope on any target of interest; 2) focusing of sampling on the CDR loops, keeping the framework sequence and structure close to a user-specified highly optimized therapeutic antibody framework; and 3) sampling of alternative rigid-body placements of the designed antibody with respect to the epitope. We hypothesized that given the diversity and quality of interfaces RFdiffusion can design, it should be possible to develop specialized versions capable of designing de novo antibodies, given that the underlying thermodynamics of interface formation are the same. RoseTTAFold2 and RFdiffusion (which trains from an earlier version of RF2) are trained on the entire Protein Data Bank (PDB[13]) which helps overcome the problem that the PDB contains relatively few antibody structures (~8,100 antibody structures versus >200,000 total structures) which complicates the training of large machine learning models. We set out to develop versions of RFdiffusion and RoseTTAFold2 specialized for antibody structure design and structure prediction by fine-tuning on native antibody structures. For simplicity, in this work, we henceforth refer to the original RFdiffusion network as "vanilla RFdiffusion", and the antibody-specific variant we describe here simply as "RFdiffusion".

Fine-tuning RFdiffusion for antibody design

RFdiffusion uses the AlphaFold2[14]/RF2 frame representation of protein backbones comprising the Cα coordinate and N-Cα-C rigid orientation for each residue. During training, a noising schedule is used that, over a set number of "timesteps" (T), corrupts the protein frames to distributions indistinguishable from random distributions (Cα coordinates are corrupted with 3D Gaussian noise, and residue orientations with Brownian motion on SO3). During training, a PDB structure and a random timestep (t) are sampled, and t noising steps are applied to the

structure. RFdiffusion predicts the de-noised (pX_0) structure at each timestep, and a mean squared error (m.s.e.) loss is minimized between the true structure (X_0) and the prediction. At inference time, translations are sampled from the 3D Gaussian and uniform rotational distributions (X_T) and RFdiffusion iteratively de-noises these frames to generate a new protein structure.

To explore the design of antibodies, we fine-tuned RFdiffusion predominantly on antibody complex structures (Fig. 1; Methods). At each step of training, an antibody complex structure is sampled, along with a random timestep (t), and this number of noise steps are added to corrupt the antibody structure (but not the target structure). To permit specification of the framework structure and sequence at inference time, the framework sequence and structure are provided to RFdiffusion during training (Fig. 1B). Because it is desirable for the rigid body position (dock) between antibody and target to be designed by RFdiffusion along with the CDR loop conformations, the framework structure is provided in a global-frame-invariant manner during training (Fig. 1C). We utilize the "template track" of RF/RFdiffusion to provide the framework structure as a 2D matrix of pairwise distances and dihedral angles between each pair of residues (a representation from which 3D structures can be accurately recapitulated)[15], (Extended Data Fig. 1A). The framework and target templates specify the internal structure of each protein chain, but not their relative positions in 3D space (in this work we keep the sequence and structure of the framework region fixed, and focus on the design solely of the CDRs and the overall rigid body placement of the antibody against the target). In vanilla RFdiffusion, de novo binders can be targeted to specific epitopes at inference time through training with an additional one-hot encoded "hotspot" feature, which provides some fraction of the residues the designed binder should interact with. For antibody design, where we seek CDR-loop-mediated interactions, we adapt this feature to specify residues on the target protein with which CDR loops interact (Fig. 1D).

With this training regime, RFdiffusion is able to design antibody structures that closely match the structure of the input framework structure, and target the specified epitope with novel CDR loops (Extended Data Fig. 1). After the RFdiffusion step, we use ProteinMPNN to design the CDR loop sequences. The designed antibodies make diverse interactions with the target epitope and differ significantly from the training dataset (Fig. 2E).

Fine-tuning RoseTTAFold2 for antibody design validation

Design pipelines typically produce a wide range of solutions to any given design challenge, and hence readily computable metrics for selecting which designs to experimentally characterize play an important role. An effective way to filter designed proteins and interfaces is based on the similarity of the design model structure to the AlphaFold2 predicted structure for the designed sequence (this is often referred to as "self-consistency"), which has been shown to correlate well with experimental success[16,17]. In the case of antibodies, however, AlphaFold2 fails to routinely predict antibody-antigen structures accurately[18], preventing its use as a filter in an antibody design pipeline.

We sought to build an improved filter by fine-tuning the RoseTTAFold2 structure prediction network on antibody structures. To make the problem more tractable, we provide information

during training about the structure of the target and the location of the target epitope to which the antibody binds; the fine-tuned RF2 must still correctly model the CDRs and find the correct orientation of the antibody against the targeted region. With this training regimen, RF2 is able to robustly distinguish true antibody-antigen pairs from decoy pairs and often accurately predicts antibody-antigen complex structures. Accuracy is higher when the bound (holo) conformation of the target structure is provided (Extended Data Fig. 2); this is available when evaluating design models, but not available in the general antibody-antigen structure prediction case. At monomer prediction, the fine-tuned RF2 outperforms the previously published IgFold network (which can only model antibody monomer structures)[19], especially at CDR H3 structure prediction (Extended Data Fig. 3).

When this fine-tuned RF2 network is used to re-predict the structure of RFdiffusion-designed VHHs, a significant fraction are confidently predicted to bind in an almost identical manner to the designed structure (Extended Data Fig. 4A). Further, in silico cross-reactivity analyses demonstrate that RFdiffusion-designed VHHs are rarely predicted to bind to unrelated proteins (Extended Data Fig. 4B). VHHs that are confidently predicted to bind their designed target are predicted to form high quality interfaces, as measured by Rosetta ddG (Extended Data Fig. 4C). The fact that many of the designed sequences generated by our RFdiffusion antibody design pipeline are predicted by RF2 to adopt the designed structures and binding modes suggested that RF2 filtering might enrich for experimentally successful binders.

Design and biochemical characterization of designed VHHs

We initially focused on the design of single-domain antibodies (VHHs) based on the variable domain from heavy-chain antibodies produced by camelids and sharks[20]. The smaller size of VHHs makes genes encoding designs much easier to assemble and cheaper than single chain variable fragments (scFv; where linker choice is a critical factor[21]) or fragment antigen-binding regions (Fab; where an interchain disulfide bond is required for proper folding[22]). VHHs are readily "humanized"; so far, two VHH-based therapies are approved by the FDA with many clinical trials ongoing[20]. Despite having fewer CDR loops (three) than conventional Fvs (six), the average interaction surface area of a VHH is very similar to that of an Fv[23], suggesting a method capable of VHH design could also be suitable for Fv design. Indeed, in silico metrics for scFvs and VHHs showed similar qualities of interfaces, as assessed by Rosetta[5] and fine-tuned RF2 (Extended Data Fig. 6).

We chose a widely used humanized VHH framework (h-NbBcII10FGLA; [ref [24]]) as the basis of our VHH design campaigns, and designed VHHs to a range of disease-relevant targets: *Clostridium difficile* toxin B (TcdB), influenza H1 hemagglutinin (HA), respiratory syncytial virus (RSV) sites I and III, SARS-CoV-2 receptor binding domain (Covid RBD) and IL-7Rα. ProteinMPNN[25] was used to design the sequences of the CDR loops (but not the framework) in the context of the target. We then filtered designs with the fine-tuned RoseTTAFold2 network (Methods) described above. Designs were screened either at high-throughput by yeast surface display (9000 designs per target; RSV sites I and III, Covid RBD, Influenza HA) or at lower-throughput with E. coli expression and single-concentration surface plasmon resonance (95 designs per target; TcdB, IL-7Rα and influenza HA–the latter was screened using both methods).

In the case of influenza HA, glycan N296, located along the HA-stem epitope, exhibited varying degrees of overlap with the approach angle of several of our designed VHHs. To best align the experimental design conditions with the computational parameters employed during design (i.e., excluding consideration of the glycan shield), affinity measurements were conducted using a commercially produced monomeric HA product expressed in insect cells (Extended Data Fig. 9). Insect cells express a truncated paucimannose glycan shield, which - relative to a natively expressed HA trimer - more closely resembles the fully deglycosylated HA monomeric PDB model used for VHH design. Of the HA binders tested against the insect-cell produced HA monomer, the highest affinity binder was measured to have a Kd of 78nM, (Fig. 2), with other binders having affinities of 546nM, 698nM, and 790nM.

The highest affinity binders to RSV site III, Influenza HA, Covid RBD, and TcdB are shown in Fig. 2A,B,C,E respectively (see also Extended Data Fig. 8 for all the SPR traces of confirmed VHH binders identified in this study). The CDR loops are distinct from VHHs observed in nature, indicating significant generalization beyond the training dataset (Fig 2E, Extended Data Fig. 5). For TcdB, there are no antibodies or VHHs targeting this site in the PDB. For the best designed VHH from both Covid RDB (Kd = 5.5μM; Fig. 2C) and TcdB (Kd = 262nM; Fig 2D) binding was confirmed to be to the desired epitope: binding was completely abolished upon addition of a previously designed, structurally characterized de novo binder to that epitope (AHB2, PDB: 7UHB[26] for Covid RBD and Fzd48 [manuscript in preparation] for TcdB) (Fig. 2C,D; Extended Data Fig. 7). For TcdB, the interactions were specific, with no binding observed to the highly related *Clostridium sordellii* toxin L (TcsL) (Extended Data Fig. 7B). These data demonstrate the ability of RFdiffusion to design VHHs making specific interactions with the target epitope. Surprisingly, design success rates were not significantly higher for filtered designs vs unfiltered designs with the RF2 settings we used (providing 100% of interface hotspots, although there was some signal with more stringent settings where we provided 0% or 10% of interface hotspots during prediction). However, given the small dataset, more extensive datasets will be necessary to evaluate more conclusively how best to use and fine-tune RF2 for design filtering.

Cryo-electron microscopy reveals atomically accurate VHH design against a natively glycosylated viral glycoprotein

Given the success of RFdiffusion at generating moderate affinity VHHs against diverse epitopes, we sought to evaluate design accuracy by cryo-EM structure determination of the designed anti-HA VHHs in complex with natively glycosylated, trimeric influenza HA glycoprotein (strain A/USA:Iowa/1943 H1N1), which retains the conserved stem epitope used during computational VHH design and upstream biochemical screening. The VHHs were combined with Iowa43 HA at a 3:1 molar excess ratio (VHH:HA monomer) at a concentration of 15μM and promptly prepared for cryo-EM grid freezing. Cryo-EM data processing revealed one VHH design effectively bound to the fully glycosylated HA trimer (out of the four tested), denoted hereafter as *VHH_flu_01* (Fig. 3). 2D classification of all particles in the dataset (Fig. 3A) and the solved 3.0Å structure of the complex (Fig. 3B) identified approximately 66% of HA particles bound to a maximum of two VHHs per trimer (Fig. 3A-H). This partial occupancy is likely attributable to the N296 glycan, which in unbound subunits partially occludes the target epitope but reorients when bound to *VHH_flu_01* (see Fig. 3H).

The structure of influenza HA bound to two copies of *VHH_flu_01* (Figure 3B,C, Extended Data Fig. 10) reveals a VHH approach angle which closely matches the predicted model (Fig. 3D), and a VHH backbone which is very close to the RFdiffusion design, with a calculated R.M.S.D. of 1.45Å (Fig. 3E). The CDR3 structure is also very similar between the cryo-EM structure and the computational model (R.M.S.D. = 0.8Å) (Fig. 3F), with residues V100, V101, S103, and F108 in the de novo designed CDR3 loop interacting with the influenza HA stem epitope in the cryo-EM structure, as designed by RFdiffusion and re-predicted with RF2 (Fig. 3G). Notably, the design is highly dissimilar from the closest antibody/VHH binding to this epitope in the PDB (Extended Data Fig. 5G,H). Taken together, these results highlight the ability of a de novo designed VHH with a novel CDR3 to accurately bind a natively glycosylated epitope with atomic-level precision.

Discussion

Our results demonstrate that computational de novo design of antibodies is now possible. The high resolution cryo-EM structure of our designed VHH to influenza HA demonstrates atomically accurate design of a VHH complex (including the highly variable H3 loop and the overall binding orientation).

With further improvements, computational de novo design of antibodies using our RFdiffusion and related approaches could revolutionize antibody discovery and development. Our RFdiffusion approach enables targeting specific epitopes of interest on the target antigen, and, when success rates increase, should be far faster and cheaper than immunizing an animal or screening a random library. By taking a structure-based approach to antibody design, the optimization of critical pharmaceutical properties such as aggregation, solubility, and expression level[27] may be tuned in a structurally aware manner (avoiding mutations which would disrupt the antibody-target interface or which would destabilize the antibody). Furthermore, the ability to explore the full space of CDR loop sequences and structures from the start, particularly for CDR1 and CDR2 which are natively limited to the space of sequences encoded by germline V genes prior to somatic hypermutation, should simplify both the optimisation of the developability features and the targeting of non-immunodominant epitopes[28]. Finally, every antibody designed by RFdiffusion has a strong structural hypothesis (further validated by RoseTTAFold2), which should enable the rational design of antibody function, by targeting specific target conformational states, for example.

Although our results demonstrate successful de novo design of VHHs, there is considerable room for improvements, as the binding affinities are modest (although comparable to affinities of de novo miniprotein binders without experimental optimization when this challenge was first solved[11]), and the success rates are still quite low. For the backbone design step, incorporating recent architectural improvements[29] or newer generative frameworks, such as flow-matching[30,31] may yield design models with higher designability and diversity. RoseTTAFold2 and vanilla RFdiffusion have also recently been extended to model all biomolecules (rather than just proteins)[32], and incorporating this capability into the antibody design RFdiffusion should permit the design of antibodies to epitopes containing non-protein atoms, such as glycans. Indeed, the sub-stoichiometric binding observed for *VHH_flu_01* could be explained by the presence of nearby glycan N296, which was not considered during the initial design of this VHH. ProteinMPNN was not modified in this current work, but designing sequences that

more closely match human CDR sequences would be expected to reduce the potential immunogenicity of designed antibodies[33]. Directly optimizing developability properties in ProteinMPNN is another future direction. Finally, improvements in RoseTTAFold2 antibody prediction methods should improve experimental success rates, and allow better in silico benchmarking of upstream design methods.

Altogether, we expect this work to be the foundation of a new era of structure-based antibody design.

Acknowledgements

Author Contributions

N.R.B. and J.L.W. conceived the study, and may change the order of their names for personal pursuits to best suit their own interests. N.R.B. and J.L.W. trained RFdiffusion and fine-tuned RoseTTAFold2. R.J.R. led the experimental work, with help from D.L.S., R.B., E.L.S., P.J.Y.L., B.H., I.G., R.A. and S.V.T.. A.J.B. led the nsEM and cryo-EM characterization work, with help from C.W. and K.C.. J.L.W., N.R.B, D.L.S., R.A., C.C., H.M.K. made designs. D.L.S. and B.S. contributed additional code. S.C. purified target proteins. J.L.W. and D.B. co-managed the project. N.R.B., J.L.W., D.B., R.J.R. and A.J.B. wrote the manuscript. All authors read and contributed to the manuscript.

Correspondence to Joseph L. Watson (jwatson3@uw.edu) and David Baker (dabaker@uw.edu)

Competing Interests

N.R.B., J.L.W., R.J.R., A.J.B., C.W., P.J.Y.L., B.H., and D.B. are co-inventors on U.S. provisional patent number 63/607,651 which covers the computational antibody design pipeline described here.

Bibliography

1. Lyu, X. *et al.* The global landscape of approved antibody therapies. *Antib. Ther.* **5**, 233–257 (2022).

2. Wilson, P. C. & Andrews, S. F. Tools to therapeutically harness the human antibody response. *Nat. Rev. Immunol.* **12**, 709–719 (2012).

3. Liu, X. *et al.* Computational design of an epitope-specific Keap1 binding antibody using hotspot residues grafting and CDR loop swapping. *Sci. Rep.* **7**, 41306 (2017).

4. Sormanni, P., Aprile, F. A. & Vendruscolo, M. Rational design of antibodies targeting specific epitopes within intrinsically disordered proteins. *Proc. Natl. Acad. Sci. U. S. A.* **112**, 9902–9907 (2015).

5. Leaver-Fay, A. *et al.* ROSETTA3: an object-oriented software suite for the simulation and design of macromolecules. *Methods Enzymol.* **487**, 545–574 (2011).

6. Xie, X. *et al.* Antibody-SGM: Antigen-Specific Joint Design of Antibody Sequence and Structure using Diffusion Models.

7. Eguchi, R. R. *et al.* Deep Generative Design of Epitope-Specific Binding Proteins by Latent Conformation Optimization. 2022.12.22.521698 Preprint at https://doi.org/10.1101/2022.12.22.521698 (2022).

8. Shanehsazzadeh, A. *et al. Unlocking de Novo Antibody Design with Generative Artificial Intelligence*. http://biorxiv.org/lookup/doi/10.1101/2023.01.08.523187 (2023) doi:10.1101/2023.01.08.523187.

9. Watson, J. L. *et al.* De novo design of protein structure and function with RFdiffusion. *Nature* **620**, 1089–1100 (2023).

10. Torres, S. V. *et al.* De novo design of high-affinity binders of bioactive helical peptides. *Nature* (2023) doi:10.1038/s41586-023-06953-1.

11. Cao, L. *et al.* Design of protein-binding proteins from the target structure alone. *Nature* **605**, 551–560 (2022).

12. Gainza, P. *et al.* De novo design of protein interactions with learned surface fingerprints. *Nature* **617**, 176–184 (2023).

13. Berman, H. M. *et al.* The Protein Data Bank. *Nucleic Acids Res.* **28**, 235–242 (2000).

14. Jumper, J. *et al.* Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583–589 (2021).

15. Yang, J. *et al.* Improved protein structure prediction using predicted interresidue orientations. *Proc. Natl. Acad. Sci. U. S. A.* **117**, 1496–1503 (2020).

16. Wang, J. *et al.* Scaffolding protein functional sites using deep learning. *Science* **377**, 387–394 (2022).

17. Bennett, N. *et al.* Improving de novo Protein Binder Design with Deep Learning. *Nat. Commun.* **14**, 2022.06.15.495993 (2023).

18. Yin, R. & Pierce, B. G. Evaluation of AlphaFold Antibody-Antigen Modeling with Implications for Improving Predictive Accuracy. *bioRxiv* 2023.07.05.547832 (2023) doi:10.1101/2023.07.05.547832.

19. Ruffolo, J. A., Chu, L.-S., Mahajan, S. P. & Gray, J. J. Fast, accurate antibody structure prediction from deep learning on massive set of natural antibodies. *Nat. Commun.* **14**, 2389 (2023).

20. Jin, B., Odongo, S., Radwanska, M. & Magez, S. Nanobodies: A Review of Generation, Diagnostics and Therapeutics. *Int. J. Mol. Sci.* **24**, 5994 (2023).

21. Hao, C. *et al.* Effects of different interchain linkers on biological activity of an anti-prostate cancer single-chain bispecific antibody. *Theor. Biol. Med. Model.* **12**, 14 (2015).

22. Gani, K., Chirmade, T., Ughade, S., Thulasiram, H. & Bhambure, R. Understanding unfolding and refolding of the antibody fragment (Fab) III: Mapping covalent and non-covalent interactions during in-vitro refolding of light chain, heavy chain, and Fab. *Biochem. Eng. J.* **187**, 108644 (2022).

23. Mitchell, L. S. & Colwell, L. J. Analysis of nanobody paratopes reveals greater diversity

than classical antibodies. *Protein Eng. Des. Sel.* **31**, 267–275 (2018).

24. Vincke, C. *et al.* General Strategy to Humanize a Camelid Single-domain Antibody and Identification of a Universal Humanized Nanobody Scaffold. *J. Biol. Chem.* **284**, 3273–3284 (2009).

25. Dauparas, J. *et al.* Robust deep learning-based protein sequence design using ProteinMPNN. *Science* **378**, 49–56 (2022).

26. Hunt, A. C. *et al.* Multivalent designed proteins neutralize SARS-CoV-2 variants of concern and confer protection against infection in mice. *Sci. Transl. Med.* **14**, eabn1252 (2022).

27. Kim, J., McFee, M., Fang, Q., Abdin, O. & Kim, P. M. Computational and artificial intelligence-based methods for antibody development. *Trends Pharmacol. Sci.* **44**, 175–189 (2023).

28. Shrock, E. L. *et al.* Germline-encoded amino acid-binding motifs drive immunodominant public antibody responses. *Science* **380**, eadc9498 (2023).

29. Wang, C. *et al.* Proteus: pioneering protein structure generation for enhanced designability and efficiency. 2024.02.10.579791 Preprint at https://doi.org/10.1101/2024.02.10.579791 (2024).

30. Yim, J. *et al.* Fast protein backbone generation with SE(3) flow matching. Preprint at http://arxiv.org/abs/2310.05297 (2023).

31. Bose, J. *et al.* SE(3)-Stochastic Flow Matching for Protein Backbone Generation. in (2023).

32. Krishna, R. *et al.* Generalized biomolecular modeling and design with RoseTTAFold All-Atom. *Science* eadl2528 (2024) doi:10.1126/science.adl2528.

33. Gao, S. H., Huang, K., Tu, H. & Adler, A. S. Monoclonal antibody humanness score and its applications. *BMC Biotechnol.* **13**, 55 (2013).

34. Altschul, S. F. *et al.* Gapped BLAST and PSI-BLAST: a new generation of protein

database search programs. *Nucleic Acids Res.* **25**, 3389–3402 (1997).

35. Dunbar, J. *et al.* SAbDab: the structural antibody database. *Nucleic Acids Res.* **42**, D1140-1146 (2014).

36. Jäger, M., Gehrig, P. & Plückthun, A. The scFv fragment of the antibody hu4D5-8: evidence for early premature domain interaction in refolding. *J. Mol. Biol.* **305**, 1111–1129 (2001).
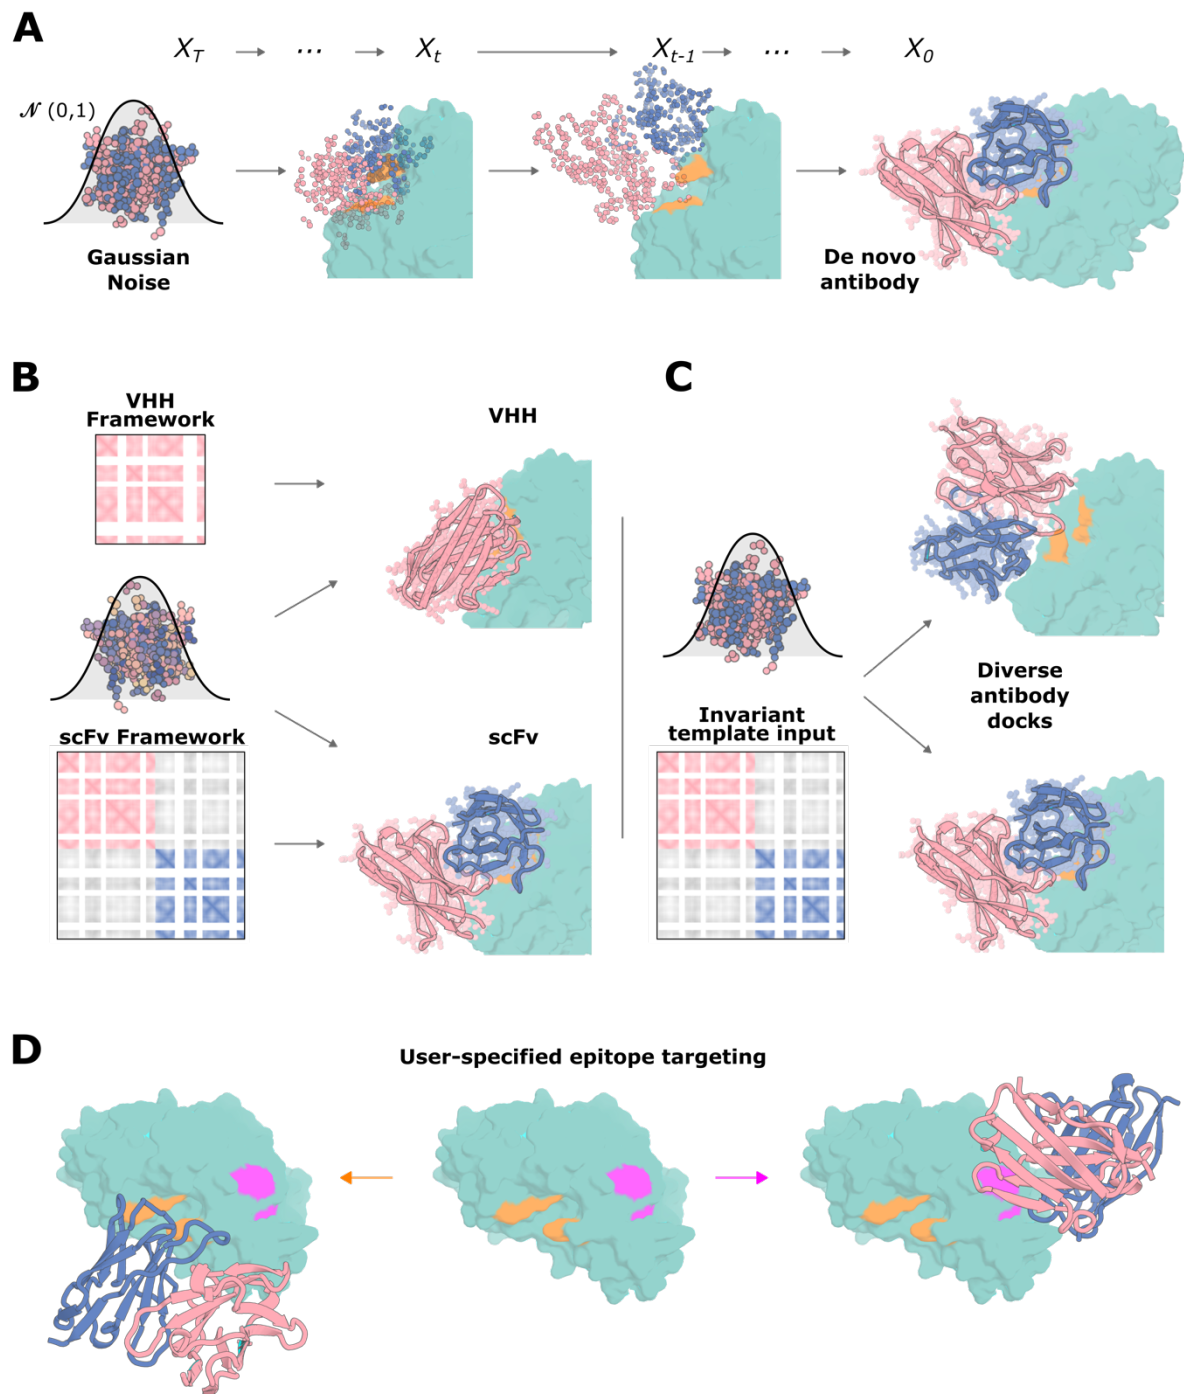
Figures

Figure 1: Overview of RFdiffusion for antibody design

A) RFdiffusion is trained such that at time T, a sample is drawn from the noise distribution (3D Gaussian distribution for translations, and uniform SO3 distribution for rotations), and this sampled noise is then "de-noised" between times T and 0, to generate an (in this case) scFv binding to the target structure through its CDR loops. B) Control over which framework is used is provided through input of a framework "template", which specifies the pairwise distances and dihedral angles between residues in the framework. The sequence of the framework

region is also included. For example, provision of a VHH framework generates a VHH (top row), whereas provision of an scFv framework generates a scFv (bottom row). C) Diversity in the antibody-target dock is achieved through the pairwise framework representation, which, because the framework structure is provided on a separate template to that of the target, does not provide information about the rigid body framework-target relationship. Hence, diverse docking modes are sampled by RFdiffusion. D) The epitope to which the antibody binds can be specified by provision of input "hotspot" residues, which direct the designed antibody (compare orange, left vs pink, right).

Figure 2: Biochemical characterization of designed VHHs

A-**B**) 9000 designed VHHs were screened against RSV site III and influenza hemagglutinin with yeast surface display, before soluble expression of the top hits in *E. coli*. Surface Plasmon Resonance (SPR) demonstrated that the highest affinity VHHs to RSV site III and Influenza Hemagglutinin bound their respective targets with 1.4μM and 78nM respectively. **C**) 9000 VHH designs were tested against SARS-CoV-2 receptor binding domain (RBD), and after soluble expression, SPR confirmed an affinity of 5.5μM to the target. Importantly, binding was to the expected epitope, confirmed by competition with a structurally confirmed de novo binder (AHB2, PDB: 7UHB). **D**) 95 VHH designs were tested against the *C. difficile* toxin TcdB. The highest affinity VHH bound with 262nM affinity, and also competed with an unpublished, structurally confirmed de novo binder to the same epitope (right). See also Extended Data Fig. 7 for quantification of the competition shown in **C** and **D**. **E**) Designed VHHs were distinct from the training dataset. Blastp[34] was used to find hits against the SAbDab[35], and the similarity of the CDR loops in the top blast hit were reported for all VHHs experimentally tested in this study. Note also that the 28 VHHs confirmed to bind their targets by SPR do not show enhanced similarity to the training set (red lines).

Figure 3: Cryo-EM structure of a de novo designed VHH binding to influenza hemagglutinin.

A) Labeled cryo-EM 2D class averages of a designed VHH, *VHH_flu_01*, bound to influenza HA, strain A/USA:Iowa/1943 H1N1. **B**) A 3.0Å cryo-EM 3D reconstruction of the complex viewed along two orthogonal axes shows *VHH_flu_01* bound to H1 along the stem in two of the three protomers. **C**) Cryo-EM structure of *VHH_flu_01* bound to influenza HA. **D**) The cryo-EM structure of *VHH_flu_01* in complex with HA closes matches the design model. **E**) cryo-EM reveals the accurate design of *VHH_flu_01* using RFdiffusion (R.M.S.D. to the RFdiffusion design of the VHH is 1.45 Å). **F**) Superposition of the designed VHH CDR3 predicted structure as compared to the built cryo-EM structure (R.M.S.D = 0.84Å). **G**) Comparison of predicted CDR3 rotamers compared to the built 3.0Å cryo-EM structure. **H**) Examination of apo HA protomers juxtaposed with those bound to the designed VHH unveils a notable repositioning

and accommodation of glycan N296 to allow for binding of the designed VHH to the HA stem. In each structural depiction panel, the designed VHH predicted structure is showcased in gray, while the cryo-EM solved structure of the designed VHH is depicted in purple. Additionally, the HA glycoprotein is represented in tan, and the HA glycan shield is illustrated in green.

Extended Data Figures

A **RFdiffusion fine-tuning is required for accurate framework recapitulation**

B **RFdiffusion fine-tuning is required for efficient hotspot targeting**

C

Extended Data Figure 1: Fine-tuning is required for antibody design with RFdiffusion

A) To test whether existing vanilla RFdiffusion models were capable of designing VHHs/scFvs, we explored means of providing the antibody template. For VHHs (left), we used RFdiffusion variant trained to condition on sequence alone[10] and provided the VHH framework sequence (gray). This version, as compared to the fine-tuned version described in this work (pink), was significantly worse at recapitulating the native VHH framework structure. For scFvs (right), we additionally tried providing fold-level information into the appropriate vanilla RFdiffusion model[9] (dark gray), but found that this was also insufficient to get accurate recapitulation of the scFv framework. Fine-tuning (pink) yields near-perfect recapitulation of the scFv framework structure. **B**) Although vanilla RFdiffusion is trained to respect "hotspots", for VHHs (left) and scFvs (right) we find this to be less robust (grays) than after fine-tuning on antibody design (pink). **C**) Examples depicting the results of (**A**) and (**B**). In all cases, the "median" accuracy example (by framework recapitulation) is shown. Left to right: i) without fine-tuning, vanilla RFdiffusion does not target "hotspot" residues (orange) effectively, and does not recapitulate the VHH framework accurately (gray vs yellow). ii) After fine-tuning on antibody design, RFdiffusion targets "hotspots" with accurately recapitulated VHHs (pink vs yellow). iii) Providing only the scFv sequence, vanilla RFdiffusion does not target "hotspots" (orange) accurately nor accurately recapitulates the VHH framework (gray vs yellow). iv) Providing additional fold-level information is insufficient to get perfect framework recapitulation (dark gray vs yellow). v) After fine-tuning on antibody design, RFdiffusion can design scFvs with accurate framework structures (blue/pink vs gray) targeting the input "hotspots" (orange).
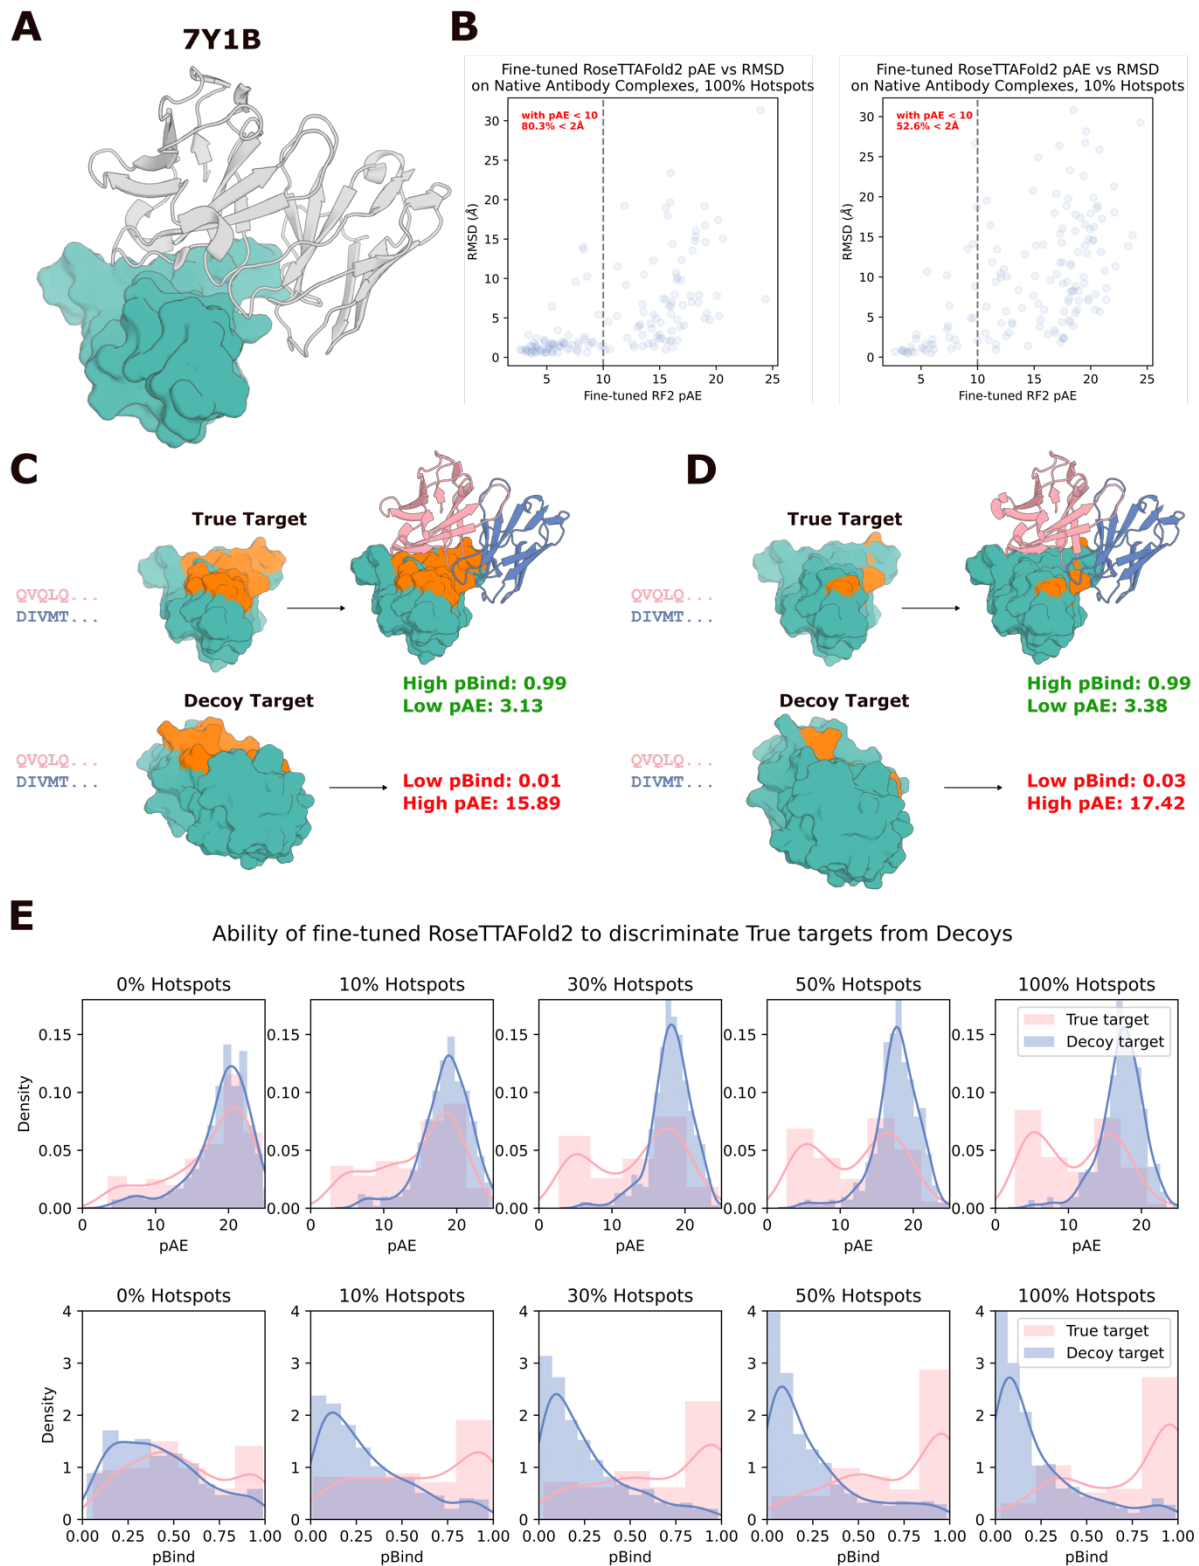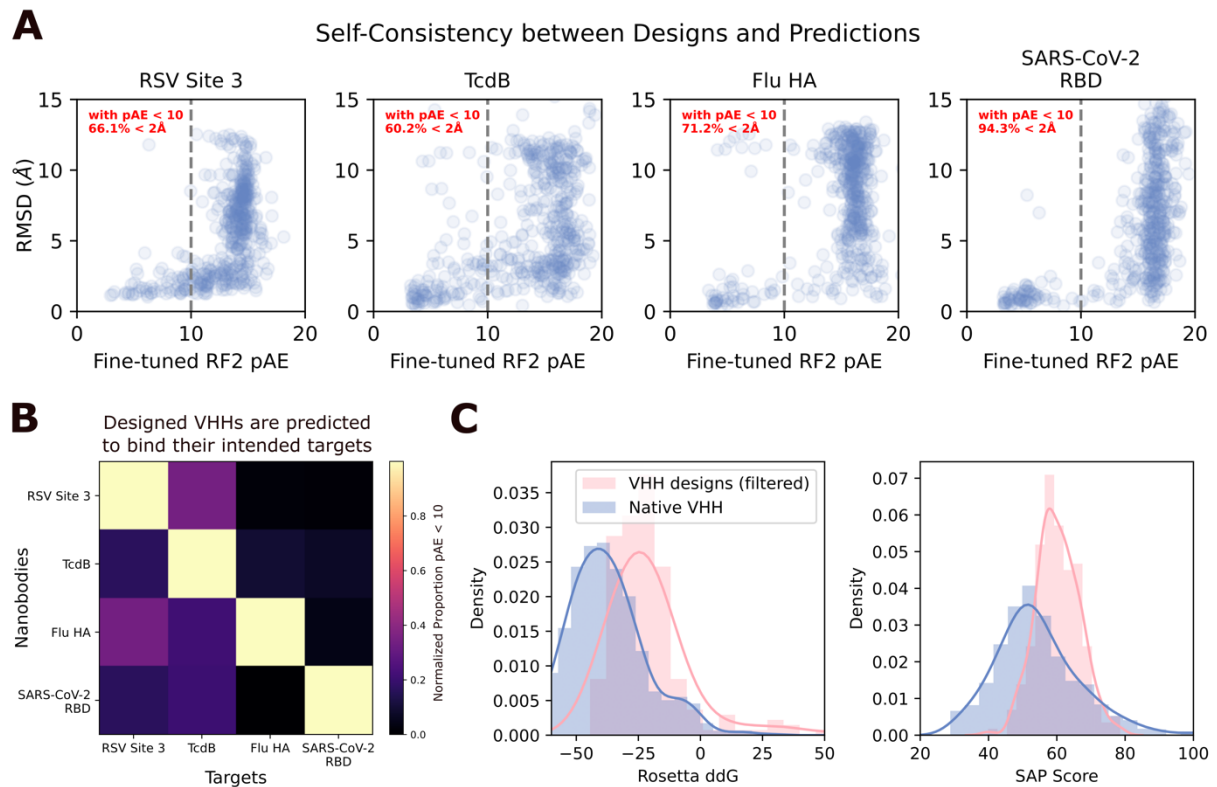
Extended Data Figure 2: Fine-tuned RoseTTAFold2 can distinguish true complexes from decoy complexes

A) An example antibody structure from the validation set used in this figure, which shares < 30% sequence similarity on the target (teal) to anything in the RoseTTAFold2 fine-tuning training dataset. **B**) Fine-tuned RoseTTAFold2 quite reliably predicts its own accuracy. Correlation between RF2 pAE and R.M.S.D. to the native structure with 100% (left) or 10% (right) of "hotspot" residues provided. With pAE < 10, 80.3% of structures are within 2Å when 100% of "hospots" are provided (along with the holo target structure), with this falling to 52.6% when only 10% of hotspots are provided. **C-D**) Cherry-picked example of RoseTTAFold2 correctly distinguishing a "true" from a "decoy" complex. The sequence of antibody 7Y1B was provided either with the correct (PDB: 7Y1B) or decoy (PDB: 8CAF) target. Both with 100% (**C**) or 10% (**D**) of "hotspots" provided, RF2 near-perfectly predicts binding (top row) or non-binding (bottom row). **E**) Quantification of the fine-tuned RF2's ability to distinguish true targets from decoy targets with both pAE (top row) and pBind (bottom row). Note that this ability depends on the proportion of "hotspots" provided. Without any "hotspots" provided, RF2 is hardly predictive, because RF2 without privileged information is quite rarely confident or accurate in its predictions.
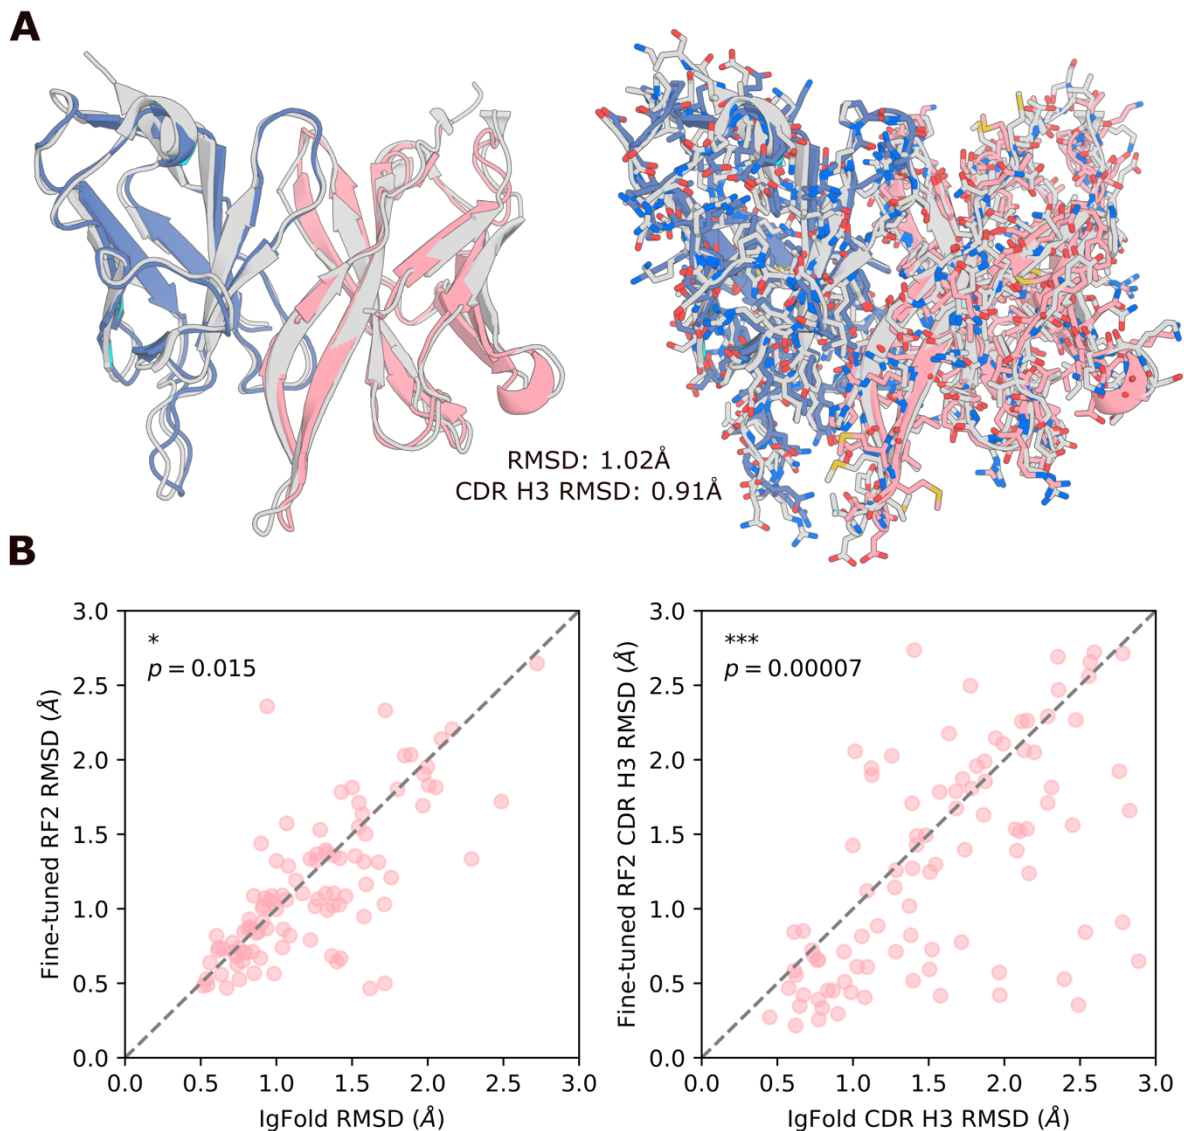
Extended Data Figure 3: Comparison of fine-tuned RoseTTAFold2 to IgFold on antibody monomer prediction

A) 104 antibodies released after the RF2 (and IgFold) training dataset date cutoff (January 13th, 2023) that share < 30% target sequence similarity to any antibody complex released prior to this date were predicted as monomers with either fine-tuned RF2 or IgFold (IgFold cannot predict antibody-target complexes). Shown is the median Fv quality prediction (by overall RMSD) of fine-tuned RF2, of PDB 8GPG, with (right) and without (left) sidechains shown. While the backbone R.M.S.D. is close to the true structure, some sidechains are incorrectly positioned. **B**) Fine-tuned RF2 slightly outperforms IgFold at prediction accuracy. Overall prediction accuracy is slightly improved in fine-tuned RF2 vs IgFold (*p=0.015*, Wilcoxon Paired Test), with greater improvements in CDR H3 prediction accuracy (*p=0.00007*, Wilcoxon Paired Test).

Extended Data Figure 4: Fine-tuned RoseTTAFold2 recapitulates design structures and computationally demonstrates specificity of VHHs for their targets

A) Comparison of RF2 pAE and R.M.S.D of the prediction to the design model. A significant fraction of designs are re-predicted by RF2 (given 100% of "hotspots"), and pAE correlates well with accuracy to the design model. **B**) RF2 can be used to assess quality of designed VHHs. Providing the VHH sequence with the true target structure (used during design) leads to higher rates of high-confidence predictions than predicting the same sequence with a decoy structure (not used in design), as assessed by the fraction of predictions with pAE < 10 (normalized to the fraction of predictions with pAE < 10 for that target with its "correct" VHH partners). In these experiments, the true or decoy target was provided along with 100% of hotspot residues, with those hotspot residues derived from the target with its "true" designed VHH bound. **C**) Orthogonal assessment of designed VHHs with Rosetta demonstrates that the interfaces of RF2-approved (R.M.S.D. < 2Å to design model, pAE < 10) VHH designs have low ddG (top; only slightly worse than native VHHs) and slightly higher SAP score as compared to natives (bottom).

Extended Data Figure 5: Alignment of VHH Design Models to Complexes in the PDB

For each of the highest affinity VHHs identified for each target, and the structurally characterized influenza HA VHH, the closest complex in the PDB is shown. Designed VHHs (pink) are shown in complex with their designed target (teal and tan). The closest complex was identified visually (Methods). **A**) Designed TcdB VHH aligned against 3 VHHs from 6OQ5 (shades of blue). The designed TcdB VHH binds to a site for which no antibody or VHH structure exists in the PDB. **B**) Designed RSV Site III VHH aligned against VHH from 5TOJ (blue). **C**) Designed SARS-CoV-2 VHH aligned against VHH from 8Q94 (blue). **D**) Designed SARS-CoV-2 VHH aligned against Fab from 7FCP (shades of blue). **E**) Highest affinity designed influenza HA VHH aligned against Fv from 8DIU (shades of blue). **F**) Highest affinity designed influenza HA VHH aligned against VHH from 6YFT (blue). **G**) Structurally characterized designed influenza HA VHH aligned against Fv from 8DIU (shades of blue). **H**) Structurally characterized designed influenza HA VHH aligned against VHH from 6YFT (blue).

Extended Data Figure 6: In silico evaluation of RFdiffusion scFv designs

A) RFdiffusion was used to generate scFv designs using the framework from Herceptin (hu4D5-8), which has been used to make scFvs previously[36]. Five targets were chosen (IL10 Receptor-α, TLR4, β-lactamase, TcdB and SARS-CoV-2 (omicron) RBD (PDBs: 6X93, 4G8A, 4ZAM, 7ML7, 7WPC). Shown are five examples with close agreement between the design model and the fine-tuned RF2 prediction (R.M.S.D. (Å): 0.60, 0.56, 0.46, 0.43, 0.61; pAE: 4.73, 4.10, 4.49, 3.52, 3.65). Gray: designs, Pink: RF2 prediction. **B**) Against the four targets to which VHHs were successfully designed, fine-tuned RF2 predicts good specificity to the designed target vs decoy targets. **C**) Against the five targets shown in (**A**), fine-tuned RF2 similarly predicts high specificity to the designed target vs decoy targets. **D**) Orthogonal assessment of designed scFvs with Rosetta demonstrates that the interfaces of RF2-approved (R.M.S.D. < 2Å to design model, pAE < 10) scFv designs have low ddG (top; only slightly worse than native Fabs) and lower SAP score as compared to natives (bottom).

Extended Data Figure 7: Analysis of SPR Competition Assays

The average response during VHH injection normalized to the response immediately preceding VHH injection for **A**) TcdB VHH competition with Fzd48. **B**) TcdB VHH does not bind to the closely related *Clostridium sordellii* TcsL toxin, indicating that it is binding through specific interactions. **C**) SARS-CoV-2 RBD VHH competition with AHB2. For the competition experiments, in the miniprotein binder-only trace, no VHH is injected and the average response over the corresponding period is plotted as a baseline. (**A**) and (**C**) are the quantification from the rightmost panels of Fig. 2C-D.
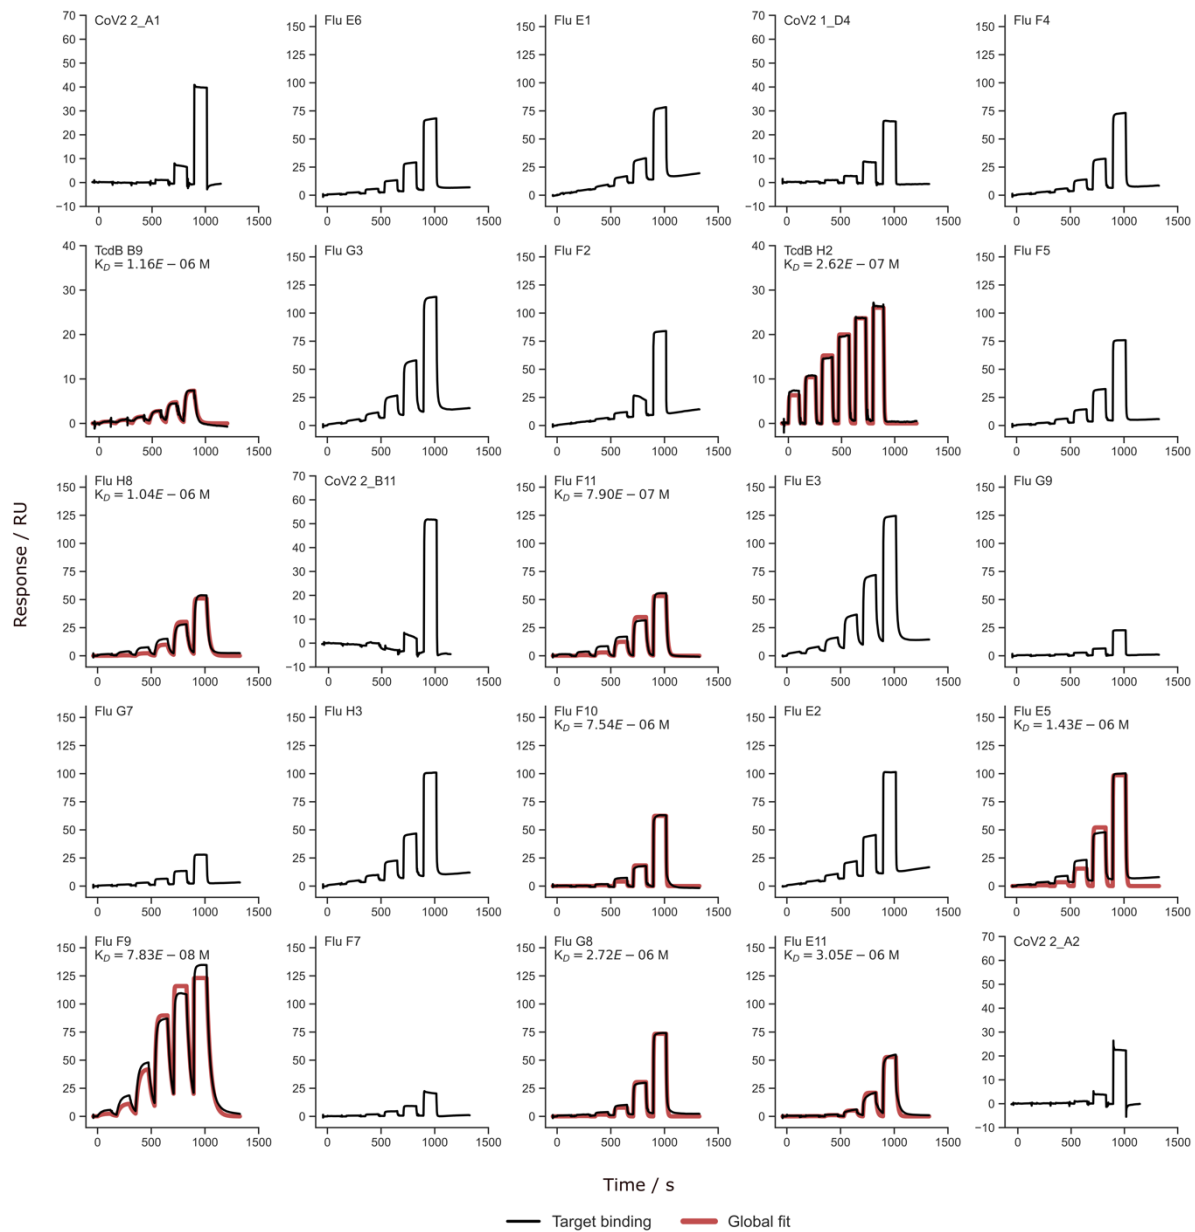
Extended Data Figure 8: SPR traces of experimentally validated VHHs

SPR traces of the experimentally validated VHH hits described in this study. For traces where confident Kd estimates could be fit, we display these on the figure panels. Designs TcdB H2 and Flu F9 are reproduced from Fig. 2.

Extended Data Figure 9: Negative-stain electron microscopy analysis of influenza HA antigens

A) Raw nsEM micrograph, **B**) 2D class averages showing a predominance of HA monomer species in the sample, and **C**) a representative predicted 3D model of this commercially produced monomeric HA antigen expressed in insect cells (adapted from PDB: 8SK7). This construct was used for screening VHH binders via yeast surface display and surface plasmon resonance. Insect-cell-produced glycoproteins exhibit a truncated glycan shield compared to those produced in mammalian cells. **D**) Raw nsEM micrograph, **E**) 2D class averages showing a clear abundance of HA trimers, and **F**) a representative 3D model of this in-house produced, trimeric Iowa43 HA antigen expressed in mammalian cells (adapted from PDB: 8SK7). This antigen is fully and natively glycosylated, and is the trimeric form of HA. Together these features make Iowa43 suitable for Cryo-EM structural studies of de novo designed VHHs and their capacity to bind to natively glycosylated glycoproteins of therapeutic interest.
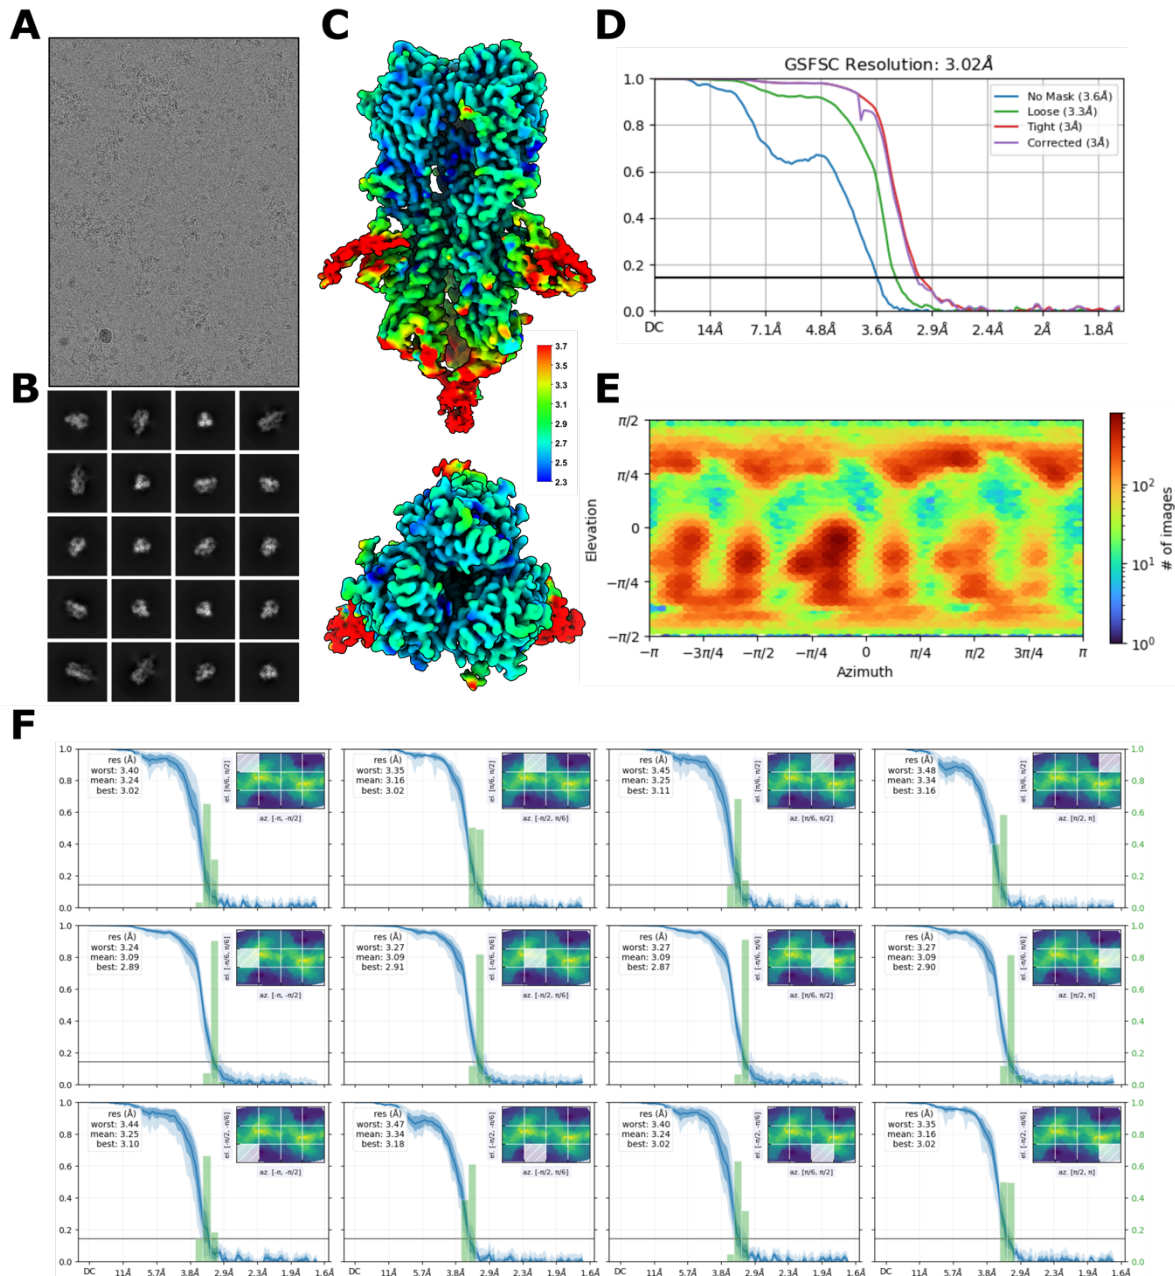
Extended Data Figure 10: Cryo-EM structure determination statistics for a de novo designed VHH bound to an influenza HA trimer

A) Representative raw micrograph showing ideal particle distribution and contrast. **B**) 2D Class averages of Influenza H1+designed VHH with clearly defined secondary structure elements and a full-sampling of particle view angles. **C**) Cryo-EM local resolution map calculated using an FSC value of 0.14 viewed along two different angles. Local resolution estimates range from ~2.3Å at the core of H1 to ~3.7Å along the periphery of the designed VHH. **D**) Global resolution estimation plot. **E**) Orientational distribution plot demonstrating complete angular sampling. **F**) Orientational diagnostics data.