

©Copyright 2026

Isaac Sappington

Expanding the Designable Landscape of Protein Interfaces
through Computational Binder Design

Isaac Sappington

A dissertation submitted
in partial fulfillment of the requirements
for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

David Baker, Chair

Ning Zheng

Frank Dimaio

Program Authorized to Offer Degree:

Department of Biochemistry

University of Washington

Abstract

Expanding the Designable Landscape of Protein Interfaces through Computational Binder Design

Isaac Sappington

Chair of the Supervisory Committee:

David Baker

Department of Biochemistry

Protein binders can be used to recognize, inhibit, activate, localize, or deliver molecules in biological systems. Antibodies perform these tasks naturally, but designing new binders from scratch remains difficult, especially when target surfaces are polar, flexible, recessed, or conformationally heterogeneous. In this work, we set out to expand the range of protein interfaces that can be targeted by computational binder design. We developed and tested deep-learning-based design strategies for generating and optimizing protein binders against challenging target surfaces. First, RFDiffusion interface conditioning was used to design binders that form β -strand-pairing interactions with target edge strands, enabling the design of high-affinity and specific binders to polar surfaces.. Second, AlphaFold2-based sequence optimization methods were developed to improve existing binder-target interfaces and to optimize binders across multiple related targets. These approaches produced experimentally validated binders and optimized variants across diverse systems, including receptors, toxins, fertilization proteins, immune targets, and viral antigens.

Table of Contents

Acknowledgments	1
General Introduction: Protein Folding, Function, and Design	2
0.1 Protein sequence determines protein structure and function	2
Chapter 1. Computational Design of Protein Binders	6
1.1 Why design protein binders, and how are they biologically functional?	7
1.2 How do we define and identify target-binding by protein binders?	9
1.3 How can we computationally design protein binders?	11
1.4 Example: EndoTags and pH-sensitive LYTAC binder design	13
1.5 RFdiffusion enables de novo backbone and binder generation	15
1.6 Designing binders against topologically complex protein targets	18
Chapter 2. Diffusion Models for Protein Scaffold Generation and Interface Conditioning	20
2.1 Diffusion model background for protein design	20
2.1.1 Why RFdiffusion trajectories often favor α -helical binders	21
2.1.2 Designing binders with edge β -strand interfaces to expand targetable surfaces	22
2.1.3 β -strand interface conditioning improves binder designs	23
2.1.4 Designing specific binders to BMP receptors ALK-2 and ALK-3	33
2.1.5 Targeting exposed edge β -strands in α -CTX, FCRL5, and NRP1	34
2.1.6 Targeting type III receptor tyrosine kinases KIT and PDGFR α	35
2.1.7 Stability of β -strand pairing binders	37
2.1.8 Specificity of β -strand pairing binders	38
2.1.9 KITbp-KIT crystal structure confirms design accuracy	41
2.1.10 Functional testing of FCRL5 and PDGFR α binders	45
2.1.11 Engineered ALK-2/ALK-3 binder agonists activate BMP signaling	49
2.1.12 Conclusions from the β -strand pairing binder study	50
References for Section 2.1	51
2.2 α -CTXbp and designed three-finger toxin binders for neutralizing snake venoms	58

	4
References for Section 2.2	61
2.3 Additional interface-conditioning tests: JUNO and GPCRs	61
2.3.1 JUNO fertilization receptor as a convex-interface design target	62
2.3.2 Binding and functional characterization of JUNO binders	63
References for Section 2.3	64
2.4 Concave-interface design for receptor-penetrating GPCR binders	65
References for Section 2.4	68
Chapter 3. Improving Methods for Binder Sequence Design	68
3.1 Why sequence optimization is needed after binder design	68
3.2 ProteinMPNN and partial diffusion for local binder optimization	69
3.3 AlphaFold2 backpropagation for direct sequence optimization	70
3.4 Interface-focused AF2 optimization from existing binder complexes	75
3.5 AF2-optimized interface mutations recapitulate SSM landscapes	78
3.6 Experimental validation of AF2-optimized binders	79
3.7 Multitarget AF2 optimization for cross-reactive binder design	84
3.8 Cross-reactive KIT optimization across mouse and human KIT	85
3.9 Functional characterization of KITopt binders	87
3.10 Cross-reactive HA stalk binder design optimization across H1, H3, and H5	89
References for Chapter 3	92
Chapter 4. Discussion	94
4.1 Designing binders to structurally heterogeneous targets	94
4.2 CD8 β binder design highlights flexible target-state uncertainty	94
4.3 TGF β R1/ALK5 binder design highlights loop-state uncertainty	95
4.4 Tau fibril binder design highlights polymorph and sample heterogeneity	96
4.5 Why computational design remains valuable despite difficult targets	99
4.6 Future applications of designed protein binders	100
References for Chapter 4	101

List of Figures

Figure 1.1.1. Proteins are polypeptides with unique sequences that fold into particular shapes based on energetic interactions.

Figure 1.1.2. The Protein Data Bank hosts standardized protein and biomolecule structural data that enables training of deep-learning models for protein structure prediction.

Figure 1.2.1. Proteins carry out functions that influence biological and molecular processes by surface recognition.

Figure 1.2.2. Design of bispecific IGF2R binders that induce an endocytotically active receptor conformation.

Figure 1.2.3. EndoTags and pH-sensitive LYTACs enable lysosomal degradation of soluble and membrane-associated targets.

Figure 1.2.4. De novo design of protein-binding proteins by RFdiffusion, ProteinMPNN, and AlphaFold2.

Figure 2.1.1. Design of β -strand pairing scaffolds with RFdiffusion.

Figure 2.1.2. Design of β -strand pairing binders.

Figure 2.1.3. Success rates from different RFdiffusion methods across several predictive in silico metrics.

Figure 2.1.4. Binding screens suggest preference for strand pairing binders at target sites.

Figure 2.1.5. Size-exclusion chromatography purification of β -strand pairing binders shows binders are largely monodisperse.

Figure 2.1.6. Binder design models and biophysical characterization.

Figure 2.1.7. The two PDGFR α binder design models are predicted to induce different PDGFR α conformations.

Figure 2.1.8. Designed binders are highly specific for their targets.

Figure 2.1.9. Analysis of strand pairing binder interface polarity in context of previously published de novo binders.

Figure 2.1.10. Structural analysis of the KITbp-KIT complex.

Figure 2.1.11. Surface hydrophobicity representations of the KITbp interface.

Figure 2.1.12. Functional activity of FCRL5 and PDGFR α binding proteins.

Figure 2.1.13. Gating strategy used for quantifying surface EGFR levels.

Figure 2.1.14. Agonism of BMP signaling pathways by ALK-2bp and ALK-3bp fusion to ActRIIAbp.

Figure 2.2.1. De novo designed protein binders neutralize three-finger snake venom toxins.

Figure 2.3.1. De novo designed binders target the JUNO fertilization receptor.

Figure 2.4.1. Computational design and characterization of receptor-penetrating GPCR miniprotein binders.

Figure 3.1.1. AlphaFold backpropagation enables protein design in relaxed sequence space.

Figure 3.1.2. Comparison of de novo AF2-backpropagation binder design and interface-focused binder optimization.

Figure 3.1.3. AF2-based interface optimization recovers experimentally enriched SSM mutations.

Figure 3.1.4. Interface-focused AF2 optimization improves binder affinity and functional inhibition.

Figure 3.1.5. Benchmarking computational workflows for affinity maturation of an NKG2D binder.

Figure 3.1.6. Cross-reactive AF2-based interface optimization using shared binder sequence logits.

Figure 3.1.7. Cross-reactive optimization of a human KIT binder for mouse and human KIT binding.

Figure 3.1.8. Functional characterization of cross-reactive KITopt binders.

Figure 3.1.9. Cross-reactive optimization of an HA stalk binder across H1, H3, and H5.

Figure 4.1.1. Binder design challenges arising from target conformational uncertainty and structural heterogeneity.

List of Tables

Table 2.1.1. Strand-pairing binders show improved yeast surface display binding (YSD) success rates.

Table 2.1.2. Crystallographic data collection and refinement statistics for the determined structure of the KITbp-KIT123 complex

Acknowledgments

I would like to thank my advisor, David Baker, for mentoring me during my PhD. His continual investment in my research interests and projects was extremely motivating. Additionally, he connected me with numerous collaborators, and he is really the primary individual actor that integrates each member of his lab into a larger “communal brain”. I am exceedingly grateful for the connections I’ve made in the lab that has simultaneously made each experiment easier and more impactful. This lab cohesiveness results in greater accomplishments that regularly advance the entire field of computational protein design. David’s leadership somehow balances levity with rigor in a way that makes the lab a tremendously fun and fulfilling environment for science.

I received a great deal of training and advice from several mentors who volunteered their time and knowledge when I needed it. Namely, Brian Coventry, Buwei Huang, and Xinru Wang taught me a great deal about protein design and biochemistry that provided a foundation for my research. Many other colleagues and friends provided me with invaluable advice and support through grad school: Edin Muratspahić, Xiangli Qu, David Kim, Joe Watson, Wei Yang, Caixuan Liu, Derrick Hicks, Susana Vázquez Torres, David Lee, Clara McCurdy, Adam Chazin-Gray, Jeremiah Sims, Preetham Venkatesh, Jake Gershon, and so many others. I also give thanks to the staff at the Institute for Protein Design who contributed so much experimentally and organizationally to my work, especially Inna Goreschnik, Dionne Vafeados, Brian Belmont, and Luki Goldschmidt.

Finally, my family and friends encouraged me over many years to be curious, pursue studies in biochemistry, and enjoy it all. My family (both Sappingtons and Vezies), Jeremiah Jurich, Samuel Curtis, Clara Merlino, and my wife Karen were all hugely influential in the 30+ years that led to this accomplishment. In the grand scheme of things, they each deserve a bit of credit.

General Introduction: Protein Folding, Function, and Design

Proteins are the primary functional molecules for biological systems, acting as molecular machines that perform tasks in cells. The two fundamental protein functions are enzymatic catalysis and binding to other molecular surfaces. These two functions are highly variable in nature, catalyzing thousands of unique chemical reactions and many more specific surface binding interactions to specifically recognize other molecules. Millions of protein molecules perform various tasks in concert in response to their molecular environment, combining to drive almost every core process in the cell (such as cellular transportation, communication, and maintenance). These complex cellular processes then give rise to tissue and organ systems which underpin our lived experience. As such, leveraging protein functions holds the potential for improving human health at a fundamental level. This dissertation explores the computational design of functional binding proteins, a relatively nascent field of science and engineering with the goal of creating new proteins that can be used to direct biological functions.

0.1 Protein sequence determines protein structure and function

Protein functions are governed by their surfaces, which perform functions by spontaneously binding to shape- and chemically-complementary molecular surfaces. Because a protein's function is tightly coupled to its surface complementarity with other molecules, the features that determine a specific protein's structure and shape are what makes it functionally unique. Here, we will explore how a protein's sequence is the fundamental feature that determines how a protein folds into its structure.

Each protein has a sequence that determines its structural shape and resulting function. This sequence refers to a specific amino acid sequence of 20 canonical amino acids ordered in a polymeric string by peptide bonds (Figure 1.1.1a). These polypeptides can range in size from very small (such as aspartame, the 2 amino-acid "dipeptide") to thousands of amino acids long.

This combination of protein length and amino acid possibility (Figure 1.1.1b) at each position results in the vastness of possible amino acid sequence space. For a protein of length L , the number of possible amino acid sequences is 20^L . Thus, for $L=300$, there are approximately $20^{300} \approx 2 \times 10^{390}$ possible sequences.

Furthermore, each unique polypeptide sequence folds into a very precise 3-dimensional structure mainly as a result of the energetic contributions of many non-covalent interactions between the flexible polypeptide backbone, rotatable amino acid side chains, and surrounding solvent molecules¹ (Figure 1.1.1c-d). These non-covalent interactions include hydrophobic effects, positive and negative charge complementarity, and hydrogen bonding. While we have an understanding of the relative energetics of these individual interaction types, calculating a protein's fold by physical modeling remains intractable, largely due to the dual nature of polypeptides which are both flexible in adopting many intermediate structures during folding while ultimately forming highly stable and precise final structures. As such, designing proteins from scratch with user-defined folds and functions remained a longstanding challenge. For functional proteins, structural features often require angstrom level precision (within one 10 billionth of a meter) of design models, which was not achievable by physics-based methods implemented by early protein designers.




Figure 1.1.1. Proteins are polypeptides with unique sequences that fold into particular shapes based on energetic interactions. (a) A polypeptide bond between amino acid residues has backbone bond angles (ϕ , ψ , and ω) that define a protein fold. (b) Each residue is identifiable by one of 20 canonical amino acid side chains that have unique chemical properties and interaction profiles. (c) The en masse interactions of a polypeptide's residues contribute to spontaneous folding into specific, low-energy shapes (Image credit: Ian Haydon). (d) The fundamental non-covalent interactions between chemical moieties of residues, as well as the interactions with aqueous solvent in biological systems, energetically drive protein folding.

Historically, understanding protein folding relied heavily on experimental determination of protein structures. Since Kendrew and colleagues reported the first low-resolution three-dimensional model of myoglobin in 1958², structural biologists have deposited more than 250,000 experimentally determined biomolecular structures into the Protein Data Bank (PDB)³. This vast resource of high quality data served as the foundation for studying the fundamental principles of protein folding (Figure 1.1.2a). Recently, new statistical deep-learning methods that efficiently discover patterns that generalize across massive datasets were applied to encode the patterns of association between protein sequences and structures present in the PDB (Figure 1.1.2b). Deep-learning structure-prediction models, especially AlphaFold2, substantially improved the

accuracy of protein structure prediction and helped enable a new generation of computational protein design methods⁴.




Figure 1.1.2. The Protein Data Bank hosts standardized protein and biomolecule structural data that enables training of deep-learning models for protein structure prediction. (a) Protein structures are visualized and studied by rendering their three dimensional coordinates. Representative “Molecules of the Month” illustrations prepared by Janet Iwasa, David Goodsell, Stephen K. Burley, and Christine Zardecki for PDB-101 (PDB101.rcsb.org). **(b)** Structures for protein sequences are determined by purification, experimental measurement, and 3D modeling of the protein sequence to match the observed measurements. All reputable scientific journals require deposition of structural data and models into the PDB, providing the highest quality dataset for training deep-learning models on the relationships between protein sequences and their folded structures. Illustration by Maria Voigt for PDB-101 (PDB101.rcsb.org).

This thesis describes my research in using these structure prediction models and other deep learning methods for the design and characterization of proteins that specifically bind the surfaces of intended target proteins. Chapter 1 will describe the rationale, theoretical principles, and standard computational approaches for protein binder design and characterization. Chapter 2 explores the design of protein binder backbones with optimal interface complementarity to arbitrary targets, including examples for designing binders that match target edge strands,

convex/concave surface, and flexible motifs. Chapter 3 addresses methods for sequence design of binders to computationally optimize towards desired binder functions. Finally, Chapter 4 summarizes findings from exploring the limits of protein-protein interface design. Each chapter will include an overview of relevant methods, as well as experimental results from various binder design campaigns.

1. Anfinsen CB. Principles that govern the folding of protein chains. *Science*. 1973;181:223-230. doi:10.1126/science.181.4096.223.
2. Kendrew JC et al. A three-dimensional model of the myoglobin molecule obtained by X-ray analysis. *Nature*. 1958;181:662-666. doi:10.1038/181662a0.
3. Berman HM et al. The Protein Data Bank. *Nucleic Acids Research*. 2000;28:235-242. doi:10.1093/nar/28.1.235.
4. Jumper J et al. Highly accurate protein structure prediction with AlphaFold. *Nature*. 2021;596:583-589. doi:10.1038/s41586-021-03819-2.

Chapter 1.

Computational Design of Protein Binders

This chapter introduces the core logic of modern protein binder design as it is typically practiced in computational structural biology: (1) generate structural hypotheses for how a binder might engage a target surface, (2) choose sequences that are compatible with those structures, (3) filter designs using structure-prediction oracles, and (4) validate and improve the best candidates experimentally. This design logic has existed for decades, but early successes relied heavily on testing large numbers of computational designs experimentally¹. More recently, deep-learning methods such as AlphaFold2 improved the ability to predict whether a designed protein sequence will fold into a desired functional shape, enabling more reliable computational filtering before experimental testing^{2,3}. Later in this chapter, I describe how advances in structure

prediction influenced similar advances in scaffold generation and sequence design. First, however, it is useful to define why protein binders are valuable, how their binding is measured, and how modern design workflows identify candidates likely to bind as intended.

1.1 Why design protein binders, and how are they biologically functional?

Binding is one of biology's core "verbs," and binders can be developed into therapeutics, diagnostics, and research tools. A binder is a folded protein that recognizes a specific molecular surface with defined specificity and affinity. This ability to spontaneously find and attach to target surfaces makes binders useful for altering the functions of target molecules, often other proteins. Binding can block a target surface from engaging its natural interaction partners, stabilize a target protein in a particular active or inactive conformation, or bring two molecules into proximity to spatially couple their functions (Figure 1.2.1). These principles underlie many binder-based technologies, including receptor agonists, antagonists, bispecific proteins, and targeted degradation strategies⁴. Even enzymes can be viewed as specialized binders that recognize substrates, transition states, and products. When multiple recognition events are connected together, proteins can assemble into multi-component machines or signaling pathways that transmit information through biological systems.




Figure 1.2.1. Proteins carry out functions that influence biological and molecular processes by surface recognition. (a) A functional protein (blue) can block the interaction of a target protein (gray) with its native interacting partner ligands (orange). (b) Proteins can bind to multi-conformational proteins to influence their conformational dynamics towards a particular state. (c) Proteins can enzymatically catalyze chemical reactions by binding and positioning a substrate molecule into a transition-state conformation, enabling chemical modification of the substrate into a new molecular product. (d) A single protein can act as a molecular machine that spatially localizes multiple functional domains.

In addition to their therapeutic applications, specific and high-affinity binders are broadly useful as affinity reagents for diagnostics, proteomic analysis, purification, imaging, and research tools⁵. Binders can also be immobilized on affinity matrices to purify targets from complex mixtures, or used to perturb signaling pathways in living systems. Overall, these applications make functional binding proteins valuable both as therapeutic agents and as tools for discovering biological mechanisms.


A core aspect of this thesis is that computationally designed binders have distinct advantages over experimentally optimizing naturally occurring binding proteins. By pre-screening sequences in silico, computational design can reduce dependence on very large high-throughput experimental screens. This is especially useful for developing binders against integral

membrane proteins or conformationally sensitive targets, which can be difficult to purify and screen at large scale. Computationally designed proteins can also be engineered for stability, small size, bacterial expression, and compatibility with modular fusion architectures. Finally, computational methods can incorporate multiple design goals simultaneously, such as targeting a specific epitope, limiting protease-sensitive sites, controlling net charge, or scaffolding multiple functional motifs within a single protein.

1.2 How do we define and identify target-binding by protein binders?

Determination of a protein's binding tightness, or affinity, involves the measurement of multiple individual binding properties. In a simple two-state binding model, the interface dissociation constant (K_D) reflects the balance between bound and unbound populations at equilibrium, and it is related to binding free energy (ΔG) by standard thermodynamic relationships. In addition to ΔG , the K_D can be described in terms of the rates of binding and unbinding. The association rate k_{on} and dissociation rate k_{off} , describe how quickly binding occurs and how long complexes persist. k_{on} can be limited by diffusion, electrostatic effects, and conformational selectivity; whereas k_{off} may reflect the total energy of interactions that must break simultaneously to unbind. In many settings (especially pharmacology), slow dissociation (small k_{off}) can be as important as low K_D because it increases residence time on the target.

Binding affinity describes the equilibrium between bound and unbound states. For a simple two-state interaction between binder B and Target T,



the equilibrium dissociation constant is:

$$K_D = [B][T]/[BT]$$

where lower K_D values indicate tighter binding. The same interaction can also be described kinetically by the association rate k_{on} and dissociation rate k_{off} :

$$K_D = k_{off}/k_{on}$$

Thus, tight binding can result from rapid association, slow dissociation, or both. In many biological settings, slow dissociation is especially important because it increases residence time on the target.

Binding affinity is related to standard binding free energy by:

$$\Delta G^\circ = RT \ln(K_D/C^\circ)$$

where R is the gas constant, T is temperature, and C° is the 1 M standard concentration. At room temperature, a 10-fold change in K_D corresponds to approximately 1.36 kcal/mol in binding free energy. This provides useful intuition for protein binder design: improving affinity often requires only a small number of favorable atomic interactions, but those interactions must compensate for conformational restriction, desolvation, and any unfavorable contacts formed during binding.

Because k_{on} and k_{off} often determine a binder's biological behavior, the most informative characterization methods directly measure binding and dissociation over time. Surface plasmon resonance (SPR) and biolayer interferometry (BLI) were used to characterize many of the binder affinities presented in this dissertation. Both methods involve immobilizing one binding partner on a sensor surface, exposing that surface to titrated concentrations of the other binding

partner, and monitoring a real-time optical signal proportional to binding. SPR detects binding-associated changes in refractive index near a sensor surface, whereas BLI measures changes in optical interference from light reflected at a biosensor tip. Both approaches can provide real-time association and dissociation traces that are fit to estimate k_{on} , k_{off} , and K_D ^{6,7}.

1.3 How can we computationally design protein binders?

We have addressed the motivations for computational protein design, introduced methods for determining binding kinetics and affinity, and discussed how binding affinity arises from ensembles of molecular interactions. This final background section describes the computational approaches used to identify binding sites, generate binder structures, design binder sequences, and filter designs with structure-prediction oracles.

In 2022, Cao et al. described a generalizable Rosetta-based approach for designing de novo protein binders from target structure alone¹. In this workflow, scaffold libraries were docked against a target binding site, followed by multiple rounds of interface sequence design, Rosetta scoring, and docking refinement. Experimentally optimized binders from this approach reached nanomolar to picomolar affinities and, in several cases, showed functional effects such as blocking receptor engagement¹. This work established important design principles that remain relevant: successful binders are most likely when a designed scaffold makes extensive shape-complementary contacts with a target surface, especially at exposed hydrophobic or otherwise energetically favorable patches.

AlphaFold2 then enabled accurate structure prediction from sequence information alone, accelerating the use of structure-prediction models as design filters². Along with predicted

coordinates, AlphaFold2 outputs confidence metrics such as predicted local distance difference test (pLDDT), which estimates local residue-level confidence, and predicted aligned error (pAE), which estimates uncertainty in the relative placement of residue pairs or domains. For de novo binder design, Bennett et al. showed that AlphaFold2 and RoseTTAFold metrics could identify two major design failure modes: failure of the designed sequence to fold into the intended monomer structure, and failure of the binder to form the intended complex with the target. In particular, pAE_interaction strongly discriminated experimentally successful binders from non-binders, and AF2 “initial guess” enabled complex evaluation using a provided starting design model³. AlphaFold-Multimer extended this framework to protein complexes, enabling prediction of binder-target assemblies⁸. More recent models, including AlphaFold3, further expanded structure prediction to general biomolecular complexes containing proteins, nucleic acids, small molecules, ions, and modified residues⁹.

Deep-learning models trained on the Protein Data Bank were also adapted for generative design tasks, including the creation of new protein folds and scaffolds. Early approaches inverted structure-prediction networks to generate new proteins by “hallucination,” in which sequences were optimized to produce confident predicted structures¹⁰. Related inpainting approaches started from partially specified sequence or structure information and filled in missing regions around fixed functional motifs¹¹. For sequence design, ProteinMPNN was trained to recover plausible amino acid identities from protein backbone geometry, enabling rapid sequence design for generated scaffolds¹². Together with RFdiffusion for backbone generation and AlphaFold2 for filtering, these methods enabled practical workflows for designing plausible protein binder folds and sequences against arbitrary target structures^{10,11,12}.

1.4 Example: EndoTags and pH-sensitive LYTAC binder design

This subsection summarizes the design and validation of IGF2R-targeting EndoTags, especially EndoTag3/IRigid1 and EndoTag4, as reported by Huang et al.¹³ It also includes selected pH-sensitive LYTAC data from Ahn et al. to illustrate how pH-independent lysosomal receptor binding can be combined with pH-triggered target release to enable substoichiometric extracellular protein degradation¹⁴.

This approach was used to scaffold multiple binding motifs against insulin-like growth factor 2 receptor (IGF2R), with the goal of inducing receptor endocytosis. Rigid IGF2R EndoTags, including IRigid1/EndoTag3, were designed by scaffolding and fusing binding motifs from two previously designed binders to IGF2R domains 6 and 11 into a single protein that approximates the receptor geometry induced by native IGF2 binding (Figure 1.2.2a,b). In these designs, binding helices from the individual domain binders were treated as partial motifs, and RFInpainting was used to connect them into single scaffolds predicted to enforce productive receptor conformations. These designs demonstrated that computationally generated binders could lock receptors into functional conformations by bispecifically binding domains 6 and 11 simultaneously (Figure 1.2.2c). Fusion of IRigid1/EndoTag3 to other target-binding proteins produced protein-based lysosome-targeting chimeras (pLYTACs) that promoted uptake and lysosomal degradation of soluble or membrane-associated targets (Figure 1.2.3a–c)^{4,13}.




Figure 1.2.2. Design of bispecific IGF2R binders that induce an endocytotically active receptor conformation. (a) Schematic of the IGF2R EndoTag design goal, in which bispecific binding approximates the native IGF2-induced arrangement of IGF2R domains 6 and 11 to promote receptor-mediated endocytosis. (b) Design of EndoTag3/IRigid1 and EndoTag4. From top to bottom: domains 6 and 11 of IGF2R extracted from the IGF2:IGF2R complex structure; domain 6 and domain 11 binder design models aligned to the active receptor conformation; minimal binding motifs extracted from the individual binder models; RFInpainting used to combine domain 6 and domain 11 interface motifs into one binder, followed by ProteinMPNN redesign and AlphaFold2 prediction. (c) Binding kinetics for EndoTag3 and EndoTag4 measured against IGF2R domain 6 and domain 11 by BLI. EndoTag3 maintains binding at pH 5.5, consistent with retention of IGF2R binding under acidic endosomal/lysosomal conditions. Adapted from Huang et al.¹³

EndoTag3/IRigid1 was designed to remain bound to IGF2R in acidic endosomal and lysosomal environments. This pH-independent receptor binding enabled EndoTag constructs to remain associated with IGF2R during endosomal trafficking and receptor recycling. In later pH-sensitive LYTAC architectures, IRigid1 was fused to binders engineered to release their target proteins at acidic pH, enabling the degrader to recycle with the lysosomal trafficking receptor while releasing the target cargo for degradation¹⁴ (Figure 1.2.3d-f).




Figure 1.2.3. EndoTags and pH-sensitive LYTACs enable lysosomal degradation of soluble and membrane-associated targets. (a) Schematic of pLYTAC designs for membrane-associated and soluble target degradation using an EndoTag lysosomal trafficking receptor binder fused to a protein-of-interest binder. (b) Western blot analysis of total EGFR in wild-type or IGF2R-knockout HeLa cells after treatment with EGFR-targeting constructs fused to IGF2R EndoTags. (c) Depletion and lysosomal colocalization of IgG-AF647 following treatment with protein G-IGF_EndoTag3. Left: remaining IgG-AF647 in Jurkat cell supernatant after treatment. Middle: confocal imaging of IgG-AF647 colocalization with lysosome in HeLa cells. Right: control imaging in IGF2R-knockout HeLa cells. (d) Schematic of pH-sensitive LYTAC architectures that combine pH-independent lysosomal trafficking receptor binding with pH-triggered target release, enabling degrader recycling and multi-turnover target capture. (e) Western blot analysis of total EphA2 in HCT116 and MDA-MB-231 cells after treatment with designed proteins. (f) Internalization of biotinylated IL-6 precomplexed with AF647 streptavidin after treatment with LYTACs at substoichiometric ratios. Panels a-c adapted from Huang et al.; panels d-f adapted from Ahn et al.^{14,15}

1.5 RFdiffusion enables de novo backbone and binder generation

This subsection summarizes RFdiffusion-based binder design results reported by Watson et al.;

Figure 1.2.4 reproduces and rearranges selected panels from that work¹⁵.

The next major advance in protein scaffold generation was RFdiffusion, a denoising diffusion model for generating protein backbones conditioned on user-specified structural or functional constraints¹⁵. RFdiffusion builds on the RosettaFold network architecture and generates protein

structures by starting from noised residue-frame coordinates, then iteratively denoising them toward protein-like backbones that satisfy the provided conditioning information (Figure 1.2.4a). When combined with ProteinMPNN for sequence design and AlphaFold2 for filtering, RFdiffusion enabled rapid generation and triage of plausible binder designs against target structures^{2,12,15}.

Prior to these deep-learning methods, many thousands of binder designs were often required to identify a true target-binding hit, and design campaigns were most reliable when accurate target structures were available. RFdiffusion, ProteinMPNN, and AlphaFold2 substantially improved this workflow by enabling rapid generation, sequence design, and filtering of candidate binders. This improvement was demonstrated across targets including influenza A H1 hemagglutinin, IL-7R α , insulin receptor, PD-L1, and TrkA, where experimentally validated binders were obtained from comparatively small design sets¹⁵ (Figure 1.2.4b-h).




Figure 1.2.4. De novo design of protein-binding proteins by RFdiffusion, ProteinMPNN, and AlphaFold2. (a) RFdiffusion generates protein binders given a target and specification of interface hotspot residues. (b) De novo binders were designed to five protein targets: influenza A H1 hemagglutinin, IL-7Ra, insulin receptor, PD-L1, and TrkA. Hits with BLI response greater than or equal to 50% of the positive control were identified for all targets. For IL-7Ra, insulin receptor, PD-L1, and TrkA, RFdiffusion had success rates roughly two orders of magnitude higher than the original design campaigns. (c) For IL-7Ra, insulin receptor, PD-L1, and TrkA, the highest-affinity binder is shown above a BLI titration series. Reported K_D values are based on global kinetic fitting with fixed global R_{max} . (d) The highest-affinity HA binder, HA₂₀, binds with a K_D of 28 nM. In c,d, target or hotspot residues are shown in yellow or orange, design models in gray, and AF2 predictions in purple. (e) Cryo-EM 2D class averages of HA₂₀ bound to influenza HA, strain A/USA:lowa/1943 H1N1. (f) Cryo-EM 3D reconstruction of the HA₂₀-HA complex. (g) The cryo-EM structure of the HA₂₀ complex closely matches the design model. (h) Structure of the HA₂₀ binder alone superimposed on the design model. Adapted from Watson et al.¹⁵

1.6 Designing binders against topologically complex protein targets

Ultimately, these models did not solve all protein design problems by simply being run out of the box. Rather, iterative methodological improvements have been needed to apply them to increasingly difficult and specialized design challenges. This thesis describes several such improvements, focused on expanding the range of target surfaces and optimization objectives accessible to computational binder design. The following chapters explore these limits from both scaffold-design and sequence-optimization perspectives.

1. Cao L, Coventry B, Goreshnik I, et al. Design of protein-binding proteins from the target structure alone. *Nature*. 2022;605:551–560. doi:10.1038/s41586-022-04654-9.
2. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*. 2021;596:583–589. doi:10.1038/s41586-021-03819-2.
3. Bennett NR, Coventry B, Goreshnik I, et al. Improving de novo protein binder design with deep learning. *Nature Communications*. 2023;14:2625. doi:10.1038/s41467-023-38328-5.
4. Banik SM, Pedram K, Wisnovsky S, et al. Lysosome-targeting chimaeras for degradation of extracellular proteins. *Nature*. 2020;584:291–297. doi:10.1038/s41586-020-2545-9.
5. Taussig MJ, Stoevesandt O, Borrebaeck CAK, et al. ProteomeBinders: planning a European resource of affinity reagents for analysis of the human proteome. *Nature Methods*. 2007;4:13–17. doi:10.1038/nmeth0107-13.
6. Myszka DG. Kinetic analysis of macromolecular interactions using surface plasmon resonance biosensors. *Current Opinion in Biotechnology*. 1997;8:50–57. doi:10.1016/S0958-1669(97)80157-7.

7. Abdiche Y, Malashock D, Pinkerton A, Pons J. Determining kinetics and affinities of protein interactions using a parallel real-time label-free biosensor, the Octet. *Analytical Biochemistry*. 2008;377:209–217. doi:10.1016/j.ab.2008.03.035.
8. Evans R, O'Neill M, Pritzel A, et al. Protein complex prediction with AlphaFold-Multimer. *bioRxiv*. 2022. doi:10.1101/2021.10.04.463034.
9. Abramson J, Adler J, Dunger J, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*. 2024;630:493–500. doi:10.1038/s41586-024-07487-w.
10. Anishchenko I, Pellock SJ, Chidyausiku TM, et al. De novo protein design by deep network hallucination. *Nature*. 2021;600:547–552. doi:10.1038/s41586-021-04184-w.
11. Wang J, Lisanza S, Juergens D, et al. Scaffolding protein functional sites using deep learning. *Science*. 2022;377:387–394. doi:10.1126/science.abn2100.
12. Dauparas J, Anishchenko I, Bennett N, et al. Robust deep learning-based protein sequence design using ProteinMPNN. *Science*. 2022;378:49–56. doi:10.1126/science.add2187.
13. Huang B, Abedi M, Ahn G, Coventry B, Sappington I, et al. Designed endocytosis-inducing proteins degrade targets and amplify signals. *Nature*. 2025;638:796–805. doi:10.1038/s41586-024-07948-2.
14. Ahn G, Coventry B, Haefner E, Sadre S, Hu J, Van M, Huang B, Sappington I, et al. Computational design of pH-sensitive binders. *bioRxiv*. 2025. doi:10.1101/2025.09.29.678932.
15. Watson JL, Juergens D, Bennett NR, et al. De novo design of protein structure and function with RFdiffusion. *Nature*. 2023;620:1089–1100. doi:10.1038/s41586-023-06415-8.

Chapter 2.

Diffusion Models for Protein Scaffold Generation and Interface Conditioning

2.1 Diffusion model background for protein design

Denoising diffusion probabilistic models (DDPMs) generate complex data by learning to reverse a gradual noising process. In standard diffusion notation, a clean structure x_0 is progressively corrupted through a Markov chain $q(x_{1:T}|x_0)$, where each noised state depends only on the previous timestep. The amount of noise added at each timestep is controlled by a noise

schedule $\{\beta_t\}_{t=1}^T$. If $\alpha_t = 1 - \beta_t$, then $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ describes the cumulative fraction of signal retained after t noising steps. In simplified coordinate notation, the noised structure at timestep t can be written as:

$$x_t = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$$

where $\epsilon \sim \mathcal{N}(0, I)$. This equation describes the conceptual DDPM noising process; RFDiffusion applies this general idea to protein residue-frame representations rather than simply diffusing raw Cartesian coordinates².

The reverse process learns to invert this corruption by denoising a noisy structure toward a clean protein-like structure. Conceptually, the neural network denoiser can be written as:

$$\hat{x}_\theta^{(0)} = f_\theta \left(x^{(t)}, t, c \right)$$

where $x(t)$ is the noised structure at timestep t , c represents conditioning information, and $\hat{x}_\theta^{(0)}$ is the model's prediction of the denoised protein structure. Equivalently, the denoising model can be viewed probabilistically as $p_\theta(x_0 | x_t, t, c)$, a learned distribution over clean structures x_0 given a noisy input x_t , timestep t , and conditioning information c . In score-based terms, the model learns a vector field related to $\nabla_x \log p(x | c)$, which points noisy structures toward

regions of structure space that are more protein-like and more consistent with the conditioning information.

In RFdiffusion, the conditioning variable c can include partial sequence information, fixed structural motifs, symmetry constraints, target hotspot residues, or secondary-structure and adjacency specifications (SS/ADJ). These conditioning features are provided during denoising and influence each structural update, so that the final generated backbone is not only protein-like, but also consistent with the user-specified design constraints².

2.1.1 Why RFdiffusion trajectories often favor α -helical binders

RFdiffusion greatly improved de novo protein scaffold generation, but in practical binder-design campaigns, unconstrained or sparsely constrained trajectories often favor compact α -helical scaffolds. The exact origin of this bias is not known, but a plausible explanation is that α -helices are locally defined by near-neighbor backbone geometry and repeated hydrogen-bonding patterns, making them comparatively easy to recover early during denoising. In contrast, β -sheets require coordinated long-range residue pairings, strand register, and global topology. If the model is given only sparse hotspot constraints, early locally helical solutions may therefore be easier to generate than β -sheet-rich folds or extended strand-pairing interfaces.

This practical bias matters because many protein target surfaces are not ideally complemented by helical binders. Hydrophilic edge β -strands, for example, expose backbone hydrogen-bond donors and acceptors that must be precisely matched to compensate for desolvation. A designed binder that forms an extended β -strand pair with the target could satisfy these polar atoms in a sequence-independent manner, but such interfaces are unlikely to emerge reliably from hotspot conditioning alone. This motivated the development of explicit interface

conditioning strategies that bias RFDiffusion trajectories toward binder-target β -strand pairing before sequence design and structure-prediction filtering.

2.1.2 Designing binders with edge β -strand interfaces to expand targetable surfaces

Despite advances in deep learning methods, design of high affinity binders to hydrophilic regions of a target surface remains challenging since the exposed hydrogen bond donors and acceptors must be complemented with precisely positioned acceptors and donors on the designed binder to compensate for the loss of interactions with water. While RFDiffusion excels at generating backbones that are shape complementary to the targeted region of a protein surface, these solutions do not always provide detailed complementarity of hydrogen bonding donors and acceptors⁷. Many therapeutically relevant target proteins have β -sheets with unpaired and exposed β -strands; these often have non-canonical structures to reduce the tendency for aggregation^{8,9}. β -strand targeted binders have been generated using pre-deep learning Rosetta methodology¹⁰, but this approach has limited ability to match the diversity of edge β -strand geometries. A generalizable deep learning method for designing β -strand pairing based binders to complement arbitrary target β -strand twists¹¹, bends¹², bulges¹³, and other irregular features^{14,15} could improve the design of binders to polar target surfaces.

Here, we demonstrate that RFDiffusion generation of binder backbone mediated hydrogen bonding interfaces with the target can yield designs with polar groups nearly perfectly complementing those on the target. We develop a general approach for guiding RFDiffusion denoising-diffusion trajectories towards such β -strand centric interfaces, and we experimentally validate this binder design approach.

2.1.3 β -strand interface conditioning improves binder designs

During training of RFdiffusion, a subset of training examples were provided with secondary structure and secondary structure block-adjacency (SS/ADJ) information; this was found to enable the conditioning of the model towards user defined protein monomer folds (as previously described by Watson, et al.)². We explored whether providing interface conditioning information indicating a binder β -strand pairing with a target edge β -strand at the point of inference could yield designable strand pairing complexes. RFdiffusion takes as input desired residue secondary structures as well as an $N \times N$ block-adjacency matrix block-adjacency matrix (N is the length of the designed protein) which specifies desired adjacencies between secondary structure blocks. These conditioning tensors can specify binder residues having desired secondary structure identities (helix, strand, loop, or masked) as well as the binder residue interaction identities (interacting, non-interacting, masked) with other binder residues (used for fold conditioning), or with target residues (used here for interface conditioning). Given the length of the designed interacting β -strand (L) and the identity of the target β -strand to be bound (T), RFdiffusion generates binders assigning a random set of L consecutive binder residues to form a β -strand pair with the target strand T , while the remainder of the binder output residues are not explicitly assigned a secondary structure or target interaction. We found that outputs from conditioned β -strand pairing runs indeed contained the user-specified β -strand interfaces (Figure 2.1.1, Figure 2.1.2a-b).




Figure 2.1.1. Design of β -strand pairing scaffolds with RFdiffusion. (a) To leverage the diverse generative potential of RFdiffusion while encouraging β -strand interfaces in design outputs, we implemented an interface conditioning algorithm that generates SS/ADJ conditioning tensors given simple user inputs. The model understands fold conditioning in the form of tensors that label the secondary structures (blue) of each residue (a, top and left), and the adjacencies of these secondary structure blocks (a, yellow center). User-specified parameters specify information about the binder-interfacial secondary structure block (a β -strand in this case), the length of this block (b, cyan block in binder tensor L), and the target site residues to which the binder block is adjacent (b, cyan block in target tensor T). From these predefined parameters, the algorithm randomly samples the location of the binder-interfacial secondary structure block in residue index space, while maintaining a defined adjacency (green) to the specified target site residues. This user-defined conditioning tensor guides diffusion outputs towards a β -strand pairing binder-target interface (c). Previously, RFdiffusion interface design calculations could be directed to specific residues specified as target “hotspots” to specify the intended target residues to bind to, this new inter-chain SS/ADJ conditioning enables user indication of “ β -strand hotspots” or “ α -helical hotspots” during binder scaffold generation. Binder scaffold outputs conditioned on expanded binder-target SS/ADJ tensors enabled design of user-specified β -strand interface types.

We supplied RFdiffusion with β -strand interface conditioning tensors for larger scale binder design campaigns against protein targets containing edge β -strands sufficiently exposed for binder access. To evaluate the generality of the approach, we selected targets that span a range of edge β -strand geometries. The seven selected targets also have considerable therapeutic relevance. Activin receptor-like kinases 2 and 3 (ALK-2 and ALK-3) are both Type I

bone morphogenetic protein receptors that regulate growth of bone, vasculature, hair follicles, enamel as well as wound healing and tumor suppression in various soft tissues¹⁶⁻²⁰.

Platelet-derived growth factor receptor α (PDGFR α) and Mast/stem-cell factor receptor (also known as KIT, SCFR, CD117) are both type III receptor tyrosine kinases that play roles in cardiomyocyte proliferation and heart tissue regeneration after myocardial infarction²¹⁻²⁹. FCRL5 functions in critical signaling pathways for B-cell activation³⁰. NRP1 is a coreceptor for various growth factor signaling pathways (TGF-1, EGF, VEGF, PI3K, HGF, and PDGF)³¹, a viral entry factor³², and plays a role in RAS/MAPK signaling in various cancers^{33,34}. α -Cobratoxin (α -CTX), the lone non-receptor target, is a prominent toxin derived from elapid snakes³⁵ that acts by blocking muscle and nerve acetylcholine signaling. All except α -CTX are therapeutic targets for different cancers³⁶⁻⁴⁸. Designed protein binders against these targets could be useful for antagonizing native signaling, targeting drug conjugates and other therapeutics to tumors, designing novel agonists⁴⁹⁻⁵¹ and targeted receptor degradation⁵², inhibitors¹, triggering cargo endocytosis⁵², or target therapeutics towards particular cell types⁵³.




Figure 2.1.2. Design of β -strand pairing binders. (a) Representation of β -strand interface conditioning information provided as a tensor to RFdiffusion to generate β -strand pairing binders. (b) Example RFdiffusion denoising trajectories without (top) and with (bottom) β -strand interface conditioning. Conditioning indicates that part of the binder scaffold should be a β -strand (gold) that contacts the indicated target edge strand (cyan). This information influences the denoising in very early trajectory timesteps (t), with the tertiary fold determined within 5 timesteps and the final output at t=50. (c) Binder design success rates using β -strand interface conditioning or RFdiffusion default settings with hotspots indicating the target edge strand of interest. (d) Structural clustering of RFdiffusion output binder scaffolds using hotspot conditioning (left) or β -strand interface conditioning (right) t-distributed Stochastic Neighbor Embedding (t-SNE) transforms of all-by-all pairwise template modeling (TM) scores₆₀ among binder scaffolds across all targets are plotted, with close proximity of points representing structural similarity. Output fold secondary structures are classified by color as indicated in the legend. Bold bordered data points indicate In silico successes (red circles, pAE interaction < 10, pLDDT > 85, $\Delta\Delta G < -30$) and experimentally validated binders (black stars).

The set includes single β -sheet targets with highly exposed edge β -strands (α -CTX, FCRL5, NRP1) and somewhat occluded edge β -strands (ALK-2 and ALK-3), as well as immunoglobulin

(Ig) fold β -sandwich domains (FCRL5, KIT, and PDGFR α). These were found to be challenging for binder design using standard hotspot-only conditioned RFdiffusion-either in silico design calculations had failed to generate designs predicted with high confidence to bind to the target or experimental testing had failed to yield binders (for ALK-3 and KIT). Because no experimentally determined structure exists for FCRL5, we designed binders against the AlphaFold 2 model⁵⁴⁻⁵⁶. To provide a stringent test of design specificity, the set includes some structurally related targets: ALK-2 and ALK-3 are members of the TGF- β superfamily⁵⁷, and KIT and PDGFR α are type-III receptor-tyrosine kinases⁵⁸.

In a first in silico benchmark experiment, we compared the diversity and quality of binders designed using our interface conditioning method to designs generated using the standard RFdiffusion target site “hotspot” directed approach. Following the generation of 500 scaffolds using each diffusion approach, ten candidate binder sequences were generated for each scaffold by ProteinMPNN⁵⁹ and each binder candidate complex was predicted with AlphaFold 2⁵⁴. Across all eight target sites across the seven target proteins (two sites were targeted for PDGFR α due to predicted structural variability), β -strand interface conditioning yielded design models with improved in silico binding metrics⁶ AlphaFold 2 interface predicted aligned error (pAE) and Rosetta $\Delta\Delta G$ (Figure 2.1.2c), with over 9.2% of β -strand interface conditioned designs meeting reasonable quality metrics, in contrast to 0.98% success rates by RFdiffusion conditioned on target hotspots alone. β -strand interface conditioning yielded in silico design successes for all targets, whereas hotspot conditioning alone did not yield any in silico successes for two targets (ALK-3 and FCRL5). β -strand interface conditioning also yielded more globular binder designs with 88.7% of output scaffolds meeting radius of gyration criteria and only 25.5% of designs from the other methods meeting the same criteria (Figure 2.1.3). β -strand interface conditioning yielded outputs that covered a distinct protein structure space compared to other methods, as quantified by all-by-all pairwise template modeling (TM) scores⁶⁰

between aligned binder scaffolds (Figure 2.1.2c), with as expected a higher fraction of β -sheet containing binders and a decrease in α -helical bundle outputs (Figure 2.1.2d).




Figure 2.1.3. Success rates from different RFdiffusion methods across several predictive in silico metrics. 40,000 total designs were generated, with equal numbers generated from different RFdiffusion scaffold generation methods across eight target sites. The methods used included hotspot-only conditioning by the default RFdiffusion model (teal, the default for hotspot conditioned binder design), hotspot-only conditioning by a SS/ADJ aware model (dark green, no secondary structure interface conditioning provided), α -helix interface and hotspot conditioning (cyan), and β -strand interface and hotspot conditioning (green). AlphaFold 2 predictions yielded two in silico filtering scores: predicted error among interface residues (interface pAE) and overall confidence in the binder fold (pLDDT). Binders with interface pAE < 10 and binder pLDDT > 85 are considered successful designs with respect to these metrics. Rosetta $\Delta\Delta G$ represents the Rosetta predicted change in energy in the binder designs and targets upon binding. Lower $\Delta\Delta G$ is predictive of higher affinity interactions, with a $\Delta\Delta G$ value of -30 used as a standard of success here. Radius of gyration is a statistic of protein globularity that is predictive of a protein's fold stability. In nature, nearly all 65 amino acid proteins (the size of all designs in this in silico benchmarking set) have a radius of gyration < 14, mostly reflecting known globular folds, such as 3+ α -helix bundles, mixed α/β proteins, and β -sandwich folds. Binders meeting each of these four criteria were considered in silico successes with a high likelihood of experimental success. Additionally, the percentage of outputs exhibiting binder-target β -strand pairing are plotted, although this was not used as a filter for success. In cases where strand

pairing binders were generated without strand interface conditioning, binder-target strand pairs tended to contain fewer binder strands and strand pairing interactions.

A second much larger campaign was carried out for each target to generate designs for experimental testing. The sets were generated using β -strand targeted RFdiffusion and with standard hotspot conditioned RFdiffusion, followed by ProteinMPNN, and selection based on AlphaFold 2 and Rosetta metrics; partial diffusion scaffold optimization^{4,61} and additional sequence sampling was carried out to generate designs with improved metrics in several cases. β -strand pairing designs greatly outperformed other scaffold types in this in silico filtering stage for most targets with the exception of the PDGFR α target site, which was likely more amenable to helical bundle binder design success due to its flat, hydrophobic surface. Structurally diverse designs with high AlphaFold 2 confidence (interface pAE < 8 and pLDDT > 85) designs were selected for experimental screening by yeast surface display. While we ordered libraries including thousands of designs for some targets (KIT, PDGFR α , FCRL5, ALK-2, and ALK-3), binder success rates in our screen suggested that it would be possible to design strand pairing binders at smaller scales (Table 2.1). Indeed, design sets for NRP1 and α -CTX yielded binders at 96 well scales. The α -CTX binder originally introduced by Vázquez Torres, et al., referred to in that work as LNG, will be referenced as α -CTX binding protein (α -CTXbp) in this work. α -CTXbp had been designed with the interface-conditioning method introduced here, but limited details on the method were revealed previously⁶². For PDGFR α , where both helical bundle designs and β -strand pairing designs passed in silico filters, significantly more binders were obtained with the strand directed approach than with standard RFdiffusion at this selection stage (Figure 2.1.4).




Figure 2.1.4. Binding screens suggest preference for strand pairing binders at target sites. **(a)** β -strand pairing PDGFR α binders are more enriched for target binding in yeast surface display (YSD). Two design approaches, β -strand interface conditioned (N=5616) and hotspot only conditioned (N=10026), were utilized to target PDGFR α . This library containing designs from both approaches were expressed in vectors to display binders extracellularly and with a Myc-tag. Libraries were pooled, sorted for expression by binding to fluorescent Anti-Myc-FITC. The expression positive fraction was initially sorted for PDGFR α binding by incubating cells with biotinylated PDGFR α and Streptavidin phycoerythrin prior to sorting. This parent sort was repeated with titrated PDGFR α concentrations as labeled, and the fold enrichment of each binder in this sort was plotted relative to the NGS counts in the parent sort. **(b)** Affinity kinetics of YSD enriched binders against PDGFR α . The 96 designs that were most enriched during YSD were bacterially expressed and purified for binding measurements against biotinylated PDGFR α via SPR. Binders with K_D values $> 1 \mu\text{M}$ and exhibiting half the observed maximal SPR response are shown. **(c)** YSD binding data for the KIT binder design library. The binding fraction is represented as the fraction of overall sorting events being FITC (surface expressed binder) and Streptavidin-PE (associated with biotinylated receptor) positive. Lower binding signal in the presence of native ligand confirms the KIT binder library demonstrates the binders largely bind at the targeted Stem Cell Factor (SCF) ligand-binding site.

The most enriched designs for binding affinity and expression from this assay were cloned, expressed in *E. coli*, and purified (Figure 2.1.5).




Figure 2.1.5. Size-exclusion chromatography purification of β -strand pairing binders shows binders are largely monodisperse. Each binder was expressed and purified in high yields ($>1\text{mg}$) from 500 mL *E. coli* cultures. Each trace is within the expected range for monomeric species. There is minimal evidence of significant aggregation during purification at room temperature, with aggregate peak areas ($<10\text{ mL}$ retention volume) comprising less than 10% of the overall eluate for 7/8 designs. The mg/mL concentrations of collected fractions from purifications were quantified as follows: 1.0 (ALK-2bp), 3.7 (ALK-3bp), 24.352 (α -CTXbp), 1.3 (KITbp), 2.8 (FCRL5bp), 6.5 (NRP1bp), 2.0 (PDGFRabp-7LBF), and 0.7 (PDGFR α -model).

SPR revealed mid- to sub-nanomolar affinity binders for our targets (Figure 2.1.6). In the following sections we describe the active designs for each target in turn.




Figure 2.1.6. Binder design models and biophysical characterization. (a) Design models of binding complexes for each target. Left: Entire complex design models with binders indicated in green/gold and targets in teal/cyan. β -strand interface conditioned scaffold design yielded binders with interfacial β -strands (gold) forming strand-pairing hydrogen bonding interactions with target edge β -strands (cyan). Middle: Close up view of design model backbone hydrogen-bond interactions with putative hydrogen bonds shown in white. Right: schematic representation of strand pairing interactions showcase the diversity of sequence-independent β -strand pairing interactions. (b) Circular dichroism thermal melts. Full spectrum analyses (left) performed at 25 °C (cyan), 75 °C (green), and 95 °C (gold) assess the overall binder fold at these three temperatures, while single wavelength thermal melts (right) were measured at 217 nM to calculate binder T_m values. (c) SPR measurement (pink) of binding kinetics at 600pM, 4nM, 30nM, 200nM, 1.5 μ M, and 10 μ M (association phases 1-6 on the X-axis). Fits for K_d determination (green) excluded the 10 μ M data excluded due to signal aberrations at this high binder concentration. Binders were reproduced and similar SPR kinetics were fitted with $n \geq 2$ for each binding protein.

2.1.4 Designing specific binders to BMP receptors ALK-2 and ALK-3

ALK-2 and ALK-3 are two of four known Type I Bone Morphogenetic Protein Receptors (BMPRI). Targeting BMPR therapeutics towards specific tissues and pathologies could have considerable utility, but achieving binding specificity for ALK-2 versus ALK-3 presents a difficult challenge due to their structural homology. The ALK-2 and ALK-3 ectodomains have 30% sequence identity and high structural similarity. On both targets, there is an edge β -strand with five consecutive residues with backbone atoms in very similar atomic positioning (RMSD 0.07), and each strand is partially occluded by an N-terminal coil. Prior to this study, multiple design campaigns – including campaigns utilizing RFdiffusion, ProteinMPNN, and AlphaFold 2 – had failed to yield binders for ALK-3. The difficulty of designing binders to these targets is demonstrated in their lower success rates during our yeast surface display assay, with less than 1% of designs showing strong binding and enrichment (Table 2.1.1).

Library	Entirely Helical Designs	Strand containing Designs	Designs with Strand Interfaces	Parallel Strand Pairing Designs	Antiparallel Strand Pairing Designs	Strand Pairing Binders	Entirely Helical Binders	Top 50 Binders (Strand Pairing)	Top 50 Binders (Entirely Helical)	Top 50 Binders (Parallel Strand Pairing)	Top 50 Binders (Antiparallel Strand Pairing)
ALK-2	4207	5257	5039	18	5018	2	0	2	0	2	0
ALK-3	1735	4263	4196	241	3879	10	0	10	0	0	10
KIT	0	1298	1298	165	1065	1197	0	50	0	7	43
FCRL5	1	4840	4839	421	4236	118	0	50	0	2	48
PDGFR α	9922	5720	5497	2953	2433	73	29	36	14	31	5

Table 2.1.1. Strand-pairing binders show improved yeast surface display binding (YSD) success rates. Binder designs passing in silico filters against each target were ordered and screened for binding by YSD and sorting cells binding to fluorophore complexed target receptors. For ALK-2, ALK-3, and PDGFR α some designs that lacked target strand-pairing interactions passed in silico filters, but were less enriched in the same YSD binding experiments compared to designs with strand-pairing interfaces in the same library. For KIT and FCRL5, designing helical binders with quality in silico metrics was exceedingly difficult. Observed preferences for strand-pair orientation varied for each target, both parallel and antiparallel strand-pairing orientations were observed to show binding signal for three out of five of these targets. Binders are defined as designs with estimated affinities (SC_{50}) less than $1\mu M$. For ALK-2 and ALK-3, we purified and characterized all the binders identified in this screen. From the many KIT, FCRL5, and PDGFR α binders identified by YSD, smaller libraries were selected for individual purification and characterization ($N=27$, $N=12$, and $N=96$, respectively). These designs were selected by ranking designs by calculated YSD SC_{50} , and selecting the top N

designs. For PDGFR α , a larger library was considered for testing to better compare strand-pairing and entirely helical binders.

We conditioned RFdiffusion on generating β -strand pairings to these five edge β -strand residues that were not occluded by the N-terminal coil. We identified binding proteins, ALK-2bp and ALK-3bp, with K_D values of 96 nM and 528 pM, respectively (Figure 2.1.6c) by surface plasmon resonance (SPR). While both binders were intended to strand pair with the homologous target edge β -strand, the designed interaction footprints of their validated were quite distinct, with ALK-2bp forming a parallel β -strand interaction with the target β -strand and ALK-3bp pairing to the same strand in an antiparallel manner. ALK-3bp makes additional strand-pairing like contacts with the C-terminal coil of ALK-3 in a non-canonical hydrogen bonding pattern, with the binder β -strand curvature nearly perfectly complementing a bulge in the coil motif; these additional interactions may account for the four orders of magnitude greater binding affinity. This additional hydrogen bonding network with the C-terminal coil was not prespecified; RFdiffusion simply generated the secondary binder β -strand given the context of the initial β -strand motif specified by interface conditioning.

2.1.5 Targeting exposed edge β -strands in α -CTX, FCRL5, and NRP1

Three of the targets have more exposed edge β -strands. The acetylcholine receptor antagonist α -CTX is a three-finger toxin that consists of a single β -sheet with extended loops that bind a hydrophobic receptor pocket. We aimed to design a β -strand pairing binder that would sterically interfere with this interaction. FCRL5 weakly binds IgG to modulate B-cell activation. While the exact IgG binding site is unknown⁶³, the receptor consists of several Ig domains. In the AlphaFold 2 model, the N-terminal Ig domain contains one three stranded sheet and one five stranded sheet, leaving an exposed β -strand available to target during design. Finally, we targeted the discoidin domain of NRP1, which contains a highly twisted β -sheet with a potential β -strand binding site where the edge β -strand twists about 90 degrees with exposed backbone

polar atoms before entering into the protein core. While quite exposed, these target edge strands have distinctive structures that we hypothesized could allow for specific β -strand complementarity.

β -strand interface conditioned RFdiffusion resulted in a 5-fold improved in silico success rate for these targets compared to hotspot directed RFdiffusion which, even without conditioning, generated β -strand pairing binder designs with up to 20% of outputs (β -strand pairing designs were generated <5% of the time for the less exposed edge strand containing targets). Binders for α -CTX (originally introduced by Vazques-Torres, et al.) and NRP1 (1.9 nM and 101 nM respectively) were obtained by testing the 96 designs based solely on in silico metrics. The best binder to FCRL5, with a 82 nM K_D , was obtained from a yeast display library of 4841 designs. The FCLR5 and NRP1 target β -strands both had similar conformational twists, but FCLR5bp and NRP1bp utilize quite distinct antiparallel and parallel hydrogen bonding binding modes, respectively (Figure 2.1.6a, right). Similar to FCLR5bp, α -CTXbp forms a mostly canonical antiparallel strand pair with a slight irregularity as the binder β -strand adapts to complement a small target strand bulge.

2.1.6 Targeting type III receptor tyrosine kinases KIT and PDGFR α

The type III RTK family receptors KIT and PDGFR α play roles in angiogenesis, tissue regeneration, and aberrant cancer signaling. Both receptor ectodomains are comprised of five Ig-like folds, with native ligands – stem cell factor and platelet derived growth factor – activating cellular signaling pathways by binding Ig domain 2 of KIT and PDGFR α , respectively, to induce receptor dimerization and intracellular cross-phosphorylation. Binders designed to occupy the ligand-binding pocket could act as antagonists to prevent aberrant signaling, and when oligomerized⁵⁰, high-affinity binders could function as strong signaling agonists for tissue repair

therapeutics. As such, we targeted the ligand binding sites on domain 2 of each receptor. As in the case of ALK-3, previous attempts to de novo design binders against KIT had failed.

For KIT, a 65 nM binder (KITbp) was identified from a yeast surface library screening of 1298 designs. KITbp was designed to bind domain 3 strand on KIT which is part of the stem cell factor ligand binding site; the binding interface features an extensive canonical antiparallel hydrogen bond network, with 8 consecutive hydrogen bonds contributing a β -strand pair of 18 total binder and target residues (Figure 2.1.6a, left). For PDGFR α , two binders were identified – PDGFRabp-7LBF (K_D = 137 pM) was designed to bind the cryo-EM structure (PDB accession code: 7LBF), while PDGFRabp-AF2model (K_D , 4 nM) was designed to bind the AlphaFold 2 model. They were enriched from 5427 and 189 member design libraries via yeast display, without any further experimental sequence optimization. Despite significant disagreement between the cryo-EM structure and the AlphaFold 2 model regarding the conformation of the target domain 2 Ig fold, design against both conformations yielded high-affinity binders. The binders may selectively induce the cognate PDGFR α conformation upon binding, as their design model conformations are not cross-compatible (Figure 2.1.7). PDGFRabp-7LBF forms a complex h-bond network in which the binder strand forms both parallel and antiparallel interactions with two different target strands of the target domain 2 Ig fold. This precisely complementary strand pairing h-bond network highlights the power of β -pairing conditioned RFdiffusion to design complex β -strand architectures to perfectly complement idiosyncratic target topologies.




Figure 2.1.7. The two PDGFRα binder design models are predicted to induce different PDGFRα conformations. Our PDGFRαbp-model binder (lime green) designed against the PDGFRα AlphaFold 2 model (cyan) binds to target site residues (orange) in a conformation that significantly varies from the available cryo-EM structure (PDB ID 7LBF, teal). When the AlphaFold 2 model binder (dark green) is aligned to the cryo-EM model, the binder interface clashes with the target site residues (red inset). Similarly, aligning the PDGFRα bp-7LBF binder (dark green) designed against the cryo-EM structure to the AlphaFold 2 model results in a target site clash. Structure agnostic (i.e., only sequences are provided) AlphaFold prediction of each binder in complex with PDGFRα reflects the conformational variability, with the AlphaFold model conformation favored in PDGFRαbp-AF2model predictions while the cryo-EM conformation is favored when predicted with PDGFRα bp-7LBF.

2.1.7 Stability of β-strand pairing binders

No obvious trends were observed between the size of the edge β-strand interface and aggregation propensity of the binders. All binders could be purified at high yields, and analysis

of size exclusion chromatography curves reflect elution patterns consistent with monomeric binders being the most prominent purified species (Figure 2.1.5). Circular dichroism thermal melts (Figure 2.1.6b) obtained for each binder indicate that the binder folds remain intact at high temperatures, even those with significant edge β -strand content. ALK-3bp was stable and monomeric up to 95 °C during a circular dichroism thermal melt experiment despite having four exposed edge β -strands. The binder with the second most β -strand content, ALK-2bp, was also thermostable with a measured T_M of 81.2 °C. Binders with majority α -helical support for their β -strand interfaces (KITbp, α -CTXbp, and PDGFRabp-AF2model) did not seem to be intrinsically more thermostable with measured T_M s of 63.2 °C, 95 °C, and 95 °C respectively. The binders with α/β secondary structure content (2-3 β -strand sheets with 2 buttressing α -helices; FCRL5bp, NRP1bp, and PDGFRabp-7LBF) had T_M s of >95 °C, 77.2 °C, and 71.1 °C, respectively. The excellent solution behavior of the designs despite having edge β -strands clearly available for intermolecular interactions suggests that the same idiosyncratic features that enable high affinity and specificity (see below) target binding disfavor self-self interactions⁹.

2.1.8 Specificity of β -strand pairing binders

We next investigated whether the identified binders were specific for their designed target and did not form off target high affinity β -strand pairing interactions. To test this, we performed an all-by-all SPR experiment where each binder was tested for binding affinity against each of the receptor targets in the test set. At 1.5 μ M and 200nM concentrations, all of the binders showed strong SPR response for their intended target receptor compared to off-target receptors (Figure 2.1.8a). There was no evidence for off-target binding of the ALK-2, ALK-3, KIT, and PDGFR α binding proteins to related family members (e.g., ALK-2bp did not bind strongly to ALK-3).




Figure 2.1.8. Designed binders are highly specific for their targets. (a) Designed binder SPR response traces against all targets. For each binding protein, immobilized cognate target protein yielded the strongest SPR binding response at various binder concentrations (7-fold increasing binder concentrations ranging from 600pM to 1.5μM, from left to right in each trace). (b,c) Average response units for the (b) 1.5μM and (c) 200nM binder concentrations. Grid colors indicate response units relative to the maximum response for each binder.

This high specificity may arise because the binders for each structurally similar target pair (i.e., ALK-2 and ALK-3, KIT and PDGFRα, FCRL5 and NR1), were designed with different β-strand pairing hydrogen bonding arrays, and each binder makes additional non-strand-pairing contacts with the intended target. With the exception of ALK-3bp which forms exclusively β-strand contacts with ALK-3, all designed binders form additional α-helical contacts that are complementary to neighboring target surface topologies. Overall, the interface side chain contacts resemble natural protein interfaces, with binder side chains making electrostatic and shape complementarity interactions with target side chains. The polar interactions made by strand-pairing lead to lower overall hydrophobicity, as measured by Spatial Aggregation Propensity (Figure 2.1.9b).




Figure 2.1.9. Analysis of strand pairing binder interface polarity in context of previously published de novo binders. **(a)** Number of hydrogen bonds among 52 binder design models with validated binding data (e.g. biolayer interferometry or SPR). Examples include binders introduced by Cao., et al. along with several recently developed de novo binders made with deep learning methods base methods (i.e. ProteinMPNN, RFdiffusion, AlphaFold 2). By definition, strand pairing binders form hydrogen bonds with their targets, leading to interfaces with more hydrogen bonds than typical de novo binder interfaces. Additionally, KITbp forms 21 hydrogen bonds with KIT, three more than any previous de novo designed binder, to our knowledge. **(b)** Change in Spatial Aggregation Propensity (Δ SAP) of binders and targets upon complex dissociation. The Δ SAP statistics from binders identified in three previous works are also shown, with 13 binder complexes from Cao, et al., and 64 from Watson, et al. The targets for these works include CD3 δ , EGFR, FGFR, H3, IGF1R, IL7Ra, Influenza HA, Insulin Receptor, PDGFR β , PDL1, TGF- β , Tie-2, TrkA, and VirB8. Binders designed with β -strand interface conditioning (three of which were introduced by Vazques-Torres, et al.) are labeled with their target protein name. The binder footprints for SHRT (Vazques-Torres, et al.), KIT, and FCRL5 are all more hydrophilic (by SAP score) than any of the sites targeted in these two studies, while KITbp, FCRL5bp, and NRP1bp are all among the least hydrophobic designed binder interfaces from these studies.

2.1.9 KITbp-KIT crystal structure confirms design accuracy

To assess the accuracy of our design method, we solved the structure of KITbp in complex with domains 1-3 of KIT at 2.8 Å resolution (Figure 2.1.10a, Table 2.1.2). The crystal structure has near exact agreement with the computational design model with 1.9 Å all-atom RMSD between the design model and the structure. Over the designed binder alone, the RMSD over all atoms between the crystal and design is 2.0 Å, and over the backbone, 0.98 Å RMSD (Figure 2.1.10b). The design model and crystal structure align with atomic level accuracy over nearly all interfacial side chain residue atoms. The KITbp binding site on KIT overlaps with that of the KIT native ligand stem cell factor (SCF), as all the binder designs targeted this site. Consistent with this, saturating concentrations of SCF (60 μ M) reduced KIT library binding sort counts by 99% in a yeast surface display binding assay (Table 2.1.2).

Crystal data	
Space group	$P2_12_12_1$
Unit cell parameters a, b, c [Å]	41.273, 78.643, 291.008
Unit cell parameters α, β, γ [°]	90, 90, 90
Data collection	
Wavelength [Å]	0.9762
Resolution range [Å]	69.18-2.80 (2.90-2.80)
Total reflections	250 107 (37 292)
Unique reflections	24 388 (3 846)
Multiplicity	10.3 (9.7)
Completeness [%]	99.8 (99.4)
Mean I/σ [I]	5.24 (0.69)
Wilson B-factor	51.56
R_{meas} [%]	43.5 (263.9)
$CC_{1/2}$	99.0 (65.2)
Data statistics for anisotropy-corrected data	
Resolution range [Å]	69.18-2.80 (2.90-2.80)
Unique reflections	15 541 (626)
Completeness [%]	63.60 (26.31)
Wilson B-factor	34.48
Refinement statistics for anisotropy-corrected data	
Reflections used in refinement	15 541 (626)

Reflections used for R-free	735 (42)
R_{work} [%]	22.54 (30.42)
R_{free} [%]	28.24 (36.73)
<i>Number of non-hydrogen atoms</i>	
Total	5 639
Macromolecules	5 472
Ligands	136
Solvent	31
Protein residues	682
RMS (bonds) [Å]	0.001
RMS (angles) [°]	0.40
Ramachandran favored [%]	95.40
Ramachandran allowed [%]	4.60
Ramachandran outliers [%]	0.00
Rotamer outliers [%]	0.16
Clashscore	4.12
<i>B-factors [Å²]</i>	
Average	48.76
Macromolecules	48.92
Ligands	47.47
Solvent	27.28
Number of TLS groups	11

Table 2.1.2. *Crystallographic data collection and refinement statistics for the determined structure of the KITbp-KIT123 complex (PDB ID: 9H71). Statistics for the highest-resolution shell are shown in parentheses.*

The 29 residue binder solvent exposed surface area (SASA) of the binder interface (1164 Å²) consists of 8 strand-pairing binder residues (from Gly21 to Asp28, 291 Å²) and 21 additional interface residues (873 Å²). By heavy atom SASA, the polar atoms comprise 65% of the strand-pairing interface, a somewhat higher proportion compared to the rest of the interface surface which was only 45% polar. There are 21 sidechain-sidechain and backbone-backbone hydrogen bonds between KITbp and KIT (calculated by HBPLUS⁶⁴), many more than in previous designed binders (Figure 2.1.9a) and in most native protein-protein interfaces with similar interface sizes⁶⁵ (Figure 2.1.10). The eight strand pairing binder residues participate in nearly half (9) out of the total (21) hydrogen-bonds in the 29 residue binder interface. The

majority of these hydrogen bonds were predicted accurately in our design model (16 out of 21). At the center of the interface, a network of four binder and KIT tyrosine residues (Y38, Y39, Y259, Y269) form a highly complementary interface between the binder interfacial helix and core-boundary of the KIT domain 2 lg fold. Lining the perimeter of the interface are numerous polar interactions, including the designed β -pairing interface as well as several polar side-chain interactions.




Figure 2.1.10. Structural analysis of the KITbp-KIT complex. (a) Crystal structure (PDB ID: 9H71) of KITbp (cyan) in complex with KIT (green) and 2Fo-Fc electron density contoured at +1.0 RMSD (gold). (b) The KITbp crystal structure superimposes on the binder design model (grey) with subangstrom backbone atom RMSD, and 2.0 Å all-atom RMSD. (c) Close-up views of the binder interface reveal high shape complementarity of the binder-target complex (top). The binder interface consists of hydrophobic and polar interactions between binder core-boundary residues and KIT receptor domain 2 core-boundary residues (middle). The binder forms an extensive hydrogen-bond β -pairing interaction with the targeted domain 2 edge strand (bottom). There are multiple side-chain interactions between binder and KIT tyrosine residues (Y38 and Y39 in the binder; Y259 and Y269 in KIT) in the core-boundary interaction, and between binder glutamate residues (E14, E28, E24, E35) and KIT hydrogen bond donors (Q256, K258, N260, W262, H263, R271). The binder also makes electrostatic interactions with both the opposite domain 2 edge strand and KIT domain 3.

KITbp contains an unpaired β -strand that pairs with the target, with three buttressing α -helices that tether the β -strand in place. Foldseek66 alignment of KITbp against the PDB did not identify significant matches to known protein structures. We were unable to find examples of lone β -strands in natural proteins, except those found in protein-protein complexes where a loop in one partner forms a β -strand conformation that extends the β -sheet of the other partner⁶⁷. A similar induction of β strand formation upon binding may occur with KITbp; Alphafold 2 prediction of the monomeric KITbp without KIT target places the β -strand pairing residues in a coil conformation that better shields the binder's hydrophobic core residues and hydrogen bonding atoms of the β -strand interface (Figure 2.1.11); such KIT-dependent conformational switching may contribute to the observed binding specificity.




Figure 2.1.11. Surface hydrophobicity representations of the KITbp interface for the monomeric (top) KIT-bound (middle) AlphaFold2 predictions of KITbp. The crystal structure representation (bottom) and the AlphaFold2 multimer prediction display more of KITbp hydrophobic residues than the AlphaFold 2 monomer prediction. The hydrogen bonding atoms (gold) are observed in a β -strand conformation amenable to target β -strand pairing in the KIT-bound conformations, while these same atoms are occluded in the AlphaFold 2 monomer prediction.

2.1.10 Functional testing of FCRL5 and PDGFR α binders

We next sought to assess the biological functionality of the designs. FCRL5 is internalized and transits through the endocytic degradation pathway upon binding to antibodies⁶³, and we investigated whether FCRL5bp could be similarly internalized. We incubated cells expressing FCRL5 with FCRL5bp tagged with pHrodo DeepRed, a pH sensitive (~ 5 pKa) dye that emits

655 nm fluorescence at late endocytic vesicle pH, and observed binder internalization that correlated with FCRL5 expression levels, reaching a steady state in this expression system (Figure 2.1.12a). We hypothesized that FCRL5bp could thus be useful as an Endotag⁵² for targeted protein degradation, as recently demonstrated for other designed proteins. We fused FCRL5bp to EGFRbp (EGFRn from Cao, et al.) in both terminal orientations and compared the ability of this molecule to degrade EGFR by endocytic lysosomal trafficking induced by FCRL5 binding. After treating cells with 50nM of either EGF, the binder fusions, or each binder alone, EGFR was robustly degraded in a manner dependent on doxycycline induced expression of FCRL5.





Figure 2.1.12. Functional activity of FCRL5 and PDGFRα binding proteins. (a) HeLa cells engineered to express FCRL5 (HeLa-FCRL5) receptor in a Dox-inducible manner were treated with 50nM neutravidin-labeled pHrodo DeepRed complexed FCRL5bp, and live cell imaging (top panels) was used to measure overlap of GFP (FCRL5-expressing cells) and pHrodo Red fluorescence (internalized FCRL5bp). 1000 ng/mL doxycycline results in strong GFP expression and intracellular pH activated pHrodo Red fluorescence. Internalization (measured arbitrary units calculated by overlap of red and green fluorescence divided by phase area) reached steady state in 18 hours (images are from this timepoint); values are presented as mean $\pm$ SD for $n = 3$ replicate wells. (b-c) EGFR is degraded HeLa-FCRL5 cells by de novo binders. Measuring EGFR degradation by immunostaining and Western blot (b) or flow cytometry (c); $n = 3$ independent biological replicates for western blots and flow cytometry. FCRL5bp fusion to EGFRbp resulted in comparable degradation to the native degradation mechanism (EGF signaling). Gating strategies for flow cytometry are provided in Supplementary Figure 9. (d-e) Western blot analysis of PDGFRα inhibition in Chinese Hamster Ovary cells engineered to

overexpress PDGFR α (repeated four times with similar results). Levels of phosphorylated PDGFR α , Erk, and Akt were measured by immunoblots with fluorescent antibodies. Signals were normalized by the fluorescent signal of an antibody against the constitutively expressed housekeeping proteins S6 or actin. Data are presented as mean $\pm$ SD of $n = 3$ independent biological replicates.

PDGFR α -7LBF was designed to sterically occlude the binding site of the native ligand PDGF-AA, and indeed PDGFR α -7LBF blocked activation of signaling by PDGF-AA through the Akt and Erk pathways with nanomolar inhibition of ERK phosphorylation ($IC_{50} = 19.6$ nM) and AKT phosphorylation ($IC_{50} = 38.25$ nM) in a CHO cell line engineered to stably express PDGFR α (Figure 2.1.12d-e).



Supplementary Figure 2.1.13. Gating strategy used for quantifying surface EGFR levels. Gating strategy used for analyzing live HeLa cells with stably integrated Dox-Inducible Flag-FCRL5-T2A-P2A-EGFP for FCRL5+ (EGFP) expression and subsequent analysis of EGFR surface expression.

2.1.11 Engineered ALK-2/ALK-3 binder agonists activate BMP signaling

We also tested ALK-2bp and ALK-3bp for their ability to signal via BMP receptors. These BMP receptors form heterotetrameric signaling complexes and signal via SMAD1/5/9¹⁶. As such, we designed chimeric agonists by genetically fusing ALK-2bp and ALK-3bp to human Fc and one of two de novo binding proteins for Activin Receptor IIA (ActRIIAbp), such that two copies of the ALK-2/ALK-3 and ActRIIA binders are displayed on each dimerized Fc. Titrating these agonists on an EA.hy926 BMP-responsive element (BRE)-driven mCherry reporter cell line resulted in dose-dependent activation of SMAD reporter activity with nanomolar calculated EC_{50} values (Figure 2.1.14). The level of response appeared to depend on the linker length used between the two binders, with Fc-1×GGG-ALK-3bp–ActRIIAbp2 signaling with threefold greater potency than the same components linked by 4×GGG. The native ligand BMP9 elicited responses with greater E_{max} (27,975 response units) and lower EC_{50} (0.05 nM) in our assay, suggesting that additional factors beyond ligand binding and receptor colocalization contribute to BMP signaling strength.




Figure 2.1.14. Agonism of BMP signaling pathways by ALK-2bp and ALK-3bp fusion to ActRIIAbp. Chimeric agonists generated by genetically fusing ALK-2bp and ALK-3bp to human Fc and ActRIIAbp robustly signal on EA.hy926 BRE-mCherry reporter cells. The calculated EC_{50} values were 31.4 nM (Fc 4xGGG ALK-3bp ActRIIAbp1), 175.5 nM (Fc 4xGGG ALK-2bp ActRIIAbp2), 332.2 nM (Fc 1xGGG ALK-3bp ActRIIAbp2), and 929.9 nM (Fc 4xGGG ALK-3bp ActRIIAbp2).

2.1.12 Conclusions from the β -strand pairing binder study

We show through in silico modeling and high throughput binding assays that β -strand interface conditioned RFdiffusion outperforms α -helical focused methods for the design of binders against seven edge β -strand containing targets. The binders are thermostable and target-specific, and the crystal structure of the KITbp:KIT complex shows the method has high structural accuracy. The higher success rate of binder design against polar targets containing β -sheets using the β -pairing targeted approach than with primarily α -helical designs likely reflects the difficulty for the latter of complementing the many polar NH and CO groups on edge β -strands with designed sidechain-backbone hydrogen bonds-using geometrically matched β -strands this polar interface can be achieved in less convoluted fashion.

As noted by the Richardsons, edge β strands of native proteins often have irregular geometries that they speculated functioned to suppress off-target β -strand-mediated interactions with themselves and with other proteins. The very high binding specificity and the lack of significant aggregation of our designs despite the availability of edge β -strands for intermolecular interactions supports the conjecture that non-canonical β strand geometries suppress off target pairing, and considerably extends the pre-deep learning observations of Sahtoe et al in demonstrating that the irregular edge β -strands of native proteins can be systematically targeted provided that the geometry of the binder strand is sculpted to match the target.

Our β -strand targeted RFdiffusion approach should enable facile binder design against many previously challenging protein surfaces, including the many Ig domains and other β -sheet containing domains frequent in the extracellular domains of cell surface receptors. As many of these are current therapeutic targets, the method could contribute to new therapeutic interventions. While we acknowledge that the binders presented here may not be immediately feasible as therapeutic molecules, we expect that further optimization with these binders as

starting points may improve their affinity, stability, pharmacokinetics, and other desired characteristics (e.g., mouse/human cross-reactivity, oligomerization, or protease stability) for therapeutic applications.

References for Section 2.1

1. Cao, L. et al. Design of protein-binding proteins from the target structure alone. *Nature* 605, 551–560 (2022).
2. Watson, J. L. et al. De novo design of protein structure and function with RFdiffusion. *Nature* 620, 1089–1100 (2023).
3. Gainza, P. et al. De novo design of protein interactions with learned surface fingerprints. *Nature* 617, 176–184 (2023).
4. Vázquez Torres, S. et al. De novo design of high-affinity binders of bioactive helical peptides. *Nature* 626, 435–442 (2024).
5. Cao, L. et al. De novo design of picomolar SARS-CoV-2 miniprotein inhibitors. *Science* 370, 426–431 (2020).
6. Bennett, N. R. et al. Improving de novo protein binder design with deep learning. *Nat. Commun.* 14, 2625 (2023).
7. Stranges, P. B. & Kuhlman, B. A comparison of successful and failed protein interface designs highlights the challenges of designing buried hydrogen bonds: Successful Interface Designs Avoid Polar Interactions. *Protein Sci.* 22, 74–82 (2013).
8. Watkins, A. M. & Arora, P. S. Anatomy of β -strands at protein-protein interfaces. *ACS Chem. Biol.* 9, 1747–1754 (2014).
9. Richardson, J. S. & Richardson, D. C. Natural beta-sheet proteins use negative design to avoid edge-to-edge aggregation. *Proc. Natl. Acad. Sci. U. S. A.* 99, 2754–2759 (2002).
10. Sahtoe, D. D. et al. Transferrin receptor targeting by de novo sheet extension. *Proc. Natl. Acad. Sci. U. S. A.* 118, e2021569118 (2021).

11. Fujiwara, K., Ebisawa, S., Watanabe, Y., Toda, H. & Ikeguchi, M. Local sequence of protein β -strands influences twist and bend angles: Side Chain Influences β -Strand Twisting. *Proteins* 82, 1484–1493 (2014).
12. Fujiwara, K., Ebisawa, S., Watanabe, Y., Fujiwara, H. & Ikeguchi, M. The origin of β -strand bending in globular proteins. *BMC Struct. Biol.* 15, 21 (2015).
13. Richardson, J. S., Getzoff, E. D. & Richardson, D. C. The beta bulge: a common small unit of nonrepetitive protein structure. *Proc. Natl. Acad. Sci. U. S. A.* 75, 2574–2578 (1978).
14. Fox, N. K., Brenner, S. E. & Chandonia, J.-M. SCOPe: Structural Classification of Proteins--extended, integrating SCOP and ASTRAL data and classification of new structures. *Nucleic Acids Res.* 42, D304–9 (2014).
15. Chandonia, J.-M. et al. SCOPe: improvements to the structural classification of proteins - extended database to facilitate variant interpretation and machine learning. *Nucleic Acids Res.* 50, D553–D559 (2022).
16. Bragdon, B. et al. Bone morphogenetic proteins: a critical review. *Cell. Signal.* 23, 609–620 (2011).
17. Sorkin, M. et al. Hair follicle specific ACVR1/ALK2 critically affects skin morphogenesis and attenuates wound healing. *Wound Repair Regen.* 25, 521–525 (2017).
18. Liu, M., Goldman, G., MacDougall, M. & Chen, S. BMP signaling pathway in dentin development and diseases. *Cells* 11, 2216 (2022).
19. Sugimoto, H. et al. Activin-like kinase 3 is important for kidney regeneration and reversal of fibrosis. *Nat. Med.* 18, 396–404 (2012).
20. Ruan, X. et al. Activin receptor-like kinase 3: a critical modulator of development and function of mineralized tissues. *Front. Cell Dev. Biol.* 11, 1209817 (2023).
21. Marino, F. et al. Role of c-kit in myocardial regeneration and aging. *Front. Endocrinol. (Lausanne)* 10, 371 (2019).

22. Di Siena, S. et al. Activated c-Kit receptor in the heart promotes cardiac repair and regeneration after injury. *Cell Death Dis.* 7, e2317 (2016).
23. Cimini, M. et al. C-kit dysfunction impairs myocardial healing after infarction. *Circulation* 116, 177–82 (2007).
24. Yaniz-Galende, E. et al. Stem cell factor gene transfer promotes cardiac repair after myocardial infarction via in situ recruitment and expansion of c-kit⁺ cells. *Circ. Res.* 111, 1434–1445 (2012).
25. Lutz, M. et al. Local injection of stem cell factor (SCF) improves myocardial homing of systemically delivered c-kit⁺ bone marrow-derived stem cells. *Cardiovasc. Res.* 77, 143–150 (2008).
26. Takematsu, E. et al. Transmembrane stem cell factor protein therapeutics enhance revascularization in ischemia without mast cell activation. *Nat. Commun.* 13, 2497 (2022).
27. Kim, B.-J. et al. Platelet-derived growth factor receptor- α positive cardiac progenitor cells derived from multipotent germline stem cells are capable of cardiomyogenesis in vitro and in vivo. *Oncotarget* 8, 29643–29656 (2017).
28. Kalra, K., Eberhard, J., Farbehi, N., Chong, J. J. & Xaymardan, M. Role of PDGF-A/B ligands in cardiac repair after myocardial infarction. *Front. Cell Dev. Biol.* 9, 669188 (2021).
29. Horikawa, S. et al. PDGFR α plays a crucial role in connective tissue remodeling. *Sci. Rep.* 5, 17948 (2015).
30. Ono, C. et al. Upregulated Fcrl5 disrupts B cell anergy and causes autoimmune disease. *Front. Immunol.* 14, 1276014 (2023).
31. Pandey, P. et al. New insights about the PDGF/PDGFR signaling pathway as a promising target to develop cancer therapeutic strategies. *Biomed. Pharmacother.* 161, 114491 (2023).
32. Saiz, M. L. et al. Epigenetic targeting of the ACE2 and NRP1 viral receptors limits SARS-CoV-2 infectivity. *Clin. Epigenetics* 13, 187 (2021).

33. Rizzolio, S. et al. Neuropilin-1 upregulation elicits adaptive resistance to oncogene-targeted therapies. *J. Clin. Invest.* 128, 3976–3990 (2018).
34. Tang, Y. H. et al. Neuropilin-1 is over-expressed in claudin-low breast cancer and promotes tumor progression through acquisition of stem cell characteristics and RAS/MAPK pathway activation. *Breast Cancer Res.* 24, 8 (2022).
35. de la Rosa, G., Corrales-García, L. L., Rodriguez-Ruiz, X., López-Vera, E. & Corzo, G. Short-chain consensus alpha-neurotoxin: a synthetic 60-mer peptide with generic traits and enhanced immunogenic properties. *Amino Acids* 50, 885–895 (2018).
36. Ensan, D. et al. Targeting ALK2: An open science approach to developing therapeutics for the treatment of diffuse intrinsic pontine glioma. *J. Med. Chem.* 63, 4978–4996 (2020).
37. Carvalho, D. et al. ALK2 inhibitors display beneficial effects in preclinical models of ACVR1 mutant diffuse intrinsic pontine glioma. *Commun. Biol.* 2, 156 (2019).
38. Zhou, X. P. et al. Germline mutations in BMPR1A/ALK3 cause a subset of cases of juvenile polyposis syndrome and of Cowden and Bannayan-Riley-Ruvalcaba syndromes. *Am. J. Hum. Genet.* 69, 704–711 (2001).
39. Ehata, S. & Miyazono, K. Bone morphogenetic protein signaling in cancer; Some topics in the recent 10 years. *Front. Cell Dev. Biol.* 10, 883523 (2022).
40. Corless, C. L. et al. PDGFRA mutations in gastrointestinal stromal tumors: frequency, spectrum and in vitro sensitivity to imatinib. *J. Clin. Oncol.* 23, 5357–5364 (2005).
41. Simpson, J. E. et al. Autophagy supports PDGFRA-dependent brain tumor development by enhancing oncogenic signaling. *Dev. Cell* 59, 228–243.e7 (2024).
42. Jansson, S. et al. The PDGF pathway in breast cancer is linked to tumour aggressiveness, triple-negative subtype and early recurrence. *Breast Cancer Res. Treat.* 169, 231–241 (2018).
43. Evans, E. K. et al. A precision therapy against cancers driven by KIT/PDGFR mutations. *Sci. Transl. Med.* 9, (2017).

44. Curtin, J. A., Busam, K., Pinkel, D. & Bastian, B. C. Somatic activation of KIT in distinct subtypes of melanoma. *J. Clin. Oncol.* 24, 4340–4346 (2006).
45. Küçükköse, E. et al. KIT promotes tumor stroma formation and counteracts tumor-suppressive TGF β signaling in colorectal cancer. *Cell Death Dis.* 13, 617 (2022).
46. Rubin, B. P. et al. KIT activation is a ubiquitous feature of gastrointestinal stromal tumors. *Cancer Res.* 61, 8118–8121 (2001).
47. Sakuma, Y., Sakurai, S., Oguni, S., Hironaka, M. & Saito, K. Alterations of the c-kit gene in testicular germ cell tumors. *Cancer Sci.* 94, 486–491 (2003).
48. Hirota, S. et al. Gain-of-function mutations of c-kit in human gastrointestinal stromal tumors. *Science* 279, 577–580 (1998).
49. Bryan, C. M. et al. Computational design of a synthetic PD-1 agonist. *Proc. Natl. Acad. Sci. U. S. A.* 118, e2102164118 (2021).
50. Edman, N. I. et al. Modulation of FGF pathway signaling and vascular differentiation using designed oligomeric assemblies. *Cell* 187, 3726–3740.e43 (2024).
51. Silva, D.-A. et al. De novo design of potent and selective mimics of IL-2 and IL-15. *Nature* 565, 186–191 (2019).
52. Huang, B. et al. Designed endocytosis-inducing proteins degrade targets and amplify signals. *Nature* 638, 796–804 (2025).
53. Chevalier, A. et al. Massively parallel de novo protein design for targeted therapeutics. *Nature* 550, 74–79 (2017).
54. Jumper, J. et al. Highly accurate protein structure prediction with AlphaFold. *Nature* 596, 583–589 (2021).
55. Varadi, M. et al. AlphaFold Protein Structure Database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Res.* 50, D439–D444 (2022).

56. Varadi, M. et al. AlphaFold Protein Structure Database in 2024: providing structure coverage for over 214 million protein sequences. *Nucleic Acids Res.* 52, D368–D375 (2024).
57. Hinck, A. P. Structural studies of the TGF- β s and their receptors - insights into evolution of the TGF- β superfamily. *FEBS Lett.* 586, 1860–1870 (2012).
58. Lemmon, M. A. & Schlessinger, J. Cell signaling by receptor tyrosine kinases. *Cell* 141, 1117–1134 (2010).
59. Dauparas, J. et al. Robust deep learning-based protein sequence design using ProteinMPNN. *Science* 378, 49–56 (2022).
60. Zhang, Y. & Skolnick, J. TM-align: a protein structure alignment algorithm based on the TM-score. *Nucleic Acids Res.* 33, 2302–2309 (2005).
61. Gloegl, M. et al. Target-conditioned diffusion generates potent TNFR superfamily antagonists and agonists. *bioRxiv* (2024) doi:10.1101/2024.09.13.612773.
62. Vázquez Torres, S. et al. De novo designed proteins neutralize lethal snake venom toxins. *Nature* 639, 225–231 (2025).
63. Franco, A. et al. Human Fc receptor-like 5 binds intact IgG via mechanisms distinct from those of Fc receptors. *J. Immunol.* 190, 5739–5746 (2013).
64. McDonald, I. K. & Thornton, J. M. Satisfying hydrogen bonding potential in proteins. *J. Mol. Biol.* 238, 777–793 (1994).
65. Lo Conte, L., Chothia, C. & Janin, J. The atomic structure of protein-protein recognition sites. *J. Mol. Biol.* 285, 2177–2198 (1999).
66. van Kempen, M. et al. Fast and accurate protein structure search with Foldseek. *Nat. Biotechnol.* 42, 243–246 (2024).
67. Cheng, P.-N., Pham, J. D. & Nowick, J. S. The supramolecular chemistry of β -sheets. *J. Am. Chem. Soc.* 135, 5477–5492 (2013).
68. Yuzawa, S. et al. Structural basis for activation of the receptor tyrosine kinase KIT by stem cell factor. *Cell* 130, 323–334 (2007).

69. Appleton, B. A. et al. Structural studies of neuropilin/antibody complexes provide insights into semaphorin and VEGF binding. *EMBO J.* 26, 4902–4912 (2007).
70. Kschonsak, M. et al. Structures of HCMV Trimer reveal the basis for receptor recognition and cell entry. *Cell* 184, 1232–1244.e16 (2021).
71. Chennamsetty, N., Voynov, V., Kayser, V., Helk, B. & Trout, B. L. Design of therapeutic proteins with enhanced stability. *Proc. Natl. Acad. Sci. U. S. A.* 106, 11937–11942 (2009).
72. Muratspahić, E. et al. De novo design of miniprotein agonists and antagonists targeting G protein-coupled receptors. *bioRxiv* 2025.03.23.644666 (2025)
doi:10.1101/2025.03.23.644666.
73. Benatuil, L., Perez, J. M., Belk, J. & Hsieh, C.-M. An improved yeast transformation method for the generation of very large human antibody libraries. *Protein Eng. Des. Sel.* 23, 155–159 (2010).
74. Huang, T. & Hinck, A. P. Production, isolation, and structural analysis of ligands and receptors of the TGF- β superfamily. *Methods Mol. Biol.* 1344, 63–92 (2016).
75. Wicky, B. I. M. et al. Hallucinating symmetric protein assemblies. *Science* 378, 56–61 (2022).
76. Kabsch, W. XDS. *Acta Crystallogr. D Biol. Crystallogr.* 66, 125–132 (2010).
77. Jamshidiha, M. et al. Coping with strong translational noncrystallographic symmetry and extreme anisotropy in molecular replacement with Phaser: human Rab27a. *Acta Crystallogr. D Struct. Biol.* 75, 342–353 (2019).
78. Tickle, I. STARANISO: use of a WebGL-based 3D interactive graphical display to represent and visualise data quality metrics for anisotropic macromolecular diffraction data. *Acta Crystallogr. A Found. Adv.* 75, e162–e162 (2019).
79. Vonrhein, C. et al. Advances in automated data analysis and processing within autoPROC, combined with improved characterisation, mitigation and visualisation of the

anisotropy of diffraction limits using STARANISO. *Acta Crystallogr. A Found. Adv.* 74, a360–a360 (2018).

80. McCoy, A. J. et al. Phaser crystallographic software. *J. Appl. Crystallogr.* 40, 658–674 (2007).

81. Emsley, P., Lohkamp, B., Scott, W. G. & Cowtan, K. Features and development of coot. *Acta Crystallogr. D Biol. Crystallogr.* 66, 486–501 (2010).

82. Adams, P. D. et al. PHENIX: a comprehensive Python-based system for macromolecular structure solution. *Acta Crystallogr. D Biol. Crystallogr.* 66, 213–221 (2010).

83. Joosten, R. P., Long, F., Murshudov, G. N. & Perrakis, A. The PDB_REDO server for macromolecular structure model optimization. *IUCrJ* 1, 213–220 (2014).

84. Ng, P. K.-S. et al. Systematic functional annotation of somatic mutations in cancer. *Cancer Cell* 33, 450–462.e10 (2018).

85. Mátés, L. et al. Molecular evolution of a novel hyperactive Sleeping Beauty transposase enables robust stable gene transfer in vertebrates. *Nat. Genet.* 41, 753–761 (2009).

2.2 α -CTXbp and designed three-finger toxin binders for neutralizing snake venoms

This subsection summarizes and paraphrases results from Vázquez Torres et al., with emphasis on the α -cobratoxin binder LNG, referred to previously in this document as α -CTXbp1. Figure 2.2.1 reproduces and rearranges selected panels from Figures 1, 3, and 4 of that work.

We assessed the ability of the designs to functionally neutralize α -neurotoxins in patch-clamp experiments using a human-derived rhabdomyosarcoma cell line expressing muscle-type nAChRs. The LNG design had a better neutralizing efficacy than a previously characterized α -cobratoxin nanobody (TPL1158_01_C09)², achieving full protection at a 1:1 molar ratio (toxin:binder) (Figure 2.2.1e).

We used a cytotoxicity assay to evaluate the cross-reactivity of the CYTX design against various cobra whole venoms. Immortalized human keratinocytes (N/TERTs) were exposed to venoms from seven different *Naja* species, which previous proteomic analyses suggested primarily consist (approximately 70%) of cytotoxins³. Preincubating CYTX with venom (two IC_{50} values) at a 1:5 molar ratio (toxin:binder) provided 70-90% protection against venom-induced cytotoxicity (Figure 2.2.1d). Similarly, preincubation of the cytotoxin binder with the isolated cytotoxin from *N. pallida* (two IC_{50} values) at a 1:5 molar ratio (toxin:binder) provided 85% protection against cytotoxicity (Figure 2.2.1d). However, preliminary studies indicated that the CYTX design, in 1:1, 1:2.5 and 1:5 molar ratios (toxin:binder), did not significantly decrease the size of the dermonecrotic lesions induced by intradermal *N. nigricollis* venom administration in a murine model⁴. The affinity of CYTX likely needs to be further optimized for full in vivo neutralization of cytotoxins.




Figure 2.2.1. De novo designed protein binders neutralize three-finger snake venom toxins.

(a) Structure of a three-finger toxin (3FTx; PDB 1QKD), with highly conserved cysteine residues shown as sticks and each of the three fingers indicated (I-III). **(b)** Crystal structures of 3FTx-binding proteins closely match their computational design models and bind via RFdiffusion conditioned β -strand interactions. **(c)** Concentration-response curves comparing SHRT and an anti-ScNtx VHH for preventing nicotinic acetylcholine receptor (nAChR) blockade by ScNtx at one IC₈₀. The y-axis shows acetylcholine-evoked peak current in the presence of toxin and binder, normalized to the full acetylcholine response; values are averaged across each group ($n=16$). **(d)** Neutralization of the cytolytic effects of whole venoms from seven *Naja* species and isolated cytotoxin by CYTX. Two IC₅₀ values of whole venom or toxin were preincubated with CYTX at a 1:5 toxin:binder molar ratio. This ratio was estimated assuming that 70% of whole venom consists of cytotoxins. Keratinocyte medium was used as a positive control (PC), Triton X-100 as a negative control (NC), and CYTX binder alone as a binder-only control (B). “-” denotes venom or toxin without binder and “+” denotes venom or toxin incubated with binder. Experiments were performed in triplicate; results are mean $\pm$ s.d. **(e)** Concentration-response curves comparing LNG/ α -CTXbp and an anti- α -cobratoxin VHH for preventing nAChR blockade by α -cobratoxin at one IC₈₀. **(f)** Survival of mice following lethal neurotoxin challenge ($n=5$). ScNtx or α -cobratoxin was administered at three LD₅₀ values. In preincubation experiments, toxin and the corresponding binder were mixed at a 1:10 toxin:binder molar ratio for 30 min before intraperitoneal administration. In rescue experiments, binders were administered intraperitoneally 15 or 30 min after toxin challenge at 1:10 or 1:5 toxin:binder molar ratios. Controls included toxin alone (C) and cross-treatment (CT) with the non-target binder. Signs of toxicity were monitored and deaths were recorded for 24 h. Panels reproduced and rearranged from Vázquez Torres et al.1; original BioRender elements are retained from the source figure.

Given the encouraging in vitro neutralization of our anti-neurotoxin designs, we proceeded to in vivo studies. We determined the mean lethal dose (LD₅₀) values for α -neurotoxins in male non-Swiss albino mice via intraperitoneal administration; α -cobratoxin had an LD₅₀ of 0.098 $\mu\text{g g}^{-1}$ and ScNtx had an LD₅₀ of 0.087 $\mu\text{g g}^{-1}$, in agreement with previous intravenous LD₅₀ doses of these toxins (0.1 $\mu\text{g g}^{-1}$)⁵. To evaluate the in vivo neutralization efficacy of our neurotoxin-targeting designs, we monitored survival for 24 h post-lethal neurotoxin challenge following the administration of purified toxins (three LD₅₀ values) (Figure 2.2.1f). The LNG binder exhibited comparable efficacy, completely neutralizing α -cobratoxin but not the non-target ScNtx (Figure 2.2.1f, left). In rescue assays that better mimicked a real-life snakebite scenario, complete protection (100%) was achieved when long-chain α -neurotoxin binders were administered intraperitoneally at a 1:10 molar ratio (toxin:binder) 15 min after a lethal α -neurotoxin challenge (three LD₅₀ values) (Figure 2.2.1f, middle). Administering the binder 30 min after toxin injection also conferred 60% protection against α -cobratoxin (Figure 2.2.1f,

right). All of the surviving mice showed no evidence of limb or respiratory paralysis. At a 1:5 molar ratio (toxin:binder), intraperitoneal administration of the SHRT design 15 min after toxin injection (three LD₅₀ values) resulted in 80% survival, whereas the LNG binder provided 80% protection. Mice injected with the binder alone showed no negative effects at 24 and 48 h post-injection or up to 2 weeks post-injection.

References for section 2.2

1. Vázquez Torres S, Benard Valle M, Mackessy SP, Menzies SK, Casewell NR, Ahmadi S, et al. De novo designed proteins neutralize lethal snake venom toxins. *Nature*. 2025;639:225–231. doi:10.1038/s41586-024-08393-x.
2. Jensen A. Design of Consensus Toxins and Their Use for the Discovery of Broadly Neutralizing Antibodies. PhD thesis, DTU Bioengineering; 2023.
3. Petras D, et al. Snake venomomics of African spitting cobras: toxin composition and assessment of congeneric cross-reactivity of the Pan-African EchiTAb-Plus-ICP antivenom by antivenomics and neutralization approaches. *J Proteome Res*. 2011;10:1266–1280.
4. Bartlett KE, et al. Dermonecrosis caused by spitting cobra snakebite results from toxin potentiation and is prevented by the repurposed drug varespladib. *Proc Natl Acad Sci USA*. 2024;121:e2315597121.
5. Modahl CM, Mukherjee AK, Mackessy SP. An analysis of venom ontogeny and prey-specific toxicity in the monocled cobra (*Naja kaouthia*). *Toxicon*. 2016;119:8–20.

2.3 Additional interface-conditioning tests: JUNO and GPCRs

In addition to edge β -strand pairing binders discussed, I designed other binders with β -strand interfaces to better complement concave and convex target surfaces, hypothesizing that β -strand interfaces would be more amenable to form twists and bending conformations

necessary to shape complement these curved target surfaces. At the time I started design for a convex target (Juno), and concave targets (several GPCRs), the prevailing intuitions (based on previous decades of de novo protein design) preferred helical, compact (low radius of gyration), and rigid protein folds. Designing β -strand interfaces with curved shapes for convex targets or protruding motifs for concave targets challenged these standards. As such, the stability of curved β -strand interface binders remained an outstanding research question for binder design.

2.3.1 JUNO fertilization receptor as a convex-interface design target

While hormonal contraceptives have allowed generations of women to plan their families, they are not suitable for everyone. Side effects and safety concerns can limit acceptability for some users, and there remains a need for non-hormonal approaches that prevent fertilization through orthogonal biological mechanisms¹. Protein binders targeting fertilization-specific receptor interactions could provide one such strategy. Additionally, these binders could potentially be developed as more humane, non-surgical sterilization methods for animals, compared with currently used spaying and neutering procedures². Long-acting contraceptive binders could provide pet owners with additional options and may offer scalable approaches for managing animal overpopulation, including feral cats, dogs, and horses, where surgical sterilization campaigns can be logistically difficult. Finally, the high stability often achievable with designed protein binders could make them attractive candidates for distribution in settings where cold-chain logistics are limiting, although delivery, formulation, and durability would need to be optimized for practical use.

The interaction between egg-cell JUNO and sperm-cell IZUMO1 is essential for mammalian gamete adhesion and fertilization^{3,4}, although recent AlphaFold-Multimer models suggest that this interaction may operate within a broader fertilization synapse involving additional egg and sperm proteins⁵. The IZUMO1-binding surface on JUNO is a convex target surface that

presented an ideal test case for my hypothesis. The computational binder design campaign followed the same principles employed for the design of the β -strand pairing binders: conditioning RFDiffusion with particular secondary-structure interface tensors to complement the target surface, with equal design attempts made for α -helical and β -strand interface approaches. In this case, the binders were conditioned with adjacency to convex JUNO secondary-structure motifs, rather than target edge β -strands.

2.3.2 Binding and functional characterization of JUNO binders

Using yeast surface display enrichment of designs that bound tightly to JUNO, we found that the tightest binding designs indeed exhibited curved β -strand interfaces. Furthermore, the best designs were ProteinMPNN sequence variations of the same binder scaffold shape, supporting the notion that scaffold complementarity played a central role in the binding interface. SPR validation of purified designs yielded low- to mid-nanomolar binding to mouse and human JUNO, with representative variants shown here fitted to apparent K_D values of 31 nM and 42 nM, respectively. Both binders shared similar convex β -strand interface motifs that complemented JUNO at or near the native IZUMO1-binding surface.




Figure 2.3.1. De novo designed binders target the JUNO fertilization receptor. (a) Predicted model of the mammalian egg-sperm fusion synapse, showing egg-cell JUNO and CD9 interacting with sperm-cell IZUMO1, SPACA6, and TMEM81. The model is based on the AlphaFold-Multimer-predicted five-subunit complex described by Elofsson et al. and is not an experimentally solved structure⁵. (b,c) SPR sensorgrams for purified de novo binders binding mouse JUNO (b) or human JUNO (c). Raw binding responses are shown in green and fitted curves are shown in black. Analyte concentrations were 22, 67, 200, and 600 nM. The fitted apparent K_D values were 31 nM for the mouse JUNO binder and 42 nM for the human JUNO binder. Insets show predicted binder-JUNO complexes. (d) In vitro hamster oocyte penetration/fertilization assay measuring fertilization index after treatment with a human anti-JUNO binder. The binder modestly reduced fertilization relative to the untreated control.

The human JUNO binder modestly inhibited human sperm penetration/fertilization of hamster oocytes at 10 $\mu\text{g/mL}$ in vitro. Because this effect was partial, we expected that at least an order-of-magnitude increase in functional inhibition would be required before the binders could be considered promising in vivo contraceptive leads (Figure 2.3.1). We employed multiple approaches to achieve this, which are discussed further in Chapter 4.

References for section 2.3

1. Howard SA, Benhabbour SR. Non-Hormonal Contraception. *Journal of Clinical Medicine*. 2023;12(14):4791. doi:10.3390/jcm12144791.

2. Kutzler M, Wood A. Non-surgical methods of contraception and sterilization. *Theriogenology*. 2006;66(3):514–525. doi:10.1016/j.theriogenology.2006.04.014.
3. Bianchi E, Doe B, Goulding D, Wright GJ. Juno is the egg Izumo receptor and is essential for mammalian fertilization. *Nature*. 2014;508:483–487. doi:10.1038/nature13203.
4. Aydin H, Sultana A, Li S, Thavalingam A, Lee JE. Molecular architecture of the human sperm IZUMO1 and egg JUNO fertilization complex. *Nature*. 2016;534:562–565. doi:10.1038/nature18595.
5. Elofsson A, Han L, Bianchi E, Wright GJ, Jovine L. Deep learning insights into the architecture of the mammalian egg-sperm fusion synapse. *eLife*. 2024;13:RP93131. doi:10.7554/eLife.93131.3.

2.4 Concave-interface design for receptor-penetrating GPCR binders

This section summarizes collaborator-generated results from Muratspahić et al., De novo design of miniproteins targeting GPCRs¹. The specific GPCR binders and experimental data described here were designed, screened, and characterized by collaborators. My contribution to the work was the development of computational design pipelines using RFdiffusion, rather than the design or experimental characterization of the specific binders summarized in this section.

Here, I describe strategies for designing and testing de novo protein binders against G protein-coupled receptors (GPCRs). GPCRs are central drug targets, but they are difficult targets for protein binder design because they are integral membrane proteins, conformationally dynamic, and often present recessed ligand-binding pockets whose functional states differ subtly^{1,2}. For GPCR agonist or antagonist design, binders often must engage ligand-binding pockets or extracellular epitopes that are partially occluded by flexible extracellular loops and N termini. We used multiple design approaches to accommodate these structural constraints. To design binder interfaces that protrude into narrow GPCR ligand-binding pockets, RFdiffusion

outputs were improved by β -hairpin interface conditioning tensors or by multistep motif-guided RFdiffusion approaches, in which a short pocket-compatible motif is generated first and then scaffolded into a full miniprotein (Figure 2.4.1). To model extracellular-loop and N-terminal flexibility, these receptor regions were repredicted during RFdiffusion trajectories using RoseTTAFold. This enabled design trajectories to sample interfaces involving flexible receptor elements rather than treating the receptor as a single rigid structure¹.

Together, these approaches generated miniproteins predicted to reach as deep as ~ 12 Å into class A GPCR orthosteric pockets. These approaches were used in design campaigns for GPCR targets including CXCR4, CCR5, and OXTR at approximately 10^5 design library scale. The resulting libraries were tested using several screening strategies suited to difficult membrane protein targets, including yeast surface display with either lipid nanodisc-embedded GPCRs or GPCRs displayed on mammalian cells. This combination of computational design and high-throughput screening was important because GPCRs are difficult to purify, conformationally dynamic, and often poorly represented by a single static structural model¹.

For the chemokine receptor CXCR4, collaborators generated 25,713 motif-guided RFdiffusion designs targeting the orthosteric site to antagonize CXCL12-mediated signaling. Screening identified five hits, with the best two showing IC_{50} values of 24 ± 5 nM and 120 ± 30 nM (Figure 2.4.1d). The lead binder antagonized CXCR4 signaling with a pA_2 of 7.6 ± 0.3 at 25 nM (Figure 2.4.1e), showed no agonistic activity, and was selective against related receptors including CXCR7 and CCR51. In the broader study, related design and screening approaches also produced GPCR agonists, including MRGPRX1 miniproteins with cryo-EM structures close to the computational design models¹.




Figure 2.4.1. Computational design and characterization of receptor-penetrating GPCR miniprotein binders. (a) GPCR-targeting miniprotein design workflow. Designed backbones were generated either de novo using constrained or scaffold-guided RFdiffusion, or by docking a library of native miniprotein scaffolds. After sequence design and in silico filtering, candidate binders were screened using yeast surface display, biofloating selection, or optical pooled screening-receptor diversion (OPS-RD). (b) Representative RFdiffusion trajectory for generating receptor-penetrating binders against CXCR4. CXCR4 is shown in yellow, selected hotspot residues are shown in pink, and de novo pentamer motifs used for scaffolding are shown in red. The inset shows deep insertion of the motif into the receptor pocket. (c) Computational model of the lead CXCR4 binder dCX1_001 bound to CXCR4. (d) Concentration-response curve for dCX1_001 inhibition of CXCL12-mediated signaling in a cAMP assay using CHO cells stably expressing CXCR4. Data are shown as mean $\pm$ SEM. (e) Functional cAMP assay and Schild regression analysis showing that dCX1_001 antagonizes CXCL12-mediated CXCR4 signaling. Data are shown as mean $\pm$ SEM ($n = 4$); $pA_2 = 7.6 \pm 0.3$, corresponding to 25 nM, with slope 0.68 ± 0.13 . Adapted from Muratspahić et al.¹

These results show that de novo miniproteins can be designed to engage the concave and flexible surfaces of GPCR ligand-binding pockets, producing both antagonists and agonists.

Expanding RFdiffusion beyond its default helical-interface bias using β -hairpin and loop

protrusions, together with flexible receptor modeling, made it possible to target recessed membrane-protein epitopes that would otherwise be poorly sampled by standard binder-design workflows.

References for Section 2.4

1. Muratspahić E, Feldman D, Kim DE, et al. De novo design of miniproteins targeting GPCRs. *Nature*. 2026. doi:10.1038/s41586-026-10656-8.
2. Hauser AS, Attwood MM, Rask-Andersen M, Schiöth HB, Gloriam DE. Trends in GPCR drug discovery: new agents, targets and indications. *Nature Reviews Drug Discovery*. 2017;16:829–842. doi:10.1038/nrd.2017.178.

Chapter 3.

Improving Methods for Binder Sequence Design

3.1 Why sequence optimization is needed after binder design

Although de novo protein design has enabled the generation of binders to diverse targets, initial designs often lack the performance required for downstream applications. Frequently, binding affinity of initial hits are modest, necessitating further improvement to be functionally relevant. Additional limitations like stability, specificity, and cross-reactivity can be addressed alongside affinity maturation. As a result, sequence optimization is a critical step in converting initial designs into biologically active molecules.

Traditionally, optimization has relied on experimental approaches such as site-saturation mutagenesis (SSM) and combinatorial mutagenesis. In SSM, selected positions are mutated to all possible amino acids, generating libraries that probe local sequence space. These libraries

are commonly screened using high throughput platforms such as yeast surface display, where iterative selection enriches mutations that improve binding. The degree of enrichment provides a readout of beneficial substitutions (M) at various sequence positions (N), which can then be combined in subsequent libraries, roughly scaling as M^N , where M represents the average number of allowed amino acid identities per site and N represents the number of mutagenized positions. The exponential size of combinatorial libraries (often 10^6 - 10^8 scale) required to exhaustively screen all combinations of individually beneficial mutations at an interface suffers from high DNA costs of synthesis and sequencing, experimental reagents, and time expended for multiple experimental iterations.

These limitations have motivated the development of computational approaches to sequence optimization, which prioritize promising variants prior to experimental testing. By evaluating sequences on predicted structural and biophysical properties, computational methods reduce the size of experimental libraries and focus resources on higher-confidence candidates. In practice, this enables a more streamlined workflow, with approximately one week for computational screening and ranking of designs (picking the top 100 or so from 10^6 - 10^7 scale design libraries), one week for DNA synthesis, and one week for experimental validation, substantially accelerating the optimization cycle relative to purely experimental approaches. This three week timeline also provides high-quality binding data (SPR and BLI) for purified proteins concomitantly during screening. This truncated timeline is roughly one third of the amount of time required for a binder to undergo SSM and SSM combo optimization with similar purified characterization.

3.2 ProteinMPNN and partial diffusion for local binder optimization

Early on in the deep-learning era of protein design, computational optimization was largely performed by using “partial” RFdiffusion and ProteinMPNN to iteratively improve the AlphaFold

metrics of initial binder hits^{1,2,3,4,5}. This approach works by diversifying interfaces about a known privileged energy landscape. In partial diffusion, the structure of the initial binder hit is partially noised and then denoised again, producing backbone variants that are similar to the input. This enables exploration of a local conformational neighborhood while preserving the overall fold. ProteinMPNN is highly sensitive to these backbone changes, and samples better sequences when input scaffolds are higher quality (which tends to be true in the conformational neighborhood sampled around the initial hit). This approach proved immediately useful for optimizing multiple binding proteins to targets, including the peptides glucagon and neuropeptide Y, as well as TNFR^{6,7}.

While this combined approach improves performance, it remains indirect. Neither ProteinMPNN nor partial diffusion explicitly optimize for binding metrics. Instead, they rely on broadly sampling and filtering around conformational landscapes to identify improved designs. ProteinMPNN has important limitations for affinity maturation. Its training objective is to recover plausible amino acid identities conditioned on a backbone, not to directly optimize AF2-derived interface confidence or binding metrics. As a result, improved binders are usually identified by generating many ProteinMPNN sequences and filtering them with a separate structure-prediction oracle. This limitation motivates methods that directly jointly optimize both sequence and structure toward predicted binding performance, discussed in the next section.

3.3 AlphaFold2 backpropagation for direct sequence optimization

While methods such as ProteinMPNN and partial diffusion enable efficient sampling of sequence and structural space, they do not directly optimize for binding performance. Instead, improved designs are identified indirectly through filtering or enrichment. A natural next step is to ask whether it is possible to optimize binder sequences directly with respect to predicted

structural and binding metrics. The development of AlphaFold2 (AF2) provides a framework for this approach⁴.

AlphaFold2 is a deep neural network that predicts protein structure from sequence by iteratively refining internal representations of residue relationships. The model takes as input an amino acid sequence (and optionally evolutionary information and structure templates) and encodes it into two primary representations: a multiple sequence alignment (MSA) representation, which captures per-residue sequence features, and a pair representation, which captures inter-residue relationships. These embeddings are processed through the Evoformer, a deep stack of attention-based blocks that exchange information between sequence and pair features. The refined representations are then passed to a structure module, which predicts 3D atomic coordinates using invariant point attention and geometric transformations. Importantly, this entire pipeline is end-to-end differentiable, meaning that changes in the input sequence propagate through the network to influence the predicted structure and associated confidence metrics^{4,5}. Because the structure module is formulated as a sequence of differentiable transformations, the full model defines a mapping

$$\mathbf{z} \in \mathbb{R}^{L \times 20}$$

where $\mathbf{z}$ is the $L \times 20$ matrix of sequence logits for a protein of length L . To pass this representation into AlphaFold2, the logits are converted into a relaxed amino acid distribution using a temperature-scaled softmax:

$$p_{i,a} = \frac{\exp(z_{i,a}/\tau)}{\sum_{b=1}^{20} \exp(z_{i,b}/\tau)}$$

Here, $p_{i,a}$ is the relaxed probability of amino acid a at residue position i , and τ is a temperature parameter that controls how sharp or diffuse the distribution is.

With the AlphaFold2 neural network parameters θ held fixed, the model can be treated as a differentiable function that maps the relaxed sequence representation to predicted structural outputs $\hat{\mathbf{Y}}$:

$$\hat{\mathbf{Y}} = f_{\theta}(\mathbf{p}(\mathbf{z}), \mathbf{c})$$

In this expression, $\hat{\mathbf{Y}}$ represents the predicted outputs of AlphaFold2, including atomic coordinates and confidence metrics such as pLDDT and pAE. The term $\mathbf{c}$ represents fixed contextual information provided to the model, such as chain breaks, residue indices, templates, fixed target sequence, or fixed target structural information.

A design objective can then be written as a weighted loss function over the predicted AlphaFold2 outputs:

$$\mathcal{L}(\mathbf{z}) = \sum_r \lambda_r \ell_r (f_{\theta}(\mathbf{p}(\mathbf{z}), \mathbf{c}))$$

Each term ℓ_r represents one design criterion, such as minimizing interface pAE, maximizing binder confidence, maintaining a structural motif, penalizing clashes, or encouraging desired contacts. The coefficient λ_r sets the relative importance of each loss term. This formulation allows multiple design goals to be optimized simultaneously within a single objective function.

Because the loss depends on the sequence logits through the AlphaFold2 prediction, gradients can be backpropagated from the structural loss back to the sequence representation. The sequence logits are then updated by gradient descent:

$$\mathbf{z}_{t+1} = \mathbf{z}_t - \eta_t \nabla_{\mathbf{z}} \mathcal{L}(\mathbf{z}_t)$$

Here, t indexes the optimization step, and η_t is the learning rate. If a mutation would improve the predicted structural objective, the corresponding amino acid logit can increase; if it worsens the

objective, that logit can decrease. Over many optimization steps, the logits become enriched for amino acid identities predicted to improve the design objective.

This differentiability enables a fundamentally different approach to sequence design: rather than sampling sequences and evaluating them, one can optimize sequences by backpropagating gradients through the AF2 network. This idea was made practical through the ColabDesign framework, which exposes AF2 as an optimization engine by allowing gradients of user-defined loss functions to be computed with respect to sequence representations. In this formulation, sequences are represented as continuous logits rather than discrete amino acids, allowing gradient-based updates to modify sequence probabilities at each position. These logits are iteratively updated using optimization algorithms (e.g., gradient descent), and discretized periodically to evaluate candidate sequences.

Crucially, this framework allows optimization toward specific structural objectives, such as minimizing predicted aligned error at the interface, maximizing interface confidence, or enforcing structural constraints. Rather than relying on indirect sampling, the method directly adjusts sequences to improve the predicted quality of the binder-target complex. This transforms sequence design from sampling sequences compatible with a backbone, approximately $P(\text{sequence}|\text{structure})$, into direct optimization of a user-defined loss on predicted structural outputs $f_\theta(\text{sequence})$. This type of continuous sequence optimization is closely related to relaxed sequence optimization, in which gradients are backpropagated through AlphaFold2 into a logit-like sequence representation rather than forcing each update immediately back into a one-hot amino acid sequence (Figure 3.1.1)⁸.




Figure 3.1.1. AlphaFold backpropagation enables protein design in relaxed sequence space. Schematic representation of how relaxed sequence optimization (RSO) enables gradient-descent search in continuous sequence space. In previous approaches, updated relaxed sequence representations were repeatedly converted back into discrete amino acid sequences by argmax operations, which can move the design away from the optimal gradient direction. In RSO, the updated relaxed sequence representation is returned directly to the structure-prediction network for subsequent optimization steps. Adapted from Frank et al.⁸

This de novo strategy contrasts with the interface-focused optimization approach developed here, which begins from an existing binder-target complex and restricts sequence updates to interface residues (Figure 3.1.2). Unlike RFdiffusion followed by ProteinMPNN, which first generates side-chain-free backbones and then designs sequences for those backbones, BindCraft⁹, a method utilizing AlphaFold2 backpropagation for de novo protein binder design, optimizes a binder sequence using gradients from AF2-derived structural and interface losses. Across iterative AF2 predictions, the binder sequence, predicted binder fold, and predicted interface geometry are updated together, allowing the designed binder interface to be shaped directly by the structure-prediction objective. After this AF2-based design stage, BindCraft redesigns the binder surface and core using ProteinMPNN while preserving the optimized interface, then filters designs using AF2 monomer predictions and additional structural and biophysical criteria.

Notably, this approach generated binders with nanomolar affinity from fewer than a dozen tested binder designs per target, ultimately providing an alternative to RFdiffusion+ProteinMPNN for de novo binder design at small scale. Together, these developments demonstrate that structure

prediction models can serve not only as evaluators, but as direct optimization engines for protein design. By enabling gradient-based optimization of sequence with respect to predicted structural outcomes, AF2-based methods overcome key limitations of sampling-based approaches and provide a more metric targeted route to high-affinity binder design.

3.4 Interface-focused AF2 optimization from existing binder complexes

Upon studying the results of BindCraft and other methods published around the same period, including AlphaProteo¹⁰, I sought to apply backpropagation-based methods to a related but distinct task. Rather than designing de novo binders with backpropagation, I focused on redesigning the interfaces of known binder-target complexes. Using binders that we had previously designed with ProteinMPNN, we hypothesized that focusing backpropagation optimization at binder interfaces would enable efficient affinity improvement and outperform optimization by partial RFdiffusion and ProteinMPNN alone (Figure 3.1.2).

We developed an approach for optimizing an input binder-target complex while restricting sequence updates to binder interface residues. Starting from an input complex, the pipeline first repredicts the binder-target structure with AF2-Multimer using the input structure as an initial guess, then selects AF2 models that reproduce the starting complex with low RMSD and favorable interface pAE. Optionally, an initial ProteinMPNN redesign is performed and retained only if it improves the AF2-predicted interface metrics. Interface residues are then identified from the predicted complex using a C α -C α distance cutoff between binder and target residues, typically 8 Å. During gradient-based optimization, non-interface positions are fixed, so only the continuous logits at binder interface residues are updated. The loss function can include AF2-derived terms such as interface pAE, pLDDT, contact losses, distogram losses, or RMSD, depending on the design objective.

To support optimization against cropped or discontinuous target regions, later versions of the pipeline renumbered target residues across structural chain breaks. When consecutive C α atoms in the target chain were separated by more than a distance threshold, downstream residues were offset in residue numbering, preventing AF2 from interpreting spatially separated target segments as covalently continuous. This allowed large targets to be cropped around the binding site while preserving residue-index discontinuities important for AF2 complex prediction.

Within this constrained framework, sequence optimization is performed over a continuous representation of binder sequence logits. These logits can be interpreted as an AF2-derived preference distribution over amino acids at each designable interface position. After optimization, the logits are converted into discrete candidate sequences by argmax selection, probabilistic sampling, or restricted combinatorial enumeration of high-probability amino acid identities. Thus, the method resembles a computationally guided combinatorial mutagenesis strategy: rather than testing all amino acid identities at all interface positions, AF2 backpropagation enriches a smaller set of candidate substitutions that are predicted to improve interface quality. Candidate sequences are then repredicted with AF2-Multimer and ranked by interface confidence metrics before experimental testing. If N interface positions are designable and an average of M amino acid identities are retained per position, the resulting reduced sequence space scales approximately as M^N , compared with 20^N for unrestricted combinatorial sampling.




Figure 3.1.2. Comparison of de novo AF2-backpropagation binder design and interface-focused binder optimization. BindCraft uses AF2-Multimer backpropagation to design de novo binders beginning from a target structure and an initially random or weakly specified binder sequence. During many iterative optimization steps, a relaxed sequence representation is updated to improve AF2-derived structural and interface metrics, and subsequent ProteinMPNN redesign and structure prediction are used to recover and filter realistic discrete binder sequences. In contrast, the interface-focused optimization workflow developed here begins from a known binder-target complex and restricts gradient-based sequence updates to binder residues near the existing interface. This reduces the search space from full de novo binder generation to local interface redesign, allowing shorter optimization trajectories and broad sampling of candidate sequences from the optimized interface logits. Candidate sequences are then repredicted with AF2-Multimer and ranked by interface confidence metrics before experimental testing.

This sequence-sampling strategy is conceptually similar to constructing an SSM-informed combination library, but it does not require an experimental SSM dataset as input. In a traditional SSM workflow, individual substitutions are first screened experimentally, and favorable substitutions are then recombined in a follow-up library. In the AF2-backpropagation workflow, the optimized sequence logits instead provide a computational preference distribution over amino acids at each designable interface position. Sampling from these logits generates combinations of mutations predicted to improve AF2 interface metrics, producing a computationally guided combination library that can be screened directly by structure prediction before experimental testing. This approach forms the basis for the optimization workflows

described in the following sections, where it is applied to affinity maturation and more complex objectives such as cross-reactivity.

3.5 AF2-optimized interface mutations recapitulate SSM landscapes

We observed substantial improvement in AlphaFold metrics for designs optimized by interface-focused gradient descent and next asked whether the optimized sequence preferences resembled experimentally measured mutational landscapes. To test this, we compared AF2-optimized designs to SSM datasets generated by yeast surface display for binders targeting human JUNO and TGFBR1/ALK5. TGFBR1/ALK5 is the type I receptor in the canonical TGF- β signaling pathway involved in development, wound repair, immune regulation, fibrosis, and cancer progression¹¹. Because TGF- β signaling is highly context-dependent and difficult to modulate selectively, designed extracellular binders to TGFBR1 could provide useful tools for perturbing receptor complex assembly or developing pathway-selective inhibitors. For each binder, we compared mutations present in computationally improved designs to mutations enriched in the experimental SSM datasets. AF2-based interface optimization recovered many experimentally favorable substitutions, including mutations that would be selected for SSM-informed combination-library sampling. In contrast, ProteinMPNN redesign recovered fewer of these experimentally enriched substitutions. This agreement suggests that AF2 backpropagation is not merely sampling plausible binder sequences, but is at least partially capturing sequence features associated with improved binding.




Figure 3.1.3. AF2-based interface optimization recovers experimentally enriched SSM mutations. (a) Experimentally determined single-site saturation mutagenesis (SSM) landscapes for binders targeting human JUNO and TGFBR1. Each heatmap shows the measured effect of amino acid substitutions at binder interface positions, plotted as $\log(\text{SC}_{50})$ from yeast surface display selections. Green-highlighted mutations indicate experimentally favorable interface substitutions selected as candidates for SSM-informed combination-library sampling. (b) Computational mutation landscapes from interface-focused AF2 backpropagation or ProteinMPNN redesign. Heatmap intensity represents the number of computationally generated sequences with improved AF2 interface metrics that contain each mutation. Green highlights indicate the same experimentally favorable SSM mutations shown in panel a, enabling direct comparison of computationally recovered substitutions to experimentally enriched substitutions. (c) Counts of experimentally favorable SSM mutations represented among improved computational designs at each interface position. AF2 backpropagation recovered more SSM-enriched substitutions than ProteinMPNN redesign for both JUNO and TGFBR1, suggesting that direct optimization of AF2 interface metrics can prioritize mutations similar to those enriched by experimental affinity selection.

3.6 Experimental validation of AF2-optimized binders

Importantly, the agreement between AF2-optimized mutations and SSM-enriched substitutions translated into measurable improvements in binding performance (Figure 3.1.4). Starting from moderate-affinity binders, interface-focused optimization produced improved variants with only a small number of targeted interface mutations, typically four to eight mutations in this set. Optimization of the human JUNO binder improved affinity from 42 nM to 708 pM, while

optimization of the mouse JUNO binder improved affinity from 31 nM to 12 nM. For TGFBR1/ALK5, AF2-based optimization rescued a weak binder whose initial binding was too poor for reliable kinetic fitting, yielding an optimized variant with an apparent affinity of 243 nM. Similarly, sequence-only optimization of an NKG2D binder improved affinity from approximately 100 nM to 51 nM without altering the binder backbone. The optimized mouse JUNO binder also showed improved functional activity in vitro: when formatted as an Fc fusion, it reduced mouse oocyte pronuclear formation in a dose-dependent manner, whereas the original binder-Fc showed little inhibition under the same assay conditions.




Figure 3.1.4. Interface-focused AF2 optimization improves binder affinity and functional inhibition. (a) SPR sensorgrams for optimized binders targeting human JUNO, mouse JUNO, and TGFBR1/ALK5. Raw binding responses are shown in green and kinetic fits are shown in black. Interface-focused AF2 optimization improved the human JUNO binder to an apparent K_D of 708 pM, the mouse JUNO binder to 12 nM, and rescued a weak TGFBR1/ALK5 binder to 243 nM. (b) In vitro fertilization assay measuring pronuclear formation after treatment with anti-JUNO binders in Fc format. The original mouse JUNO binder-Fc showed little inhibition of pronuclear formation across the tested concentrations, whereas the optimized mouse JUNO binder-Fc reduced pronuclear formation in a dose-dependent manner, reaching approximately 50% pronuclear formation at 50 $\mu\text{g/mL}$. Anti-JUNO antibody was included as a positive inhibition control.

NKG2D is an activating receptor expressed on natural killer cells and subsets of T cells that recognizes stress-induced ligands on infected, transformed, or otherwise damaged cells, making it a useful test case for designing binders that could modulate or redirect immune-cell recognition¹². Because cell-engager formats such as CARs and bispecific T-cell engagers depend on productive receptor clustering and cytotoxic synapse formation at the target-cell surface, improving affinity beyond ~100 nM could increase sensitivity to lower antigen density cells and provide a more stringent benchmark for optimization methods¹³.

To benchmark computational affinity-maturation strategies more rigorously, I compared multiple optimization workflows on a previously validated NKG2D binder with an initial affinity of approximately 100 nM. This test case allowed direct comparison of methods for local binder improvement rather than de novo hit identification. A total of 91 designs was selected for experimental characterization, enabling a controlled head-to-head comparison among several approaches: fixed-backbone AF2 backpropagation alone, ProteinMPNN redesign, an experimental SSM informed virtual combination library ranked by AlphaFold3, partial RFdiffusion followed by ProteinMPNN redesign, partial RFdiffusion followed by ProteinMPNN and AF2-based sequence optimization, interface-motif resampling into new scaffolds followed by ProteinMPNN and AF2-based optimization, and AF3-assisted sequence-design workflows. Ordered designs were selected by standard filtering metrics (iPTM > 0.925, interface_pAE < 7,

monomer_pLDDT > 90). Stricter filtering metrics (iPTM > 0.94, interface_pAE < 6) were required for the partial RFdiffusion + ProteinMPNN + AF2 backprop group (farthest left) due to the overrepresentation of designs that passed the standard filter set. Across these methods, partial diffusion followed by sequence optimization produced the strongest overall outcomes, both computationally and experimentally.

From a computational perspective, designs generated by partial diffusion followed by AF2-based interface optimization achieved the most favorable predicted interface metrics. Experimentally, this same workflow also yielded the highest success rate, with over half of tested designs expressing well and retaining NKG2D binding. Among the designs that bound, most improved on the affinity of the parent binder, and five exhibited at least twofold improvement. The best binder identified from this workflow reached an affinity of 2 nM, representing an improvement of nearly two orders of magnitude over the starting design. In contrast, fixed-backbone sequence optimization alone produced fewer high-affinity hits and did not reach the same level of performance. Other approaches, including the SSM-informed virtual combination library and partial diffusion followed by ProteinMPNN alone, also produced improved binders, but were less consistent and generally yielded weaker best-case affinities. These results are consistent with a division of labor between the methods: partial diffusion expands the local structural search space around the parent binder, whereas AF2-based interface optimization more effectively refines interface sequences within that structural neighborhood than sequence design alone.

Structural comparison of the top performing optimized design with the parent binder suggests a plausible basis for the observed affinity improvement. The optimized design preserves the global binding pose but shows local remodeling at the low pAE hotspot region of the interface. In particular, an interface helix is shifted closer to the target surface, and an F to Q substitution

is predicted to form bidentate hydrogen bonds with target backbone atoms. Together, these changes are consistent with tighter local packing and improved polar complementarity at the binding interface.




Figure 3.1.5. Benchmarking computational workflows for affinity maturation of an NKG2D binder. (a) Head-to-head comparison of multiple popular computational optimization strategies applied to a parent NKG2D binder with a 100nM initial affinity. A testable set of 91 designs was selected for experimental characterization to enable a direct comparison of methods. From left to right, the workflows are: (1) partial diffusion + ProteinMPNN + AF2 backpropagation, (2) interface motif resampling into new scaffolds + ProteinMPNN + AF2 backpropagation, (3) partial diffusion + ProteinMPNN + AlphaFold3 + ProteinMPNN + AlphaFold, (4) an SSM-informed virtual combination library ranked by AlphaFold, (5) AF2 backpropagation only, (6) AF3 and 1000× ProteinMPNN sampling. Red labels indicate the best experimentally measured affinity obtained from each workflow. Partial diffusion followed by AF2-based sequence optimization produced the strongest affinity binder and the highest overall success rate. (b) Structural comparison of the best optimized design from the partial diffusion + ProteinMPNN +

AF2-backpropagation workflow (green) and the parent binder (gray). The full complex overlay shows that the optimized design preserves the overall binding mode while introducing local interface improvements. The zoomed-in insets show the low pAE interface hotspot region that changed the most during optimization, where an interface helix shifts closer to the target surface and an F to Q mutation is predicted to form bidentate hydrogen bonds with target backbone atoms, providing a plausible structural basis for the improved affinity.

These results suggest that, while sequence optimization is effective for refining existing interfaces, the accessible sequence space around a fixed backbone is often insufficient to reach optimal binding configurations. Partial diffusion expands the accessible design space by introducing subtle but meaningful changes to backbone geometry, enabling the formation of new contacts that can then be optimized at the sequence level. Together, these findings highlight that high-affinity binder design often requires coordinated exploration of both structural and sequence space, rather than sequence optimization alone.

3.7 Multitarget AF2 optimization for cross-reactive binder design

While many binder design efforts focus on maximizing affinity toward a single target, many applications require proteins that bind multiple related targets, a property known as cross-reactivity. This is particularly important for translating findings between mouse models and human therapeutics, and for targeting variable pathogen proteins such as influenza hemagglutinin. Designing cross-reactive binders introduces additional constraints because the binder must maintain favorable interactions across homologous or structurally related targets while accommodating differences in local sequence, side-chain chemistry, and backbone geometry.




Figure 3.1.6. Cross-reactive AF2-based interface optimization using shared binder sequence logits. Schematic of the cross-reactive binder optimization workflow. Candidate design campaigns begin from multiple binder-target complex models representing related targets, such as human and mouse homologs or multiple viral subtypes. Input structures are preprocessed, interface residues are selected for optimization, and a shared binder sequence-logit tensor is initialized. In the multitarget optimization mode shown here, the shared binder logits are copied into separate AF2-Multimer models for the primary, secondary, and tertiary targets. Each model computes an AF2-derived interface loss and corresponding sequence gradient, and the target-specific gradients are averaged before a single optimizer update is applied to the shared binder logits. In a complementary single-target control mode, logits are optimized against the primary target only, but the resulting candidate sequences are still evaluated against all targets. After optimization, discrete candidate sequences are generated from the optimized interface logits and repredicted against each target. Final designs are ranked by a cross-reactive score summarizing predicted interface quality across all target complexes, enabling selection of sequences predicted to preserve binding across the full target set.

3.8 Cross-reactive KIT optimization across mouse and human KIT

To evaluate the cross-reactive optimization framework on a therapeutically relevant mouse/human target pair, I applied it to KIT/CD117, a type III receptor tyrosine kinase activated by stem cell factor (SCF)-dependent dimerization and important in hematopoietic, mast cell, melanocyte, germ cell, and interstitial cell of Cajal biology^{14,15}. Cross-reactive targeting of KIT could enable preclinical testing of human-directed designs in mouse models. Starting from a human KIT binder with approximately 10 nM affinity and no detectable mouse KIT binding, optimization produced a variant, KITopt, that bound both human and mouse KIT with apparent

subnanomolar affinity (Figure 3.1.7). Across the broader design campaign, 8 of 42 experimentally tested designs achieved cross-reactive binding with affinities better than 20 nM for both human and mouse KIT. These results show that interface-focused optimization can improve affinity while preserving multi-target constraints, although cross-reactive optimization remains more restrictive than single-target affinity maturation.




Figure 3.1.7. Cross-reactive optimization of a human KIT binder for mouse and human KIT binding. (a) Predicted binder-KIT complexes for the parent binder and optimized KITopt variant against mouse and human KIT. Structures are colored by pLDDT, showing predicted confidence of the binder-target complex. (b) Predicted interface quality for parent binders and optimized child designs evaluated against human and mouse KIT. Stars indicate parent binders and circles indicate optimized designs. Lower min pAE interaction values indicate more confident predicted interfaces. Several optimized designs shifted toward low predicted interface

error for both species, suggesting improved cross-reactive binding potential. (c) SPR validation of parent and optimized KIT binders against mouse and human KIT. The parent binder showed no detectable binding to mouse KIT but retained binding to human KIT, whereas KITopt showed concentration-dependent binding to both mouse and human KIT with apparent subnanomolar affinity. Raw binding responses are shown in green and kinetic fits are shown in black.

3.9 Functional characterization of KITopt binders

Functional assays further revealed important nuances in how cross-reactive binding translated to downstream activity (Figure 3.1.8). First, collaborator-generated preliminary biophysical data suggested that bivalent KITopt/KITag constructs can dimerize the human KIT extracellular domain in solution. Mass photometry detected a higher-molecular-weight species consistent with a 2:1 KIT:minibinder complex, supporting the intended mechanism of receptor dimerization. In cultured human cells, representative western blotting showed detectable KIT phosphorylation after stimulation with KITag constructs, consistent with receptor activation.

The optimized KIT binders were also tested as targeting modules in two application formats. In an ELP nanoparticle uptake assay using CD117-positive P815 cells, KITag14 functionalization produced up to a 2.38-fold increase in AF647-siRNA uptake, whereas a non-binding design control did not increase uptake under the same conditions. This suggests that the optimized binder can support receptor-mediated particle uptake, although performance remained lower than the best mouse SCF-functionalized particles and will likely depend on binder valency, linker geometry, and formulation conditions.

In a separate preliminary CAR T cell assay, collaborator-generated data showed that KITopt binder CARs mediated measurable mouse tumor-cell killing. Neither binder CAR reached the activity of the CD19-CAR positive control. These preliminary CAR T cell results should be interpreted cautiously because the CD19-CAR positive control targets a highly validated antigen-receptor system, whereas CD117/KIT may present different constraints related to

antigen density, epitope accessibility, receptor geometry, and immune synapse formation. Thus, the lower activity of the KITopt binder CARs does not necessarily indicate that binder affinity is insufficient, but suggests that CD117-directed CAR activity may require additional optimization of hinge length, spacer geometry, and overall CAR architecture.




Figure 3.1.8. Functional characterization of cross-reactive KITopt binders. (a) Mass photometry analysis of human KIT extracellular domain in the absence or presence of a bivalent dual KITopt binder construct. KIT alone is primarily observed near the expected monomeric mass, whereas addition of the minibinder produces a higher-molecular-weight species

consistent with formation of a 2:1 KIT:minibinder complex. (b) Representative western blot measuring KIT phosphorylation after stimulation with KITag minibinder constructs. Cell lysates were probed with anti-phospho-KIT antibody and anti-GAPDH loading control. The Ag14/KITopt construct induced detectable KIT phosphorylation across the tested concentrations relative to negative controls. (c) Schematic of receptor-targeted elastin-like polypeptide (ELP) nanoparticles for nucleic acid or gene-editor delivery. (d) Uptake of AF647-siRNA formulated with CD117-targeted V5(S)-ELP-EEP13 nanoparticles in P815 cells. KITag14 functionalization increased uptake up to 2.38-fold, with strongest activity at low targeting-domain valency, compared to a negative control. (e) Preliminary CAR T cell cytotoxicity assay using KITopt binder CARs. KITopt-mediated measurable tumor-cell killing at high effector:target ratios, but activity was lower than a CD19-CAR positive control. Panels adapted from collaborator preliminary progress materials.

3.10 Cross-reactive HA stalk binder design optimization across H1, H3, and H5

As a final test of the cross-reactive AF2 optimization framework, I applied the method to a highly ambitious therapeutic design problem: generating a minibinder that could bind hemagglutinin (HA) from H1, H3, and H5 influenza A viruses. The goal was to target a conserved HA region that could potentially block infection across seasonal human subtypes and pandemic-risk avian subtypes. Broadly neutralizing antibodies have shown that conserved HA stalk epitopes can support cross-subtype recognition across group 1 and group 2 influenza A viruses, and more recent work has identified membrane-proximal anchor-epitope antibodies that contact conserved residues near the HA fusion peptide^{16,17}. This site was attractive because it is functionally constrained and contains an exposed edge β -strand, which could in principle support sequence-independent backbone hydrogen-bonding interactions. However, it was also a difficult design target because the epitope is membrane-proximal, partially occluded, glycan-proximal, and sensitive to whether HA is modeled as cleaved HA1/HA2 or uncleaved HA0, a cleavage transition known to affect HA structure¹⁸. Moreover, the design campaign intended to design cross-reactive binders de novo, rather than simply optimizing existing binders.

Computationally, this campaign used the same cross-reactive design logic developed for KIT, but with a more difficult target set: H1, H3, and H5. I generated candidate binders against the HA stalk, filtered for designs predicted to maintain binding geometry across all three subtypes, and applied stringent final filters including AF2-Multimer iPTM > 0.85 and AF3 min-pAE interaction < 1.5. Because commercially available biotinylated HA constructs were primarily uncleaved HA0, I also re-evaluated designs against uncleaved HA models. This introduced an additional practical complication, but many computationally passing candidates retained predicted compatibility with the available experimental constructs.

Experimentally, the first screen identified a weak cross-reactive signal rather than a mature high-affinity binder. One design, HA41, showed measurable SPR binding signals to H1, H3, and H5 at 2 μ M binder. I then attempted to optimize this initial hit using the cross-reactive AF2 optimization strategy. After a second round of optimization, several HA41-derived variants showed stronger binding to H3, with apparent affinities in the low-micromolar range (Figure 3.1.9). However, the optimized mutations did not yield cleanly quantifiable binding to H1 or H5; for those subtypes, binding remained weak and was only observed at high micromolar binder concentrations.




Figure 3.1.9. Cross-reactive optimization of an HA stalk binder across H1, H3, and H5. (a) Computational evaluation of candidate HA binders against H1, H3, and H5 hemagglutinin targets. Designs generated by partial diffusion alone or by partial diffusion followed by AF2-based sequence optimization were evaluated by AF2-Multimer interface metrics. Each point represents a candidate binder-target prediction, plotted by iPTM and pAE interaction. Lower pAE interaction and higher iPTM indicate more favorable predictions. (b) SPR sensorgrams for a second-round optimized HA41-derived binder tested against H1, H3, and H5 HA. Raw binding responses are shown in green and fitted curves are shown in black. Binder concentrations were 40 nM, 200 nM, 1 μ M, and 5 μ M. The optimized binder showed the strongest binding response to H3, while H1 and H5 binding remained weak and was only clearly detectable at high binder concentrations. These results indicate partial cross-reactive binding but not the high-affinity pan-HA activity required for an antiviral lead.

Overall, this campaign was a partial success. It demonstrated that the cross-reactive optimization framework could identify and preserve weak binding to a conserved HA stalk site across a difficult subtype set, and that optimization could improve binding to at least one subtype, especially H3. However, it did not produce the strong, clean, broadly reactive binder needed for a convincing antiviral lead. This result is scientifically useful because it highlights a boundary of the current method: cross-reactive optimization can push designs toward multi-target compatibility, but for HA, additional work is likely needed in scaffold generation,

epitope accessibility modeling, HA0 versus HA1/HA2 state handling, and experimental screening scale before this approach can reliably produce infection-blocking pan-HA binders.

References for Chapter 3

1. Dauparas J, Anishchenko I, Bennett N, et al. Robust deep learning-based protein sequence design using ProteinMPNN. *Science*. 2022;378:49–56.
doi:10.1126/science.add2187.
2. Watson JL, Juergens D, Bennett NR, et al. De novo design of protein structure and function with RFdiffusion. *Nature*. 2023;620:1089–1100.
doi:10.1038/s41586-023-06415-8.
3. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*. 2021;596:583–589. doi:10.1038/s41586-021-03819-2.
4. Evans R, O'Neill M, Pritzel A, et al. Protein complex prediction with AlphaFold-Multimer. *bioRxiv*. 2022. doi:10.1101/2021.10.04.463034.
5. Bennett NR, Coventry B, Goreschnik I, et al. Improving de novo protein binder design with deep learning. *Nature Communications*. 2023;14:2625.
doi:10.1038/s41467-023-38328-5.
6. Vázquez Torres S, Leung PJY, Venkatesh P, et al. De novo design of high-affinity binders of bioactive helical peptides. *Nature*. 2024;626:435–442.
doi:10.1038/s41586-023-06953-1.
7. Glögl M, Krishnakumar A, Ragotte RJ, et al. Target-conditioned diffusion generates potent TNFR superfamily antagonists and agonists. *Science*. 2024;386:eadp1779.
doi:10.1126/science.adp1779.
8. Frank C, et al. Scalable protein design using optimization in a relaxed sequence space. *Science*. 2024;386:439–445. doi:10.1126/science.adq1741.

9. Pacesa M, et al. One-shot design of functional protein binders with BindCraft. *Nature*. 2025;646:483–492. doi:10.1038/s41586-025-09429-6.
10. Zambaldi V, La D, Chu AE, et al. De novo design of high-affinity protein binders with AlphaProteo. *arXiv*. 2024. arXiv:2409.08022.
11. Derynck R, Budi EH. Specificity, versatility, and control of TGF- β family signaling. *Science Signaling*. 2019;12(570):eaav5183. doi:10.1126/scisignal.aav5183.
12. Raulet DH, Gasser S, Gowen BG, Deng W, Jung H. Regulation of ligands for the NKG2D activating receptor. *Annual Review of Immunology*. 2013;31:413–441. doi:10.1146/annurev-immunol-032712-095951.
13. Stone JD, Aggen DH, Schietinger A, Schreiber H, Kranz DM. A sensitivity scale for targeting T cells with chimeric antigen receptors (CARs) and bispecific T-cell engagers (BiTEs). *Oncotarget*. 2012;1(6):863–873. doi:10.4161/onci.20592.
14. Lennartsson J, Ronnstrand L. Stem cell factor receptor/c-Kit: from basic science to clinical implications. *Physiological Reviews*. 2012;92:1619–1649. doi:10.1152/physrev.00046.2011.
15. Yuzawa S, Opatowsky Y, Zhang Z, Mandiyan V, Lax I, Schlessinger J. Structural basis for activation of the receptor tyrosine kinase KIT by stem cell factor. *Cell*. 2007;130:323–334. doi:10.1016/j.cell.2007.05.055.
16. Corti D, Voss J, Gamblin SJ, et al. A neutralizing antibody selected from plasma cells that binds to group 1 and group 2 influenza A hemagglutinins. *Science*. 2011;333:850–856. doi:10.1126/science.1205669.
17. Guthmiller JJ, Stovicek O, Wang J, et al. Broadly neutralizing antibodies target a haemagglutinin anchor epitope. *Nature*. 2022;602:314–320. doi:10.1038/s41586-021-04356-8.
18. Chen J, Lee KH, Steinhauer DA, Stevens DJ, Skehel JJ, Wiley DC. Structure of the hemagglutinin precursor cleavage site, a determinant of influenza pathogenicity and the

origin of the labile conformation. *Cell*. 1998;95:409–417.

doi:10.1016/S0092-8674(00)81771-7.

Chapter 4. Discussion

4.1 Designing binders to structurally heterogeneous targets

While the methods described in the previous chapters demonstrate that computational design can produce high-affinity binders for well-behaved protein targets, many biologically important targets present different challenges. These include structural heterogeneity, conformational flexibility, variable post-translational modification dependence, and poorly defined binding epitopes. To illustrate these challenges, I will describe attempts to design binders targeting CD8 β , TGF β R1/ALK5, and pathological tau fibrils (Figure 4.1.1). Binder campaigns were ultimately successful for CD8 β and TGF β R1, but only at low hit rates and with substantial dependence on high-throughput screening. These targets demonstrate that, in the deep-learning era, a major remaining bottleneck is not only designing good interfaces, but accurately modeling the target state that is actually present in the experiment.

4.2 CD8 β binder design highlights flexible target-state uncertainty

CD8 β is part of the CD8 co-receptor complex expressed on cytotoxic T cells, where it contributes to peptide-MHC class I recognition and T-cell signaling¹. Because no experimental structure was available for the relevant CD8 β target construct, I generated an AlphaFold2 model for design. Early small-scale design campaigns, totaling approximately 200 designs across multiple target sites and interface topologies, did not yield detectable binding. Binders were eventually identified only after high-throughput screening of a larger library ($n = 17,500$) that included designs targeting both the structured CD8 β extracellular domain and a flexible membrane-proximal coil. The successful designs were enriched for similar modeled coil

conformations, suggesting that the sampled coil geometry was important for productive binding. This result supports the idea that flexible target regions can be designable, but only if the relevant conformations are sampled accurately enough during scaffold generation.

4.3 TGF β R1/ALK5 binder design highlights loop-state uncertainty

TGF β R1/ALK5 presented a related but distinct challenge. TGF β R1 is the type I receptor kinase in the canonical TGF- β signaling pathway, which regulates development, tissue homeostasis, immune regulation, fibrosis, and cancer progression². Extracellular binders could provide useful tools for perturbing receptor complex assembly or modulating pathway activation. However, only one weak binder was identified from a high-throughput screen of a 5,997-member design library. That initial hit was later improved by the AF2-based interface-optimization methods described in Chapter 3, yielding an optimized binder with an apparent affinity of 243 nM. Without the initial weak hit, however, sequence optimization would not have been possible. During the original scaffold-design campaign, flexible target-site loops were either held fixed or allowed to remodel during design based on experimental evidence for local conformational flexibility.

Retrospectively, the validated binder involved only subtle target-site changes, including minor loop reorientation and side-chain rotamer changes near the interface. This suggests that some conformational sampling was useful, but that the design library likely sampled a broader range of target conformations than were experimentally relevant.

Together, the CD8 β and TGF β R1 campaigns show that binder design can become limited by target-state uncertainty. Structure-prediction models are often accurate for ordered domains, but they remain difficult to interpret for disordered regions, flexible loops, coils, and partially ordered membrane-proximal segments³. In these cases, the challenge is not simply generating a

complementary binder surface, but identifying the target conformation that should be used for design and experimental screening.

4.4 Tau fibril binder design highlights polymorph and sample heterogeneity

Tau fibrils represent an even more difficult target class. Tau is mostly disordered as a soluble monomer, but can assemble into ordered amyloid fibrils. In healthy neurons, soluble tau associates with and stabilizes microtubules, whereas pathological tau can assemble into fibrils associated with neurodegenerative disease. These fibrils can adopt distinct conformations, or polymorphs, that vary with disease state and assembly conditions. Unlike globular protein targets, tau fibrils contain both ordered cores and disordered “fuzzy coat” regions. In Alzheimer’s disease tau filaments, the ordered paired helical filament and straight filament cores are composed primarily of tau residues 306–378, while disordered N- and C-terminal regions project away from the core⁴. Binding epitopes may therefore depend on tau isoform, disease state, fibril polymorph, and post-translational modifications such as phosphorylation or lysine modification.

These features complicate both structural modeling and experimental validation.

Computationally, the large number of tau molecules present in a fibril prevents routine modeling of the full assembly, and the disordered fuzzy coat remains poorly described by current structure-prediction methods. Experimentally, targeting a specific tau fibril polymorph is challenging because routine production of fibrils that faithfully reproduce patient-derived disease polymorphs remains specialized and nontrivial, although recent recombinant assembly methods have made important progress toward this goal⁵. As a result, structurally relevant samples often require extraction from diseased human tissue. These samples are insoluble, heterogeneous,

and difficult to purify, making typical recombinant expression, purification, and binding-assay workflows poorly suited to the target.

For Alzheimer's disease tau fibrils, collaborators in Brian Kraemer's lab prepared biotinylated samples for BLI experiments because SPR microfluidics was not suitable for insoluble fibril material. Tau aggregates were extracted from Alzheimer's disease brain tissue and partially purified by centrifugation. Biotinylation was attempted by incubating the fibril preparation with separately synthesized biotinylated tau, allowing biotinylated tau to incorporate into or associate with fibril material. These samples likely still contained heterogeneous fibrils and contaminating brain-derived material, but they represented the most tractable approximation of biotinylated patient-derived Alzheimer's disease tau fibrils available for binding experiments. Binding of an anti-tau antibody produced a strong BLI signal, confirming that the samples could be captured and interrogated by BLI.

Initial computational screening of approximately 250 tau-targeting designs yielded only a small number of weak hits. Subsequent optimization strategies included interface-focused sequence optimization, partial diffusion, and charge-biased sequence sampling intended to better complement the surface shape and negative charge associated with phosphorylated pathological tau. These approaches produced several candidates with measurable BLI responses, particularly within a moderate positive-charge regime. However, the designed-binder responses were weak relative to the anti-tau antibody control, did not show clean kinetic behavior, and may reflect nonspecific electrostatic interactions or buffer artifacts rather than specific, high-affinity binding. Therefore, these ongoing results are not shown here and are best interpreted as ambiguous preliminary signal rather than validated tau-fibril binding.




Figure 4.1.1. Binder design challenges arising from target conformational uncertainty and structural heterogeneity. (a) CD8 β binder design required modeling both the structured extracellular domain and a flexible membrane-proximal coil. Multiple sampled conformations illustrate uncertainty in the target state used for design. Successful binders were identified only after screening designs that incorporated this flexible region, and the resulting hits were enriched for similar modeled coil conformations. (b) TGF β R1/ALK5 binder design was complicated by local flexibility near the target binding site. The validated binder model involved subtle target-site remodeling, including loop reorientation and side-chain rotamer changes near the interface, suggesting that accurate local conformational sampling can determine whether a designed interface is compatible with the experimentally accessible receptor state. (c) Tau fibrils present a distinct challenge because the ordered fibril core is embedded within disordered fuzzy-coat regions and exists in disease-associated polymorphic assemblies. Schematic representation of full-length Alzheimer's disease tau filaments shows disordered fuzzy-coat regions surrounding the ordered filament core; adapted from Fitzpatrick et al.⁴ Right: structural model of the ordered tau filament core used as a tractable design target.

This outcome highlights a critical limitation of current computational design approaches:

success depends not only on the ability to model protein-protein interfaces, but also on the validity of the structural model and the biophysical tractability of the target. In the case of tau fibrils, uncertainty in epitope definition, structural heterogeneity, sample preparation, and surface-charge effects complicate both design and interpretation. These challenges are not

failures of the design algorithms alone, but reflect the underlying biological and experimental complexity of the system.

4.5 Why computational design remains valuable despite difficult targets

Despite these challenges, computational design provides several key advantages when working with complex targets such as CD8 β , TGF β R1, and tau fibrils. First, it enables systematic exploration of sequence and structural space that would be difficult or impossible to search experimentally. Even when only a small fraction of designs produce weak or ambiguous signal, these results can identify starting points for further refinement or reveal limitations in the original design hypothesis. Second, computational methods allow precise control over design assumptions, such as targeting specific conformational states, enforcing structural constraints, biasing charge distributions, or focusing on defined post-translationally modified surfaces. These hypotheses can be difficult to implement cleanly in traditional library-based approaches.

Third, computational design can substantially reduce experimental burden. Rather than screening millions of variants through high throughput display systems, design pipelines can prioritize tens to hundreds of candidates with predicted structural compatibility. This is particularly important for targets such as fibrils or flexible receptors, where assay development is slow, reagents are limited, and experimental throughput is inherently constrained. Finally, even unsuccessful or ambiguous campaigns can be informative. They help define which target features are designable, which interaction modes may be favored, and where current modeling assumptions break down.

Taken together, these results suggest that computational binder design is not only a route to high-affinity binders, but also a framework for probing the designability of complex biological surfaces. In this sense, difficult design targets are useful not only as applications, but also as stress tests that reveal which aspects of protein structure, conformational heterogeneity, and assay tractability remain limiting.

4.6 Future applications of designed protein binders

The ability to design protein binders with defined affinity, specificity, and structural properties enables a wide range of applications across biology and medicine. In therapeutic contexts, binders can function as receptor agonists or antagonists, block protein-protein interactions, or serve as targeting domains for drug delivery systems. For example, binders developed in this work were explored as components of receptor-targeting molecules, cross-reactive therapeutic leads, and modular cell-targeting systems, where affinity, specificity, and binding geometry directly influence downstream function.

Binders also enable the construction of modular, multi-functional proteins. By combining targeting domains with effector domains, oligomerization modules, degradation tags, delivery vehicles, or immune-cell signaling components, it is possible to design molecules that induce receptor clustering, mediate cell-cell interactions, or selectively deliver payloads. These principles underlie emerging technologies such as synthetic cytokine mimetics, targeted protein degradation systems, receptor-targeted delivery platforms, and immune cell engagers.

In diagnostic and research settings, binders provide tools for detecting specific protein states, including conformational variants and post-translationally modified epitopes. This is particularly relevant for diseases such as neurodegeneration, where distinguishing between protein isoforms, aggregation states, or disease-associated fibril conformations is critical. Although the tau binder results presented here remain preliminary, they illustrate how computational design can be used to formulate and test molecular hypotheses about difficult disease-associated structures.

More broadly, computational binder design enables a shift from discovery-based to design-based molecular engineering. Instead of relying only on naturally occurring binding proteins or large experimental libraries, researchers can now generate bespoke binding molecules tailored to specific biological questions. The results in this dissertation show both the power and the remaining limitations of this approach. Designed binders can now be generated and optimized against diverse targets, but success still depends on accurately modeling the relevant target state, sampling compatible interface geometries, and validating designs in assays that reflect the intended biological context. As modeling accuracy improves and design methods become more robust to conformationally heterogeneous and experimentally challenging targets, computational binder design is likely to play an increasingly central role in both basic research and translational protein engineering.

References for Chapter 4

1. Bosselut R, Kubo S, Guinter T, Kopacz JL, Altman JD, Feigenbaum L, Singer A. Role of CD8 β domains in CD8 coreceptor function: importance for MHC I binding, signaling, and positive selection of CD8⁺ T cells in the thymus. *Immunity*. 2000;12:409–418.
doi:10.1016/S1074-7613(00)80193-4.

2. Derynck R, Budi EH. Specificity, versatility, and control of TGF- β family signaling. *Science Signaling*. 2019;12(570):eaav5183. doi:10.1126/scisignal.aav5183.
3. Ruff KM, Pappu RV. AlphaFold and implications for intrinsically disordered proteins. *Journal of Molecular Biology*. 2021;433(20):167208. doi:10.1016/j.jmb.2021.167208.
4. Fitzpatrick AWP, Falcon B, He S, et al. Cryo-EM structures of tau filaments from Alzheimer's disease. *Nature*. 2017;547:185–190. doi:10.1038/nature23002.
5. Lövestam S, Koh FA, van Knippenberg B, Kotecha A, Murzin AG, Goedert M, Scheres SHW. Assembly of recombinant tau into filaments identical to those of Alzheimer's disease and chronic traumatic encephalopathy. *eLife*. 2022;11:e76494. doi:10.7554/eLife.76494.