



PROTEIN DESIGN

Top-down design of protein architectures with reinforcement learning

Isaac D. Lutz^{1,2,3,†}, Shunzhi Wang^{1,2,†,*}, Christoffer Norn^{1,2,4,†}, Alexis Courbet^{1,2,5}, Andrew J. Borst^{1,2}, Yan Ting Zhao^{1,6,7}, Annie Dosey^{1,2}, Longxing Cao^{1,2,8}, Jinwei Xu^{1,2}, Elizabeth M. Leaf^{1,2}, Catherine Treichel^{1,2}, Patrisia Litvicov^{1,6}, Zhe Li^{1,2}, Alexander D. Goodson^{1,2}, Paula Rivera-Sánchez⁴, Ana-Maria Bratovianu⁴, Minkyung Baek^{1,2,9}, Neil P. King^{1,2}, Hannele Ruohola-Baker^{1,3,6,7}, David Baker^{1,2,3,*}

As a result of evolutionary selection, the subunits of naturally occurring protein assemblies often fit together with substantial shape complementarity to generate architectures optimal for function in a manner not achievable by current design approaches. We describe a “top-down” reinforcement learning-based design approach that solves this problem using Monte Carlo tree search to sample protein conformers in the context of an overall architecture and specified functional constraints. Cryo-electron microscopy structures of the designed disk-shaped nanopores and ultracompact icosahedra are very close to the computational models. The icosahedra enable very-high-density display of immunogens and signaling molecules, which potentiates vaccine response and angiogenesis induction. Our approach enables the top-down design of complex protein nanomaterials with desired system properties and demonstrates the power of reinforcement learning in protein design.

Multisubunit protein assemblies play critical roles in biology and are the result of evolutionary selection for function of the entire assembly. Therefore, the subunits in structures such as icosahedral viral capsids often fit together almost perfectly (1, 2). In contrast to direct evolutionary selection on overall system properties, de novo protein design has generated protein architectures using a “bottom-up” hierarchical approach (Fig. 1A, left) in which monomeric structures are first docked into symmetric oligomers (3–6) and then assembled into closed assemblies with tetrahedral, octahedral, or icosahedral symmetry (7–14) or open assemblies such as two-dimensional (2D) layers and 3D crystals (15–19). An advantage of this hierarchical approach is that the multiple interfaces that stabilize the assembly can be validated independently (the first by characterization of the symmetric oligomer and the second by characterization of the nanomaterial assembly from the preformed oligomer), considerably increasing the robustness of the overall design process. Although such designed assemblies are already proving use-

ful for biomedicine in immunobiology and other areas, as highlighted by the recent approval of a de novo–designed COVID vaccine (20–23), the bottom-up approach does have limitations. The properties of the assembly are limited to what can be generated from the available oligomeric building blocks, at least one of the subunit-subunit interfaces must be strong enough to stabilize a cyclic oligomeric substructure in isolation, and, more generally, there is no way to directly optimize the properties of the overall assembly.

We sought to overcome the limitations of bottom-up protein complex design by developing a top-down approach (Fig. 1A, right) that starts from a specification of the desired properties (overall symmetry, porosity, etc.) of the structure and systematically builds up subunits that pack together to optimize these properties. We reasoned that protein fragment assembly (24–28), which can generate a wide variety of monomeric protein structures, could provide a suitable mechanism for generating diversity. Previous design approaches such as SEWING have built up proteins from fragments, optimizing for monomer stability at each step (29), but we aimed instead to optimize for overall system properties, which could involve trading off monomer stability for increased subunit-subunit interaction strength and other properties. To enable such end state–based optimization, we turned to reinforcement learning (RL), which has achieved considerable success recently in different fields of artificial intelligence, such as self-driving cars (30), the AlphaGo program that defeats top human players in the game of Go (31, 32), and algorithm development (33). Monte Carlo tree search (MCTS) (34, 35) is an RL algorithm that finds optimal

series of choices within a search tree. In MCTS, choices are selected randomly at each branch point to find a path down the tree, and after exploring a path, the state is evaluated, and probabilities at each branch point back-propagated up the tree are reweighted accordingly such that subsequent iterations are more likely to lead to optimal paths.

Backbone sampling by MCTS

We sought to develop a MCTS algorithm for generating protein complexes that builds up the monomeric subunits from protein fragments directly optimizing for prespecified global structural properties. We set up the tree search such that at each step in the tree, a short protein fragment is appended at either the N terminus or C terminus of the growing chain. The number of fragments to consider at each step is a trade-off between the rapidity of learning (with a smaller number, weights on each choice can be learned more quickly) and the total diversity of structures that can be generated (which increases with the number of choices at each step). We chose to balance these factors by using as building blocks parametrically generated straight helices, which are fully described by a single parameter (the length, which we allow to vary from nine to 22 residues), followed by short loops clustered into 316 bins (derived from clustering loops in a large helical protein database; see the materials and methods). The search begins with the selection of one of the helix possibilities and then alternates between the addition of a loop or a helix choice at either terminus. Once a loop bin is chosen, we select randomly from the closely related loop backbones within the cluster (Fig. 1B, left). Although this is a far narrower set of local structures than observed in native protein structures, we found in preliminary explorations that a wide variety of compact protein shapes could be readily generated from such building blocks. Building up a 100-residue protein backbone with this approach requires about five helix and four loop additions, yielding a total number of possibilities of $\sim 1 \times 10^{17}$, with additional structural diversity from the variation in loop backbones within a bin. The size of the search tree grows exponentially with the number of structural elements, so the space of possibilities is more effectively explored for monomers with fewer helices than for larger monomers.

The search is modulated based on the specific problem specification through geometric constraints that are applied at each step in the search tree and score functions that are evaluated only after full structures are completed. Potential moves consisting of helix or loop fragments are selected at each level of the search tree only if they pass geometric constraints that can be evaluated before the

¹Department of Biochemistry, University of Washington, Seattle, WA, USA. ²Institute for Protein Design, University of Washington, Seattle, WA, USA. ³Department of Bioengineering, University of Washington, Seattle, WA, USA. ⁴BiolInnovation Institute, DK2200 Copenhagen N, Denmark. ⁵Howard Hughes Medical Institute, University of Washington, Seattle, WA, USA. ⁶Institute for Stem Cell and Regenerative Medicine, University of Washington, Seattle, WA, USA. ⁷Oral Health Sciences, University of Washington, Seattle, WA, USA. ⁸Key Laboratory of Structural Biology of Zhejiang Province, School of Life Sciences, Westlake University, Hangzhou, Zhejiang, China. ⁹School of Biological Sciences, Seoul National University, Seoul, Republic of Korea. *Corresponding author. Email: dabaker@uw.edu (D.B.); swang523@uw.edu (S.W.)

†These authors contributed equally to this work.




Fig. 1. Top-down design strategy and computational pipeline. (A) Bottom-up (left) and top-down (right) strategies to protein assembly design. (B) (Left) MCTS architecture for monomer backbone generation. During each simulation, a helix stub is initialized at a random rigid body start position, and different configurations of helices and loops are sampled and constructed sequentially with probability P_t stored in each edge to build the search tree. Each move is checked against a set of predefined geometric constraints during the expansion stage and then updates probabilities P'_t afterward. Upon successful completion of a search tree, the monomer is evaluated by score functions

and probabilities P'_t are back-propagated to update all of the search tree edges. (Right) Symmetric transformations are applied to build an icosahedral capsid in parallel with monomers using the MCTS generative algorithm. (C) Concurrent geometric check (left) is performed at every step of the expansion stage and the search tree is terminated if there are violations. Final evaluation (right) with a series of score functions is performed upon completion of a simulation for monomers and assemblies. (D) In silico RL-generated capsids (blue) occupy a distinct structural space compared with de novo-designed protein cages (red) and natural capsids (green).

assembly of the entire structure; these include internal clashes and overall shape constraints (see the materials and methods for a full list of geometric constraints). Upon selection of a move passing the geometric constraints, its probability is upweighted, as are the probabilities of all prior moves leading to this point

in the search tree. Completed backbones are evaluated using score functions that assess how well the overall generated structure satisfies the user specification of the problem to be solved (Fig. 1C and materials and methods), and the probabilities of selection of each move at each step along the search tree are

reweighted accordingly. As individual move weights become increasingly biased after many traversals through the search tree, the generated complete backbones have higher and higher scores (fig. S1). Because each iteration takes on average only tens of milliseconds, high-scoring backbones can be sampled

at scale by searching over tens of thousands of iterations. To address the classical RL problem of balancing exploration with exploitation (30–32), the search is initialized from many independent trees, and the maximum probability of any one move is capped (see the materials and methods).

We first tested the MCTS approach *in silico* at the protein monomer level, choosing as a test problem the generation of protein backbones with arbitrarily prespecified overall shapes. To our knowledge, there are no current approaches for addressing this problem. A specified build volume is represented on a grid, and the MCTS is initialized randomly within the volume. At each move, only additions that stay within the specified volume are accepted. For a range of prescribed shapes, including regular polyhedra and letters from the alphabet, the ensembles of generated structures closely fill the specified volumes, and individual backbones have the prespecified shapes (fig. S2). The average sequence length of the solutions increases through the optimization as the choices of moves and combinations of moves that lead to satisfaction of the input constraints are learned, enabling traversal further down the search tree (fig. S1A).

We next sought to generalize the MCTS to the design of symmetric nanomaterials by applying symmetry operators to generate assemblies with the desired symmetry at each step in the search tree. Each move (helix or loop addition) is assessed by considering not only the growing monomer, but also its interactions with all nearby symmetry mates, computed using transformation matrices specifying each symmetry operator; moves that introduce steric clashes are discarded (Fig. 1, B and C). We tested these capabilities *in silico* by designing cyclic assemblies with symmetries C5 through C12, as well as tetrahedral, octahedral, icosahedral, and quasisymmetric icosahedral assemblies of up to 240 subunits (figs. S3 and S4). We found that by providing different geometric constraints and score functions to guide the search, we could control properties such as shape, size, porosity, and termini position from the top down (figs. S3 to S6).

Nanopore construction using constrained symmetric MCTS

As a first experimental test of the MCTS approach, we applied it to the highly constrained design challenge of filling the space between two previously designed cyclic protein rings (6, 36) to generate disk-shaped structures with a central nanopore (Fig. 2A). Filling this substantial but irregularly shaped space such that there are no large voids between the two rings is not straightforward with previously described protein design methods. We approached this challenge with MCTS by geometrically constraining the search to the space

between the two rings, requiring dense packing such that the only large void in the resulting assembly is the pore of the inner C6 ring. Both the inner and the outer ring have C6 symmetry, and the search tree was initialized to start at the N termini of the outer ring and simultaneously build six subunits that collectively fill the empty space. We performed the MCTS for each of 2000 placements of a set of different inner rings with a range of inner pore sizes inside a constant outer ring (for each inner ring, we sampled rotations around and translations along the common cyclic symmetry axis). We selected backbones that fully filled the space between the two rings, designed sequences with ProteinMPNN (37), and selected for experimental characterization 32 designs predicted to assemble into the designed assemblies by AlphaFold (AF) (38). Of these, we found that 28 were soluble and could be purified and 11 formed particles with the expected size and shape by negative-stain electron microscopy (nsEM). nsEM 3D reconstructions for two designs had an overall shape closely consistent with that of the design models (Fig. 2B; some C7 2D class averages were also obtained; fig. S7). We obtained a cryo-electron microscopy (cryo-EM) map of a third design at 5.1-Å resolution and found it to be closely consistent with the design model: The alpha helices of the model are clearly within the contours of the density (Fig. 2C and fig. S8). The MCTS solution effectively satisfies the design criteria: The space between the two original rings is completely filled in, generating a disk-like structure with a narrow circular pore in the center. We are not aware of any previously designed or naturally occurring proteins that have this overall shape, which could be very useful for downstream nanopore-based sensing applications. More generally, these results demonstrate that the MCTS approach can solve highly constrained protein design problems.

Top-down design of mini-icosahedra

We next explored the use of MCTS to generate icosahedral assemblies by using 59 transformation matrices to compute symmetry mates for a growing monomer. We sought to design very small, closely packed capsids inaccessible by other design methods, and developed geometric constraints and score functions to specifically favor such structures (Fig. 1 and materials and methods). The end state-based score functions include measures of cage porosity and interface designability, as well as external placement of at least one terminus to enable fusion constructs (Fig. 1C). Given a specification of the length and number of helices in the monomer and the size of the overall assembly, we initialized millions of MCTS trajectories starting from a short helical fragment randomly placed

within a specified upper distance bound of the origin in a random orientation and performed 10,000 iterations for each to generate a large set of diverse structures. The MCTS generated closely packed icosahedral assemblies *in silico*, which span a structural space distinct from that of native and previous *de novo* icosahedra, with shorter sequence lengths than any previously described protein icosahedra and porosities comparable to the densely packed capsids generated by evolution (Fig. 1D).

The MCTS method rapidly generates tens of thousands of candidate icosahedral assemblies, and we experimented with approaches for rapidly designing sequences that stabilize these assemblies in a manner compatible with our overall top-down approach. In previous bottom-up nanocage design studies, the sequences and backbones of the oligomeric building blocks are pre-optimized, so only the new interface formed between the building blocks in the cage is designed, and the overall backbone is kept largely fixed (11). By contrast, with the top-down MCTS approach, the entire sequence must be designed, with backbone relaxation to optimize sequence-structure compatibility both within and between the monomers and to increase interface shape complementarity. A deep neural network trained to learn the sequence and structure relationships of native proteins was used to generate amino acid sequence profiles for each position in the newly generated backbones, which were used in turn to bias amino acid selection in the sequence design stage using Rosetta design (materials and methods and figs. S9 to S15). The resulting designs were filtered on the basis of interface contact molecular surface area (38), shape complementarity, predicted binding energy, exposed surface hydrophobicity, and AF (39) prediction similarity to the design model (see the materials and methods). The rigid body and internal degrees of freedom of the selected icosahedral assemblies were then optimized by Rosetta symmetric relaxation (40, 41), starting from both the Rosetta design model of the assembly and the AF-predicted structure of the monomer mapped back onto the assembly. To further increase sequence-structure compatibility, we repeated this design-predict-relax cycle three times, at each iteration performing sequence design on the full assemblies generated in the previous iteration, mapping back the predicted monomer structures into the assemblies, and relaxing the full structure in Rosetta. We applied this sequence design and backbone refinement procedure to 220,000 of the MCTS-generated backbones and selected 368 designs for experimental characterization (detailed filtering processes are described in the materials and methods and figs. S11 to S13).

Linear gene fragments encoding each design with hexahistidine purification tags were cloned




Fig. 2. Disk-nanopore design with symmetric MCTS. (A) Schematic illustration of MCTS-based sampling to build space-filling connectors between two concentric rings to generate disk-like structures with different nanopore inner diameters. The inner ring was placed in the center of a host outer ring, varying the rotation and vertical offset, which generates different void volumes (teal; middle panel above arrows). MCTS was then performed to densely fill these void volumes (blue). (B) Design models (left column) and nsEM 3D ab initio

reconstruction maps (right column) of two connected disk-nanopores (RNR_C6_1 and RNR_C6_2). The symmetric MCTS sampling built helices to connect the inner ring C terminus and outer ring N terminus (highlighted in red and blue, respectively, in the left column). (C) The cryo-EM map at 5.1-Å resolution for design RNR_C6_3 viewed from the top, bottom, and side is very close to the design model, with a narrow circular pore in the center of an otherwise nonporous disk-like structure.

into an *Escherichia coli* expression vector, and the proteins produced in *E. coli* in a 96-well format were purified by immobilized metal affinity chromatography (IMAC) pull-down. A total of 208 of the 368 designs were expressed and soluble as assessed by SDS-polyacrylamide gel electrophoresis. To evaluate particle formation, we performed nsEM on the IMAC elution fraction for each soluble sample. Two designs (RC_I_1 and RC_I_2, RL capsid with I symmetry, design 1 and 2) formed uniform particles with the expected size and shape (Fig. 3, A and B). Size-exclusion chromatography (SEC) of both designs yielded single peaks with an apparent molecular weight in the range expected for these assemblies (Fig. 3, C and D). The designed assemblies had the expected alpha-helical circular dichroism (CD) spectra and apparent melting temperatures

above 65°C. nsEM analysis showed that assembly morphologies were retained after 1 hour of treatment at 95°C and subsequent cooling to 25°C (Fig. 3, F and H, and fig. S17).

To evaluate the accuracy of our design strategy, we determined the structures of SEC-purified RC_I_1 and RC_I_2 capsid particles using cryo-EM (Fig. 4 and fig. S18). For RC_I_1, 3D reconstruction yielded a 2.5-Å-resolution cryo-EM atomic model that closely matched the computational design (Fig. 4, A and B, and fig. S19). The N-terminal helices of two monomers pack in an antiparallel fashion to form the primarily hydrophobic C2 interface, whereas the two helices near the C terminus form the C5 interface with their neighbors (Fig. 4, B and C). Small apertures (diameter ~13 Å) present at the C3 axes of the capsid make the N termini available for genetic fu-

sion (Fig. 4C). Over the designed monomer, the root mean square deviation (RMSD) between the cryo-EM structure and the design model is 0.76 Å (Fig. 4D); a single rotamer flip (Phe⁶³) and tilting of the C-terminal helix results in a slight expansion of the overall cage diameter, resulting in an RMSD over all 60 subunits of 3.72 Å (Fig. 4E). For RC_I_2, the 2.9-Å cryo-EM structure of design RC_I_2 was even closer to the design model (Fig. 4, F and G, and fig. S20), with RMSDs at the C2 and C5 interfaces of 0.66 and 0.27 Å, respectively (Fig. 4H). The RC_I_2 monomer adopts the designed three-helical bundle fold with a 0.59-Å RMSD to the design model (Fig. 4I), and the overall assembly is almost identical to the design model with a 1.39-Å RMSD over all 60 subunits (Fig. 4J). The C2 interface is situated near the extended C terminus of the




Fig. 3. Experimental characterization of designed capsids RC_I_1 and RC_I_2. (A and B) Representative nsEM micrographs and reference-free 2D class averages (inset) for RC_I_1 (left) and RC_I_2 (right). Scale bar, 200 nm. (C and D) A single peak was observed for each SEC elution profile near the expected elution volumes for the target complexes. (E and G) Capsid computational design models. (F and H) Circular dichroism spectra measured at different temperatures ($^{\circ}\text{C}$).

monomer, allowing for potential monomeric or dimeric genetic fusions. The C5 pentameric interface is mediated by interactions between the N-terminal helices, which point inward and enable functionalization of the interior of the capsid. With diameters of 13 and 10 nm for RC_I_1 and RC_I_2, respectively, and associated monomer lengths of 67 and 54 residues, the designed mini-capsids are considerably smaller than most viral capsids.

Applications of top-down-designed capsids

The compact size and corresponding small exterior surface area of the designed parti-

cles enables the display of 60 or 120 copies of N- and/or C-terminal fused proteins with exceptionally high density: six or more times higher than previously designed icosahedral cages. We set out to explore whether this higher density could lead to greater biological efficacy in signaling and vaccine applications. We began by exploring the robustness of the designs to substantial sequence changes and to fusion of proteins to their outward-facing termini.

To evaluate robustness to sequence changes, we used ProteinMPNN (37) to generate diverse sequences for the RC_I_1 capsid backbone, and

the designs were filtered using the AF and Rosetta metrics described above. Two of six experimentally tested ProteinMPNN designs, RC_I_1-H9 and RC_I_1-H11 (the former designed by ProteinMPNN using the working capsid backbone C α coordinates as input, the latter the idealized polyA backbone without any backbone optimization and relaxation), assembled into the designed I1 symmetric capsid as evidenced by IMAC, SEC, and nsEM. A 3-Å cryo-EM structure of RC_I_1-H11 was almost identical to the design model, with a monomeric RMSD of 0.60 Å (Fig. 5A and fig. S21) and a very low full-cage RMSD over all 60 subunits of only 0.96 Å. RC_I_1-H9 and RC_I_1-H11 have on average 46% sequence divergence from the parent capsid and 30% sequence difference from each other, including highly diverse interface residue selections (fig. S22; for example, the errant Phe⁶³ of the parent capsid was redesigned to Glu⁶³ in RC_I_1-H11, likely accounting at least in part for the closer agreement of RC_I_1-H11 with the design model). These results demonstrate that the RL approach can generate directly designable protein backbone geometries with a high degree of accuracy.

We evaluated the robustness of the designs to genetic fusion by fusing SpyTag, SpyCatcher (42), and green fluorescent protein (GFP) proteins to the RC_I_1-H11 capsid with an N-terminal (GGS)_n linker (Fig. 5B and figs. S23 to S25). In all cases, SEC elution profiles and nsEM micrographs showed monodisperse particles of the expected size and shape (see the materials and methods). The 2D class averages (inset) revealed spherical structures similar to that of the original icosahedral capsid, with additional density at the periphery of the particles, consistent with fused proteins connected to scaffolds through a flexible linker. Unlike a larger cage, nuclear localization sequence-tagged capsids fused to GFP are efficiently translocated into the nucleus, opening the door to nuclear delivery of high-valency protein and DNA-organizing constructs (fig. S26).

To assess the efficacy of the designed capsids in activating cellular signaling pathways by clustering cell surface receptors, we fused 60 copies of the angiopoietin 1 (Ang1) F domain (Fd), which binds the Tie2 receptor, to RC_I_1-H11 using SpyTag-SpyCatcher conjugation (14, 18, 43) (see the materials and methods and fig. S27). We found that the F domain-displaying capsids had very high potency in driving FOXO1 exclusion from the nucleus (Fig. 5, C and D), activating the AKT pathway (Fig. 5D and fig. S28, A to C) and stabilizing nascent blood vessels formed from human umbilical vein endothelial cells (HUVECs; Fig. 5E) (43–49). The Fd-displaying capsids (0.16 nM RC_I_1-H11-Fd) elicited stronger responses than a 10-fold greater concentration of a much




Fig. 4. Near-atomic resolution cryo-EM structures of designed capsids match design models. (A) A 2.5-Å cryo-EM reconstruction of RC_I_1 viewed along the three symmetry axes. Scale bar, 20 Å. (B) Cryo-EM structure of RC_I_1 highlighting monomer packing and interfaces along each symmetry axis. (C) Overlay and RMSD calculations for RC_I_1 compared with the design model for each symmetry interface (cryo-EM is shown in blue; design is shown in gray). (D) Overlay and RMSD calculation for a single monomer of RC_I_1. (E) Overlay and

RMSD calculation for the entire 60-mer RC_I_1 capsid. (F) A 2.9-Å cryo-EM reconstruction of RC_I_2 viewed along the three symmetry axes. Scale bar, 20 Å. (G) Cryo-EM structure of RC_I_2 highlighting monomer packing and interfaces along each symmetry axis. (H) Overlay and RMSD calculations for RC_I_2 compared with the design model for each symmetry interface (cryo-EM is shown in pink; design is shown in gray). (I) Overlay and RMSD calculation for a single monomer of RC_I_2. (J) Overlay and RMSD calculation for the entire RC_I_2 capsid.

larger F-domain-presenting icosahedral nanoparticle (I53-50) (12, 43); the elevated potency likely results from the higher surface display density [to facilitate comparison, concentrations at the bottom of Fig. 5D are in terms of Fd monomer (0.16 nM capsid $\times$ 60 Fd copies per capsid = 10 nM Fd)]. The 0.16 nM (10 nM Fd) capsid also elicited stronger responses

than 100 nM Ang1. The F domain-displaying capsid is thus an exceptionally potent Tie2-activating ligand. The designed capsid is also far easier to produce and much more stable than Ang1 and thus could be useful in stimulating differentiation and regeneration.

The high surface presentation density enabled by the designed scaffolds provides a route

to investigating the effect of packing density on the elicitation of immune responses by nanoparticle-based immunogens. As a first step in this direction, we fused trimeric influenza hemagglutinin (HA) to the N terminus of I1-capsid RC_I_1 using a (GS)₆ linker. The fusion protein was expressed and secreted from mammalian cells and clearly forms




Fig. 5. Applications of designed capsids. (A) Robustness of RC_I_1 capsid to sequence redesign using ProteinMPNN. The 3-Å-resolution cryo-EM reconstruction of RC_I_1-H11 (top) reveals a close agreement between the experimental structure (purple) and the design model (gray) with a RMSD of 0.96 Å. The RC_I_1-H11 structure is nearly identical to RC_I_1 despite considerable sequence differences [bottom; residue differences are highlighted in red (RC_I_1-H11) and teal (RC_I_1)]. (B) From top to bottom, models and representative nsEM images of spyTag-, spyCatcher-, and GFP-fused (to N terminus) RC_I_1-H11 with 2D class averages. Scale bar, 50 nm. (C and D) RC_I_1-H11-Fd activates Tie2 downstream Akt phosphorylation and FOXO1 translocation. Serum-starved HUVECs were treated with serially diluted RC_I_1-H11-Fd (1000–0.1 nM), Fd-st (100 nM), Ang1 (100 nM), I53-50 (100 nM), or phosphate-buffered saline (PBS) control for 15 min before protein lysate collection for Western blot analysis, or cells were fixed for FOXO1 antibody stain. (C) Left, representative confocal images of HUVECs immunofluorescence stained with FOXO1 antibody. Right, quantification showing the percentage of cells with nuclear FOXO1; 100 cells were counted in each biological replicate. Levels of significance were compared with PBS control in the FOXO1 graph. (D) Quantification of Western blot showing

pAKT signal normalized to RC_I_1-H11-Fd at 10 nM. RC_I_1-H11-Fd induces a significantly higher signal than the previously characterized I53-50-Fd (inset) at 100 nM Fd equivalent. (E) Quantification of vascular stability by averaging the number of nodes, meshes, and tubes calculated at the 72-hour time point using the Angiogenesis Analyzer plug-in in ImageJ (fig. S28D). In (C) to (E), *P* values were calculated using one-way ANOVA with Bonferroni's multiple-comparisons test in Prism for comparing groups of two or more; **P* < 0.05; ***P* < 0.01; ****P* < 0.001; *****P* < 0.0001; significance over PBS control is noted as # in (D). (F) Representative nsEM micrograph and 2D class averages (inset) of mammalian cell secreted RC_I_1 particle flexibly fused with M15 influenza HA (MI15-RC_I_1). Scale bar, 50 nm. (G) The RC_I_1 displayed HA is antigenically intact, reacting with both head (5J8) and stem (CR9114) anti-HA antibodies in biolayer interferometry experiments. (H) Models of RC_I_1 (top) and I53_dn5 (bottom) displaying MI15 influenza HA (left); the presentation is considerably denser in the former. Top right: Mouse immunization schedule. Bottom right: HA-specific antibody titers in immune sera. Statistical significance was determined using one-way ANOVA with Tukey's multiple-comparisons test; **P* < 0.05; *****P* < 0.0001. The RC_I_1 display format produces a higher antibody titer than the I53_dn5 nanoparticle currently in clinical trials.

HA-displaying particles according to SEC and nsEM (fig. S29 and Fig. 5F). Biolayer interferometry showed binding of both 5J8 [anti-HA head antibody (50)] and CR9114 [anti-HA stem antibody (57)] immunoglobulin G to HA capsids (Fig. 5G), indicating that the

HA remains antigenically intact when displayed on the surface of the capsids. We immunized mice with HA-displaying RC_I_1, as well as a much larger icosahedral immunogen, HA-I53_dn5 (52), which has previously been shown to elicit protective responses against

influenza and is currently being evaluated in clinical trials (53). We found that HA-displaying RC_I_1 elicited a strong antibody response against vaccine-matched HA that was greater than that produced by the clinical vaccine candidate by a small but statistically significant

amount (Fig. 5H). These results indicate that the high antigen presentation density enabled by top-down design can yield robust immune responses.

Conclusion

Our top-down RL approach enables the solution of design challenges inaccessible to previous bottom-up design methods. Cryo-EM structures confirm the design of 54- and 67-residue proteins that assemble into 60-subunit icosahedra with both internal monomer and overall assembly structure nearly identical to the computational models, and of disk-shaped nanopores generated by densely filling the space between cyclic protein rings with different diameters. Both the icosahedra and the disk designs are distinct from any previously designed or naturally occurring structures; the former have smaller subunits, smaller radii, and lower porosities, and the latter have narrow central pores within large, circular, otherwise nonporous structures. These structures could not have been built with previous bottom-up approaches. For the icosahedra, generating the shape complementarity of the interfaces requires the context of the full capsid structure, possible only through a top-down approach, and for the disks, densely filling a prescribed volume from preexisting building blocks is generally not possible. The density of protein chains and termini available for fusion to the icosahedra is considerably greater than the most compact previously designed assembly, enabling fusion to functional protein domains to generate bioactive nanoparticles. The Ang1 F domain-displaying capsids are potent activators of angiogenesis, and the influenza HA-displaying capsids elicit strong anti-HA antibody responses in mice. The capability of the MCTS approach to optimize any set of specified geometric criteria in a top-down fashion provides a route to potent, multivalent cellular receptor agonists and vaccines that are custom designed to rigidly scaffold immunogen or receptor-binding monomers and precisely position them relative to one another. More generally, our results demonstrate the power of RL for protein design, which we expect can be increased further by the incorporation of policy and value networks (30–32, 54) to further guide the search.

REFERENCES AND NOTES

- R. Zandti, D. Reguera, R. F. Bruinsma, W. M. Gelbart, J. Rudnick, *Proc. Natl. Acad. Sci. U.S.A.* **101**, 15556–15560 (2004).
- T. Douglas, M. Young, *Science* **312**, 873–875 (2006).
- J. E. Padilla, C. Colovos, T. O. Yeates, *Proc. Natl. Acad. Sci. U.S.A.* **98**, 2217–2221 (2001).
- J. A. Fallas et al., *Nat. Chem.* **9**, 353–360 (2017).
- J. Zhu et al., *Chem. Rev.* **121**, 13701–13796 (2021).
- B. I. M. Wicky et al., *Science* **378**, 56–61 (2022).
- N. P. King et al., *Science* **336**, 1171–1174 (2012).
- Y.-T. Lai, D. Cascio, T. O. Yeates, *Science* **336**, 1129 (2012).
- Y.-T. Lai et al., *Nat. Chem.* **6**, 1065–1071 (2014).
- N. P. King et al., *Nature* **510**, 103–108 (2014).
- Y. Hsia et al., *Nature* **540**, 150 (2016).
- J. B. Bale et al., *Science* **353**, 389–394 (2016).
- Y. Hsia et al., *Nat. Commun.* **12**, 2294 (2021).
- R. Divine et al., *Science* **372**, eabd9994 (2021).
- C. J. Lanci et al., *Proc. Natl. Acad. Sci. U.S.A.* **109**, 7304–7309 (2012).
- J. C. Sinclair, K. M. Davies, C. Vénien-Bryan, M. E. M. Noble, *Nat. Nanotechnol.* **6**, 558–562 (2011).
- S. Gonen, F. DiMaio, T. Gonen, D. Baker, *Science* **348**, 1365–1368 (2015).
- A. J. Ben-Sasson et al., *Nature* **589**, 468–473 (2021).
- Z. Li et al., *bioRxiv* 2022.11.18.517014 [Preprint] (2022).
- A. C. Walls et al., *Cell* **183**, 1367–1382.e17 (2020).
- P. S. Arunachalam et al., *Nature* **594**, 253–258 (2021).
- P. S. Arunachalam et al., *bioRxiv* [Preprint] (2022).
- J. Y. Song et al., *EClinicalMedicine* **51**, 101569 (2022).
- K. T. Simons, R. Bonneau, I. Ruczinski, D. Baker, *Proteins* **37**, 171–176 (1999).
- D. Gront, D. W. Kulp, R. M. Vernon, C. E. M. Strauss, D. Baker, *PLOS ONE* **6**, e23294 (2011).
- E. Verschuuren et al., *Curr. Opin. Struct. Biol.* **21**, 452–459 (2011).
- N. Koga et al., *Nature* **491**, 222–227 (2012).
- R. Pearce, X. Huang, G. S. Omenn, Y. Zhang, *Proc. Natl. Acad. Sci. U.S.A.* **120**, e2208275120 (2023).
- T. M. Jacobs et al., *Science* **352**, 687–690 (2016).
- B. R. Kiran et al., *IEEE Trans. Intell. Transp. Syst.* **23**, 4909–4926 (2022).
- V. Mnih et al., *Nature* **518**, 529–533 (2015).
- D. Silver et al., *Nature* **529**, 484–489 (2016).
- A. Fawzi et al., *Nature* **610**, 47–53 (2022).
- R. Coulom, in *Computers and Games*, H. J. van den Herik, P. Ciancarini, H. H. L. M. Donkers, Eds. (Springer, 2007), vol. 4630 of *Lecture Notes in Computer Science*, pp. 72–83; http://link.springer.com/10.1007/978-3-540-75538-8_7.
- L. Kocsis, C. Szepesvári, in *Machine Learning: ECML 2006*, J. Fürnkranz, T. Scheffer, M. Spiliopoulou, Eds. (Springer, 2006), vol. 4212 of *Lecture Notes in Computer Science*, pp. 282–293; http://link.springer.com/10.1007/11871842_29.
- J. P. Hallinan et al., *Commun. Biol.* **4**, 1240 (2021).
- J. Dauparas et al., *Science* **378**, 49–56 (2022).
- L. Cao et al., *Nature* **605**, 551–560 (2022).
- J. Jumper et al., *Nature* **596**, 583–589 (2021).
- L. G. Nivón, R. Moretti, D. Baker, *PLOS ONE* **8**, e59004 (2013).
- P. Conway, M. D. Tyka, F. DiMaio, D. E. Konerding, D. Baker, *Protein Sci.* **23**, 47–55 (2014).
- B. Zakeri et al., *Proc. Natl. Acad. Sci. U.S.A.* **109**, E690–E697 (2012).
- Y. T. Zhao et al., *EMBO Rep.* **22**, e53471 (2021).
- S. Han et al., *Sci. Transl. Med.* **8**, 335ra55 (2016).
- I. Kim et al., *Circ. Res.* **86**, 24–29 (2000).
- M. Kim et al., *J. Clin. Invest.* **126**, 3511–3525 (2016).
- C. Daly et al., *Genes Dev.* **18**, 1060–1071 (2004).
- K. L. DeCicco-Skinner et al., *J. Vis. Exp.* **91**, e51312 (2014).
- N. P. J. Brindle, P. Saharinen, K. Alitalo, *Circ. Res.* **98**, 1014–1023 (2006).
- J. C. Krause et al., *J. Virol.* **85**, 10905–10908 (2011).
- C. Dreyfus et al., *Science* **337**, 1343–1348 (2012).
- G. Ueda et al., *eLife* **9**, e57659 (2020).
- S. Boyoglu-Barnum et al., *Nature* **592**, 623–628 (2021).
- D. Silver et al., *Nature* **550**, 354–359 (2017).
- I. Lutz, Protein backbone MCTS, Zenodo, (2023); <https://doi.org/10.5281/zenodo.7709840>.
- S. Wang, C. Norn, I. Lutz, 2023 RL capsid design, Zenodo (2023); <https://doi.org/10.5281/zenodo.7758067>.
- computational design and discussion; R. Kibler, D. Feldman, L. Milles, and N. Ennist for help with the experiments; and S. Dickinson and J. Quispel for help in maintaining and operating the electron microscopes used. **Funding:** This work was conducted at the Advanced Light Source (ALS), a national user facility operated by Lawrence Berkeley National Laboratory on behalf of the Department of Energy, Office of Basic Energy Sciences, through the Integrated Diffraction Analysis Technologies (IDAT) program, supported by DOE Office of Biological and Environmental Research. Additional support comes from the National Institutes of Health project ALS-ENABLE (NIH grant P30 GM124169) and a High-End Instrumentation Grant S1000D18483. This work was also supported by the National Institute on Aging grant (grant 1U19AG065156-01 to I.D.L. and D.B.); Amgen (S.W.), the Audacious Project at the Institute for Protein Design (A.J.B., Z.L., L.C., H.R.-B., Y.T.Z., D.B., and N.P.K.); the NIH/National Institute of Dental and Craniofacial Research (NIDCR) (grant T90 DE021984 to Y.T.Z.); the National Institute of Allergy and Infectious Diseases (NIAID grant 1P01AI167966 to N.P.K.); Novo Nordisk (foundation grant NNF170C0030446 to C.N.); the Open Philanthropy Project Universal Flu Vaccine and Improving Rosetta Design (A.D., N.P.K., and D.B.); a Microsoft gift (M.B.); and the Department of Defense Peer Reviewed Medical Research Program (award W81XWH-21-1-0006 to H.R.-B. and D.B.). **Author contributions:** I.D.L. and S.W. contributed equally and the author order was chosen arbitrarily; citations on CVs, etc., will be adjusted accordingly. S.W., I.D.L., C.N., and D.B. conceptualized the research. I.D.L. developed the RL backbone generation method. S.W. and C.N. developed the sequence design pipeline. L.C. and M.B. developed the DL sequence profile prediction method. S.W. designed the original capsids and performed the screening, expression, and characterization experiments. A.C. and J.X. designed the nanopore ring connector and performed cryo-EM characterization. A.J.B. designed cryo-EM experiments and optimization of sample purification conditions, performed the initial cryo-EM screening experiments, and optimized the cryo-EM freezing conditions. A.J.B., S.W., and Z.L. prepared additional cryo-EM grids and collected cryo-EM data. A.J.B. processed the cryo-EM data and built and solved the structures for each designed capsid. S.W. and A.D. designed and characterized the fusion capsids. A.D., C.T., E.M.L., and N.P.K. designed and performed the immunization studies for the HA-capsid fusions. A.D.G. produced and characterized HA-capsid proteins. Y.T.Z., P.L., and H.R.-B. designed and performed the cell signaling assays. C.N., P.R.S., and A.M.B. designed and performed the nuclear localization experiments. All authors analyzed the data. D.B. supervised the research. S.W., I.D.L., and D.B. wrote the manuscript with the input from the other authors. All authors revised the manuscript. **Competing interests:** D.B., S.W., I.D.L., C.N., A.D., N.P.K., and A.J.B. are inventors on a provisional patent application (63/383,700) submitted by the University of Washington for the design, composition, and applications of the protein assemblies described in this work. The remaining authors declare no competing interests. **Data and materials availability:** All data are available in the main text or the supplementary materials. MCTS backbone sampling code and structural coordinates of the in silico-generated examples are available at <https://github.com/idutz/protein-backbone-MCTS> and archived at Zenodo (55). Code for the sequence design and filtering of assemblies is available at https://files.ipd.uw.edu/pub/2023_RL_capsid_design/sequence_design_pipeline.tar and archived at Zenodo (56). For RC_1_1, RC_1_2, and RC_1_H11 capsid structures, coordinates are deposited in the Protein Data Bank with the accession codes 8F54, 8F53, and 8F4X; cryo-EM density maps are deposited in the Electron Microscopy Data Bank (EMDB) with the accession codes EMD-28860, EMD-28859, and EMD-28858. **License information:** Copyright © 2023 the authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original US government works. <https://www.science.org/about/science-licenses-journal-article-reuse>

SUPPLEMENTARY MATERIALS

science.org/doi/10.1126/science.adf6591
Materials and Methods
Figs. S1 to S29
Tables S1 to S5
References (57–78)
MDAR Reproducibility Checklist

[View/request a protocol for this paper from Bio-protocol.](#)

Submitted 5 November 2022; accepted 21 March 2023
10.1126/science.adf6591



Top-down design of protein architectures with reinforcement learning

Isaac D. Lutz, Shunzhi Wang, Christoffer Norn, Alexis Courbet, Andrew J. Borst, Yan Ting Zhao, Annie Dosey, Longxing Cao, Jinwei Xu, Elizabeth M. Leaf, Catherine Treichel, Patrisia Litvicov, Zhe Li, Alexander D. Goodson, Paula Rivera-Snchez, Ana-Maria Bratovianu, Minkyung Baek, Neil P. King, Hannele Ruohola-Baker, and David Baker

Science, **380** (6642), .

DOI: 10.1126/science.adf6591

Complex architectures by top-down design

Designing artificial protein complexes based on components with established structural properties can impose limits on the properties of the final complex. Lutz *et al.* developed methods for generating complex protein architectures that adhere to preordained parameters using reinforcement learning. They demonstrated this approach by generating designs that fill arbitrary volumes, including a symmetrical connector between previously designed protein rings. A small protein designed to assemble into 60-subunit icosahedra may be useful for presenting antigens in vaccines or signaling molecules in multivalent agonist complexes, as the authors demonstrate in preliminary biological experiments. —MAF

View the article online

<https://www.science.org/doi/10.1126/science.adf6591>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

Science (ISSN) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science* is a registered trademark of AAAS.

Copyright © 2023 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works