MEVOCRAFT: MOTIF-SCOPED EVOLUTIONARY CONDITIONING FOR MULTISTATE PROTEIN DESIGN

Anonymous authors
Paper under double-blind review

ABSTRACT

Many proteins function by switching among multiple distinct conformational states, and designing proteins with such behavior could enable new applications in synthetic biology and therapeutics. A recent method, SwitchCraft, introduces a unified framework for multistate protein design that optimizes a sequence by gradient-based hallucination through a biomolecular structure predictor. Though promising, SwitchCraft performs every forward pass *without a multiple sequence alignment (MSA)*, omitting the evolutionary information that is typically useful for structure prediction. As a result, the gradient used to update the sequence is computed from the inaccurate structures. Incorporating an MSA into this hallucination process is challenging because the query sequence changes at every optimization step, meaningful homologs may not exist, and the sequence is often represented as a continuous distribution rather than a discrete sequence. To address these challenges, in this paper we propose a novel method, called motif-scoped evolutionary conditioning for multistate protein design (MevoCraft). MevoCraft conditions each forward pass of gradient-based hallucination with a motif MSA which is constructed by a single homology search on the fixed functional motif and injected only into states required to exhibit that motif. To the best of our knowledge, MevoCraft is the first method to incorporate an MSA into gradient-based hallucination for multistate protein design. Experiments show that MevoCraft outperforms SwitchCraft in 50 of 70 positive-allostery cases and 49 of 70 negative-allostery cases, while matching it in another 19 positive-allostery cases and 14 negative-allostery cases. In motif switching, MevoCraft produces $3.8\times$ the number of successful designs achieved by SwitchCraft.

1 INTRODUCTION

Many natural proteins switch among distinct conformational states in response to ligands or environmental signals, as exemplified by oxygen-bound hemoglobin and calcium-bound calmodulin (Perutz, 1970; Chin & Means, 2000; Hoeflich & Ikura, 2002). Designing proteins with such behavior, typically called multistate protein design, could enable biosensors, conditionally active therapeutics, and programmable components for synthetic biology. Multistate protein design is much more difficult than designing a single stable protein because one sequence must encode multiple structures and functions (Davey & Chica, 2012; Gonzalez et al., 2022). Early methods optimize sequences with state-specific energy functions (Allen & Mayo, 2010; Hallen & Donald, 2016; Löffler et al., 2017; Sauer et al., 2020). Though interpretable, they require predefined structural states and handcrafted energy functions, making the resulting designs sensitive to inaccuracy in those functions. Generative methods (Dauparas et al., 2022; Yi et al., 2025; Lisanza et al., 2025; Jing et al., 2025; Abrudan et al., 2026) are constrained by the scarcity of training data that capture conformational transitions. They either omit explicit ligand context or rely on post-hoc aggregation of state-specific outputs to obtain multistate prediction.

More recently, SwitchCraft (Jing et al., 2026) introduces a unified framework for multistate protein design that optimizes a sequence by gradient-based hallucination through biomolecular structure predictors such as AlphaFold3 (Abramson et al., 2024) or Boltz (Wohlwend et al., 2024). In SwitchCraft, each state is defined by its conditioning ligand context, and the design sequence is optimized by backpropagating state-specific loss through the structure predictor. This formulation circumvents the scarcity of multistate training data. Though promising, SwitchCraft performs every

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

forward pass *without a multiple sequence alignment (MSA)*, omitting the evolutionary information that is typically useful for structure prediction. In this single-sequence setting, AlphaFold-family structure predictors are substantially less accurate than when conditioned on an MSA (Del Alamo et al., 2022; Lin et al., 2023; Wayment-Steele et al., 2024). As a result, the gradient used to update the sequence are computed from the inaccurate structures.

But incorporating an MSA into the hallucination process of SwitchCraft is challenging. First, the query sequence changes at every optimization step, so a new homology search would be required at each step. Conventional sequence search methods, such as MMseqs (Steinegger & Söding, 2017) or HHblits (Remmert et al., 2012), are too slow for this setting: a single design task would require hundreds of database searches, making the computational cost prohibitive. Second, the sequences to design may lie far from natural sequence space, particularly during early optimization steps or when constructing *de novo* scaffolds. In such cases, meaningful homologs may not exist in current sequence databases. Faster embedding-based search methods, such as PLMSearch (Liu et al., 2024), TM-Vec (Hamamsy et al., 2024), and ERAST (Jiang et al., 2026), reduce only the cost of search and they still depend on homologs already existed in the database. Third, the sequence is often represented as a continuous distribution over amino-acid identities rather than as a discrete sequence. Existing MSA retrieval and generation methods, such as MSAGPT (Chen et al., 2024), require a discrete query sequence and therefore cannot be applied directly. These challenges make it difficult to incorporate an MSA into gradient-based hallucination for multistate protein design (Pacesa et al., 2025; Jing et al., 2026).

To address these challenges, in this paper we propose a novel method, called motif-scoped evolutionary conditioning for multistate protein design (MevoCraft). MevoCraft conditions each forward pass of gradient-based hallucination with a motif MSA which is constructed by a single homology search on the fixed functional motif and injected only into states required to exhibit that motif. The main contributions of this paper are outlined as follows:

- We propose a novel MSA construction method via homology search on the fixed functional motif, avoiding multiple times of search for variable and non-discrete design sequences at every optimization step. Moreover, functional motifs are often evolutionarily conserved, and hence this search strategy can identify motif homologs even when the complete design sequence has no meaningful homologs.
- We propose a novel state-conditional MSA injection method, which restricts the MSA to states required to exhibit the corresponding motif, while allowing other states to explore alternative conformations without constraint from that MSA.
- To the best of our knowledge, MevoCraft is the first method to incorporate MSA into gradient-based hallucination for multistate protein design.
- Experiments show that MevoCraft outperforms SwitchCraft in 50 of 70 positive-allostery cases and 49 of 70 negative-allostery cases, while matching it in another 19 positive-allostery cases and 14 negative-allostery cases. In motif switching, MevoCraft produces $3.8\times$ the number of successful designs achieved by SwitchCraft.

2 Related Works

Multistate protein design aims to identify a single protein sequence that simultaneously satisfies constraints across multiple structural or functional states. Existing methods can be broadly divided into three categories, including energy-based methods, generative methods, and gradient-based hallucination methods.

Energy-based methods. Early methods, including FASTER (Allen & Mayo, 2010), COMETS (Hallen & Donald, 2016), and Rosetta:MSF (Löffler et al., 2017), jointly optimize state-specific energy functions over multiple predefined structural or chemical states using heuristic or combinatorial sequence search. RECON (Sauer et al., 2020) further extends energy-based design to structural ensembles by progressively enforcing sequence convergence across different states. These methods rely on manually specified energy functions, the limited accuracy of which in modeling protein energetics typically constrain the quality and reliability of the resulting designs.

Generative methods. Generative methods formulate multistate protein design as conditional generation. Tied ProteinMPNN (Dauparas et al., 2022) and ADFLIP (Yi et al., 2025) extend the model to multistate protein design by averaging or linearly combining the sequence logits predicted for different structural states during decoding, but they are not explicitly trained on multistate data. ProteinGenerator (Lisanza et al., 2025) performs sequence-space diffusion conditioned on multiple structural states and combines state-specific sequence predictions to generate sequences compatible with different conformations. ProDiT (Jing et al., 2025) further incorporates sequence, structure, and functional information into a multimodal diffusion model for generating proteins with multiple structural or functional states. DynamicMPNN (Abrudan et al., 2026) explicitly trains an inverse-folding model on conformational ensembles to jointly learn sequence constraints shared across multiple states. But they do not explicitly consider the ligand information.

Gradient-based hallucination methods. The most related work is SwitchCraft (Jing et al., 2026), which introduces a unified framework for multistate protein design that optimizes a sequence by gradient-based hallucination through a biomolecular structure predictor. Each state is defined by its conditioning ligand context, and the design sequence is optimized by backpropagating state-specific losses through the structure predictor. Unlike energy-based methods, it replaces handcrafted energy functions with learned structural objectives; unlike generative methods, it directly optimizes each sequence rather than learning an amortized distribution over multistate proteins. Though promising, SwitchCraft performs every forward pass *without a MSA*. In this single-sequence setting, the predictor is less accurate than when conditioned on an MSA. As a result, the gradient used to update the sequence are computed from the inaccurate structures.

3 METHODS

Following SwitchCraft, we formulate multistate protein design as a gradient-based protein hallucination problem. As illustrated in Figure 1, MevoCraft consists of three components: design specification (Section 3.1), MSA scheduling (Section 3.2), and design optimization (Section 3.3). Design specification defines the fixed functional motif, the ligand-conditioned target states, and a set of loss functions that jointly specify the desired protein behavior. MSA scheduling constructs an MSA by a single homology search on the fixed functional motif and injected only into states required to exhibit that motif. Design optimization then backpropagates the multistate loss through the structure predictor to update the shared sequence representation across all target states. These three components enable MevoCraft to leverage evolutionary information to improve structure prediction while preserving its ability to design multistate proteins.

3.1 DESIGN SPECIFICATION

A multistate protein design problem specifies a set of target states $\mathcal{S} = \{1, \ldots, N_{\text{states}}\}$, where N_{states} is the number of desired states. Each state $s \in \mathcal{S}$ is associated with a folding context $\mathcal{C}_s$, which contains the fixed molecules present in that state, such as a conditioning ligand, and may be empty. All states share a trainable sequence representation $\boldsymbol{z}(t) \in \mathbb{R}^{20\times L}$ at optimization step t, where L is the length of the protein to be designed.

A binary mask $\mathbf{m} \in \{0,1\}^L$ distinguishes the positions to be optimized, where $\mathbf{m}_i = 1$ indicates that position i is optimized and $\mathbf{m}_i = 0$ indicates that position i is fixed. Let $\mathcal{A}$ denote the set of 20 canonical amino acids. An optional partial sequence $\mathbf{s}$ specifies $\mathbf{s}_i \in \mathcal{A}$ at some fixed motif positions. We denote these positions by $\mathcal{I}_{\text{m}}$ and the number of these positions by $L_m = |\mathcal{I}_{\text{m}}|$. This shared parameterization ensures that the optimization process produces only one single sequence for all target states rather than a separate sequence for each target state.

Let f_θ denote the pretrained Boltz structure predictor (Wohlwend et al., 2024) with fixed parameters θ. At step t, its prediction for state s is

$$\widehat{\mathcal{Y}}_s(t) = f_\theta(\boldsymbol{z}(t), \mathcal{C}_s, M_s(t)), \tag{1}$$

where $M_s(t)$ is the step-specific MSA input for each state s and $\widehat{\mathcal{Y}}_s(t)$ denotes the predicted structure. The multistate objective combines N_{loss} loss terms:

$$\mathcal{L}(\boldsymbol{z}(t)) = \sum_{n=1}^{N_{\text{loss}}} \mathcal{L}_n\Big(\{\widehat{\mathcal{Y}}_s(t)\}_{s\in\mathcal{S}_n}\Big), \tag{2}$$

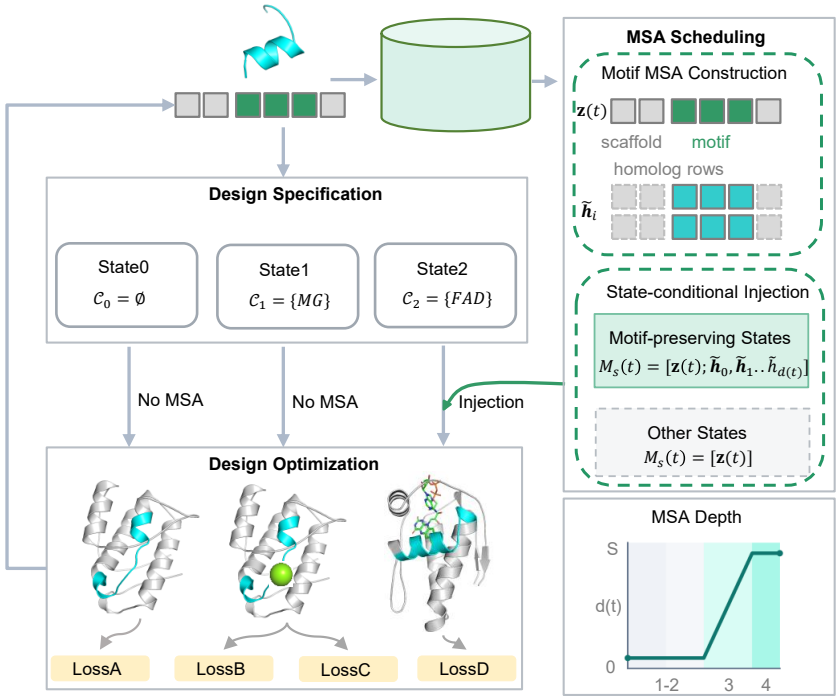


Figure 1: Overview of MevoCraft. A shared sequence $\boldsymbol{z}(t)$ is optimized across folding contexts $\mathcal{C}_s$. A motif MSA is injected only into motif-preserving states, while other states receive query-only inputs. Losses A–D are a selection from motif/anti-motif, binding/anti-binding, conformational-change, and contact losses. MSA depth $d(t)$ increases in later stages, and the multistate loss updates $\boldsymbol{z}(t)$.

where $\mathcal{L}_n$ is the n-th loss term, and $\mathcal{S}_n \subseteq \mathcal{S}$ is the subset of states on which it is evaluated. Gradients of Equation (2) are backpropagated through the fixed structure predictor to update the shared sequence representation $\boldsymbol{z}(t)$.

As in SwitchCraft, the objective contains four types of losses: (1) motif and anti-motif losses, which respectively encourage and discourage formation of the specified motif; (2) binding and anti-binding losses, which respectively promote and suppress ligand binding; (3) a conformational-change loss, which promotes structural differences between target states; and (4) a contact loss, which encourages well-formed intraprotein contact. Detailed definitions are provided in Jing et al. (2026).

3.2 MSA SCHEDULING

Given the motif specification $(\mathbf{m}, \mathbf{s})$, MSA scheduling constructs an MSA $M_s(t)$ for each target state s, which is supplied to the structure predictor at optimization step t as in Equation (1). The MSA scheduling consists of motif MSA construction and state-conditional injection. We construct an MSA by a single homology search on the fixed functional motif and injected only into states required to exhibit that motif. How MSA depth changes across optimization steps is specified in Section 3.3.

3.2.1 MOTIF MSA CONSTRUCTION

Constructing an MSA for the variable design sequence is impractical because the query sequence changes at every optimization step and is often represented as a continuous distribution over amino-acid identities. We instead construct an MSA via homology search on the fixed functional motif $(\mathbf{m}, \mathbf{s})$ before sequence optimization begins. This design addresses the challenges described in Section 1. First, because the motif is fixed, so it avoids multiple times of search for variable and non-discrete design sequences at every optimization steps. Second, functional motifs are often evolutionarily

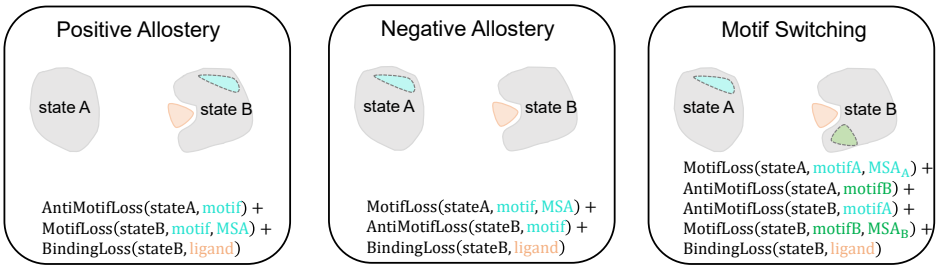


Figure 2: State-conditional MSA injection for representative multistate design objectives.

conserved, and hence this search strategy can identify motif homologs even when the complete design sequence has no meaningful homologs.

Formally, this homology search yields a set of fixed-length homologs:

$$\mathcal{H} = \{\mathbf{h}_1, \ldots, \mathbf{h}_K\}, \qquad \mathbf{h}_i \in \mathcal{A}_{\text{gap}}^{L_m}, \tag{3}$$

where $K = |\mathcal{H}|$ is the number of retained homologs and $\mathcal{A}_{\text{gap}} = \mathcal{A} \cup \{\texttt{gap}\}$ is the amino-acid alphabet augmented with a gap token. Depending on the available motif data, we use two complementary retrieval strategies: (i) a structure-based Foldseek search (Van Kempen et al., 2024), which queries the motif structure and projects each hit onto a common set of motif columns; and (ii) a sequence-based MMseqs search (Steinegger & Söding, 2017), typically using the full-length native sequence as the query, after which aligned residues corresponding to the motif positions are extracted. Hits are filtered using E-value and sequence-quality criteria, including the removal of sequences containing unresolved "X" residues, and are deduplicated before injection. Further implementation details are provided in Section A.1.

At injection time, each homolog is embedded into the full design length L by placing amino acids only at motif positions and filling all scaffold positions with `gap` tokens:

$$\tilde{\mathbf{h}}_i[j] = \begin{cases} \mathbf{h}_i[\pi(j)], & j \in \mathcal{I}_{\text{m}}, \\ \texttt{gap}, & \text{otherwise}, \end{cases} \tag{4}$$

where $\mathcal{I}_{\text{m}}$ indexes motif positions in the sequence and π maps those positions to MSA columns.

3.2.2 State-Conditional Injection

As illustrated in Figure 2, multistate design tasks such as positive allostery, negative allostery, and motif switching require a functional motif to be exhibited in some target states and disrupted in others. For example, in positive allostery, the motif should be formed in the ligand-bound holo state and disrupted in the ligand-free apo state. We inject the motif MSA only into states required to exhibit the motif, which we call *motif-preserving* states. Formally, the set of motif-preserving states is defined as

$$\mathcal{S}_{\text{motif}} = \big\{ s \in \mathcal{S} : \text{ state } s \text{ is specified to exhibit the motif} \big\}. \tag{5}$$

Prior work shows that MSA inputs can improve prediction confidence and bias structure predictors toward particular conformations (Wayment-Steele et al., 2024). We therefore inject the motif MSA only into motif-preserving states, where it can improve prediction confidence and favor the target motif conformation. We withhold the MSA from other states to avoid imposing the same evolutionary constraints where the motif should be disrupted. We evaluate state-conditional injection in Section 4.4.

For notational simplicity, the equations below use one motif and K homologs. For motif switching, each motif has its own homolog set and its own $\mathcal{S}_{\text{motif}}$. Each state receives the MSA of the motif it should exhibit.

The shared sequence representation $\boldsymbol{z}(t)$ serves directly as the query row of the state-conditional MSA. For a scheduled depth $d(t) \in \{0, 1, \ldots, K\}$, we set

$$M_s(t) = \begin{cases} \big[\boldsymbol{z}(t);\ \tilde{\mathbf{h}}_{i_1}, \ldots, \tilde{\mathbf{h}}_{i_{d(t)}}\big], & s \in \mathcal{S}_{\text{motif}} \text{ and } d(t) > 0, \\ \big[\boldsymbol{z}(t)\big], & \text{otherwise}. \end{cases} \tag{6}$$

where $\{i_1, \ldots, i_{d(t)}\}$ is a uniformly sampled subset of $\{1, \ldots, K\}$. The homologs $\{\tilde{\mathbf{h}}_i\}$ are treated as non-differentiable constants, and only $\boldsymbol{z}(t)$ is updated.

3.3 Design Optimization

Following BoltzDesign1 (Cho et al., 2025), we optimize the shared sequence representation $\boldsymbol{z}(t)$ using a 240-step, four-stage gradient schedule that gradually transforms continuous sequence logits into a discrete amino-acid sequence. At each iteration, the structure predictor evaluates all target states under their respective folding context and scheduled MSA input. The resulting multistate loss in Equation (2) updates the shared sequence representation according to

$$\boldsymbol{z}(t+1) = \boldsymbol{z}(t) - \eta_t \nabla_{\boldsymbol{z}(t)} \mathcal{L}(\boldsymbol{z}(t)), \tag{7}$$

where η_t is the step-dependent learning rate. Motif positions specified by $(\mathbf{m}, \mathbf{s})$ remain fixed, while all designable positions are jointly updated using gradients accumulated across the target states.

We couple the number of injected homologs to the progression of sequence optimization. We leave early iterations unconditioned to preserve sequence exploration and introduce stronger evolutionary conditioning in later stages. Let $\rho(t) \in [0, 1]$ denote normalized progress within Stage 3. We schedule the number of injected homologs as

$$d(t) = \begin{cases} 0, & \text{Stages 1–2,} \\ \mathrm{round}(\rho(t)K), & \text{Stage 3,} \\ K, & \text{Stage 4.} \end{cases} \tag{8}$$

Stages 1–2 optimize all states without MSA conditioning. Stage 3 progressively increases the MSA depth for motif-preserving states, and Stage 4 supplies the full MSA. After the final stage, the argmax amino-acid identity at each designable position, together with the fixed motif residues, defines the output sequence. Details of the optimization protocol and stage-specific hyperparameters are provided in Section A.2, and the complete procedure is summarized in Algorithm 1.

4 Experiments

We follow the experimental protocol of SwitchCraft (Jing et al., 2026), which considers six multistate design settings: positive allostery, negative allostery, motif switching, ligand modification, induced binding, and ligand discrimination. We focus on positive allostery, negative allostery and motif switching because these settings explicitly specify a functional motif and therefore support motif-scoped evolutionary conditioning. The other three settings do not specify a motif region. For these settings, we simply set $M_s(t) = [\boldsymbol{z}(t)]$, and the design procedure degenerates to the original single-sequence formulation.

To ensure a fair comparison, we reproduce SwitchCraft using the same computational devices as MevoCraft. The reproduced results are consistent with those reported in the original paper. All experiments are conducted on NVIDIA RTX A6000 GPUs. For each design setting, we generate 100 independent designs using the same hyperparameters and without task-specific tuning. The formal definition of the success criteria for each design task is provided in Section A.3.

4.1 Positive and Negative Allostery

In positive allostery, the functional motif is disrupted in the ligand-free apo state and formed in the ligand-bound holo state. The inverse behavior holds for negative allostery. The original benchmark contains 24 functional motifs from RFDiffusion (Watson et al., 2023). We exclude 1ycr and 1qjg because their motifs contain fewer than 10 residues, which we consider to be insufficient for constructing a meaningful motif MSA. We evaluate each motif with five ligands: the small molecules OQO and flavin adenine dinucleotide (FAD); the metal ions Zn^{2+} and Mg^{2+}; and double-stranded DNA (dsDNA) with sequence `GAATTC`. These ligands are denoted OQO, FAD, ZN, MG, and DNA respectively.

Among the remaining 22 motifs, both MevoCraft and SwitchCraft produce zero successful designs on eight motifs. We omit these cases from the visualization and show the comparison of the remaining

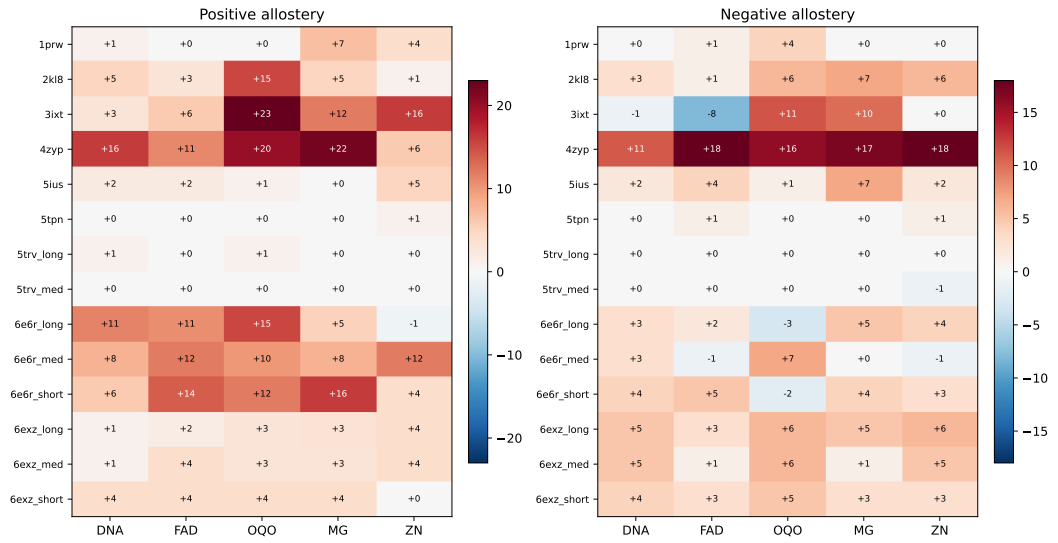


Figure 3: Difference in successful design counts for positive allostery (left) and negative allostery (right). Rows denote functional motifs, and columns denote conditioning ligands. Each cell reports the number of successful designs produced by MevoCraft minus that produced by SwitchCraft. Positive and negative values indicate more and fewer successful designs with MevoCraft, respectively, while zero indicates equal counts.

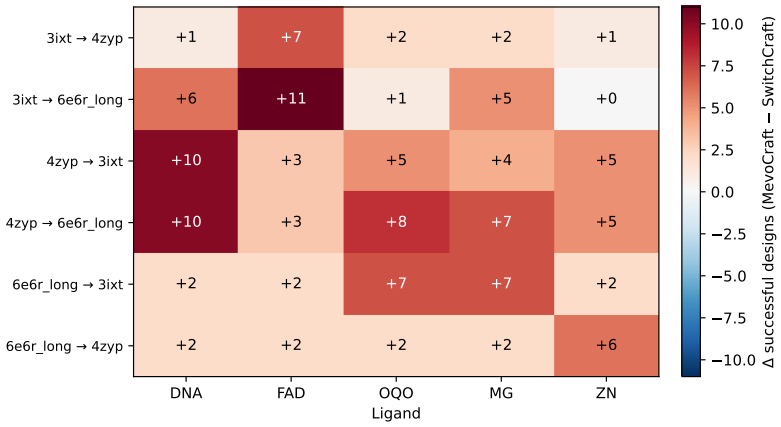


Figure 4: Difference in successful design counts for motif switching. Rows denote ordered motif pairs, with each arrow pointing from the motif required in the apo state to the motif required in the holo state. Columns denote conditioning ligands.

14 motifs in Figure 3. For positive allostery, MevoCraft produces more successful designs than SwitchCraft in 53 of the 70 displayed motif–ligand combinations, the same number of successful designs in 16 combinations, and fewer successful designs in one combination. For negative allostery, MevoCraft produces more successful designs than SwitchCraft in 46 of the 70 combinations, the same number of successful designs in 17 combinations, and fewer successful designs in seven combinations. Detailed results for all motif–ligand combinations are provided in Figures S2 and S3. Notably, SwitchCraft produces few or no successful designs for `4zyp` and `5ius` under positive allostery, and for `4zyp` and the three `6exz` scaffold-length variants under negative allostery. MevoCraft produces successful designs for these targets across multiple conditioning ligands. These cases indicate that the gains extend beyond targets on which SwitchCraft already performs well to targets on which SwitchCraft rarely succeeds. Overall, the results show that motif-scoped evolutionary conditioning can improve design success in both positive- and negative-allostery tasks.

Table 1: Comparison of success rate, pLDDT, motif RMSD, and RMSD gap between MevoCraft and SwitchCraft.

	Positive allostery		Negative allostery	
Metric	MevoCraft	SwitchCraft	MevoCraft	SwitchCraft
Success rate (%) (↑)	**15.9**	6.0	**7.3**	1.9
Unbound pLDDT (↑)	**0.711**	0.702	**0.749**	0.743
Bound pLDDT (↑)	**0.735**	0.705	**0.758**	0.745
Motif RMSD (Å) (↓)	**1.33**	2.28	**1.06**	1.65
RMSD gap (Å) ($\|gap\|$ ↑)	**1.66**	0.99	**−0.36**	−0.01

4.2 Motif Switching

Motif switching requires a single protein sequence to realize two distinct functional motifs in different ligand-dependent states. For an ordered pair $A \rightarrow B$, motif A must be structurally realized in the ligand-free apo state and disrupted in the ligand-bound holo state, while motif B must exhibit the inverse behavior.

We select `3ixt`, `4zyp`, and `6e6r_long`, the three motifs with the highest combined success rates across positive and negative allostery under MevoCraft. This selection follows the strategy of SwitchCraft (Jing et al., 2026), which paired the motifs with the highest success rates in positive and negative allostery in a single motif-switching experiment. Here, we evaluate all six ordered pairs of distinct motifs from the selected set, testing both switching directions for each pair. Each pair is evaluated with DNA, FAD, OQO, MG, and ZN, yielding 30 design settings with 100 designs per method in each setting.

Figure 4 shows the difference in successful design counts between MevoCraft and SwitchCraft, with the counts for each method provided in Figure S4 of Section B.3. MevoCraft produces more successful designs than SwitchCraft in 29 of the 30 cases and the same number in the remaining case. Notably, MevoCraft produces successful designs in all 30 settings, including the 11 settings in which SwitchCraft produces none.In total, MevoCraft produces $3.8\times$ the number of successful designs achieved by SwitchCraft. Overall, MevoCraft improves motif-switching success rates and succeeds across more design settings.

4.3 Prediction Confidence and Conformational Accuracy

To investigate how MSA conditioning contributes to design success, we analyze the five targets with the largest gains in success rates, including `6e6r_long`, `6e6r_med`, `6e6r_short`, `2kl8`, and `4zyp`. Table 1 shows the success rates, pLDDT, motif RMSD, and the unbound–bound motif RMSD gap, averaged over the five selected targets and conditioning ligands. These results reveal three patterns. First, MevoCraft achieves higher pLDDT in both apo and holo states under both allostery tasks, indicating increased prediction confidence. In the MSA-conditioned motif-preserving state, pLDDT increases from 0.705 to 0.735 for positive allostery and from 0.743 to 0.749 for negative allostery. These results support the role of MSA conditioning in improving prediction confidence. Second, motif RMSD in the motif-preserving state decreases from 2.28 Å to 1.33 Å for positive allostery and from 1.65 Å to 1.06 Å for negative allostery. The closer agreement with the reference motif geometry indicates that MSA conditioning helps the structure predictor more accurately realize the target motif conformation. Third, the RMSD gap shows greater separation between states in the direction required by each task. Defined as apo motif RMSD minus holo motif RMSD, the gap increases from 0.99 Å to 1.66 Å for positive allostery and becomes more negative, from -0.01 Å to -0.36 Å, for negative allostery. These observations suggest that higher prediction confidence and more accurate motif geometry help sequence optimization better distinguish the conformations required in different states.

4.4 Ablation Study

We conduct an ablation study on positive allostery with the `1prw` motif to examine the effect of state-conditional MSA injection and injection stage. Table 2 compares four variants with the default MevoCraft configuration.

Table 2: Ablation study of state-conditional MSA injection and the stage of MSA injection.

Variant	Unbound pLDDT (↑)	Bound pLDDT (↑)	Unbound RMSD (Å) (↑)	Bound RMSD (Å) (↓)	RMSD gap (Å) (↑)
Full MSA: Evaluation only	0.6058	0.6728	5.3959	2.5676	2.8283
Full MSA: Stage 4 only	0.6076	0.6716	5.3923	2.3663	3.0260
Full MSA: Stages 3–4	0.6031	0.6702	5.8344	2.5422	3.2922
MSA in all states	**0.6173**	0.6805	3.3389	2.3301	1.0088
MevoCraft	0.6120	**0.6911**	**5.8949**	**2.2818**	**3.6131**

First, *Full MSA: Evaluation only* tests whether MSA conditioning improves sequence optimization or merely makes prediction more readily satisfy the evaluation criteria. This variant omits MSA throughout optimization but uses the same full-depth, state-conditional MSA inputs as MevoCraft at final evaluation. Under these matched evaluation conditions, MevoCraft achieves higher pLDDT in both states, lower bound motif RMSD, and a larger RMSD gap. These results suggest that MSA conditioning improves the designed sequences through optimization rather than merely making the evaluation criteria easier to satisfy.

Second, *Full MSA: Stage 4 only* introduces the full MSA only in the final optimization stage. Compared with MevoCraft, it yields lower pLDDT in both states, higher bound motif RMSD, and a smaller RMSD gap, supporting the benefit of earlier MSA conditioning.

Third, *Full MSA: Stages 3–4* replaces progressive injection with an abrupt transition from no MSA to the full MSA at the start of Stage 3. This comparison examines whether a sudden change in conditioning could make the optimization process unstable. The full-depth variant has a higher bound pLDDT and motif RMSD than MevoCraft. These results support the benefit of progressive MSA injection.

Fourth, *MSA in all states* supplies the MSA to both unbound and bound states. Unbound pLDDT increases from 0.6120 to 0.6173, which is expected because MSA conditioning can improve prediction confidence. However, unbound motif RMSD also decreases from $5.89\,\text{Å}$ to $3.34\,\text{Å}$. This decrease is undesirable for positive allostery, where the unbound state should not exhibit the target motif. With bound motif RMSD remaining similar, the RMSD gap narrows from $3.61\,\text{Å}$ to $1.01\,\text{Å}$. These results support the role of state-conditional injection in preserving conformational differences between states.

5 Conclusion

In this paper, we propose a novel method called MevoCraft for multistate protein design. MevoCraft conditions each forward pass of gradient-based hallucination with a motif MSA constructed by a single homology search on the fixed functional motif and injected only into states required to exhibit that motif. To the best of our knowledge, MevoCraft is the first method to incorporate MSA into gradient-based hallucination for multistate protein design. Experiments show that MevoCraft outperforms SwitchCraft in 50 of 70 positive-allostery cases and 49 of 70 negative-allostery cases, while matching it in another 19 and 14 cases, respectively. In motif switching, MevoCraft produces $3.8\times$ the number of successful designs achieved by SwitchCraft. Further analysis shows higher prediction confidence, lower motif RMSD, and greater separation between target states, supporting MevoCraft as a promising method to multistate protein design.

AI use statement

In this work, we used generative AI tools only for writing assistance (language polishing and drafting of the manuscript text). We have not used generative AI tools for ideation, code generation, data processing, experimental design, or any other research component, and no other disclosure categories apply to this work. All AI-assisted text was reviewed and verified by the authors, who take full responsibility for the correctness, originality, and integrity of all content in this paper.

REPRODUCIBILITY STATEMENT

The code is available at the anonymous github repository `https://anonymous.4open.science/r/MevoCraft-BA61`.

ETHICS STATEMENT

This work is purely computational: it involves no human subjects, animal experiments, or wet-lab manipulation of microorganisms.

REFERENCES

Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016):493–500, 2024.

Alex Abrudan, Sebastian Pujalte Ojeda, Chaitanya Joshi, Matthew Greenig, Felipe Engelberger, Alena Khmelinskaia, Jens Meiler, Michele Vendruscolo, and Tuomas Knowles. Multi-state protein sequence design with dynamicmpnn. In *International Conference on Learning Representations*, 2026.

Benjamin D Allen and Stephen L Mayo. An efficient algorithm for multistate protein design based on faster. *Journal of Computational Chemistry*, 31(5):904–916, 2010.

Bo Chen, Zhilei Bei, Xingyi Cheng, Pan Li, Jie Tang, and Le Song. Msagpt: Neural prompting protein structure prediction via msa generative pre-training. *Advances in Neural Information Processing Systems*, 37:37504–37534, 2024.

David Chin and Anthony R Means. Calmodulin: a prototypical calcium sensor. *Trends in cell biology*, 10(8):322–328, 2000.

Yehlin Cho, Martin Pacesa, Zhidian Zhang, Bruno E Correia, and Sergey Ovchinnikov. Boltzdesign1: Inverting all-atom structure prediction model for generalized biomolecular binder design. *BioRxiv*, pp. 2025–04, 2025.

Justas Dauparas, Ivan Anishchenko, Nathaniel Bennett, Hua Bai, Robert J Ragotte, Lukas F Milles, Basile IM Wicky, Alexis Courbet, Rob J de Haas, Neville Bethel, et al. Robust deep learning–based protein sequence design using proteinmpnn. *Science*, 378(6615):49–56, 2022.

James A Davey and Roberto A Chica. Multistate approaches in computational protein design. *Protein Science*, 21(9):1241–1252, 2012.

Diego Del Alamo, Davide Sala, Hassane S Mchaourab, and Jens Meiler. Sampling alternative conformational states of transporters and receptors with alphafold2. *Elife*, 11:e75751, 2022.

Natali A Gonzalez, Brigitte A Li, and Michelle E McCully. The stability and dynamics of computationally designed proteins. *Protein Engineering, Design and Selection*, 35:gzac001, 2022.

Mark A Hallen and Bruce R Donald. Comets (constrained optimization of multistate energies by tree search): A provable and efficient protein design algorithm to optimize binding affinity and specificity with respect to sequence. *Journal of Computational Biology*, 23(5):311–321, 2016.

Tymor Hamamsy, James T Morton, Robert Blackwell, Daniel Berenberg, Nicholas Carriero, Vladimir Gligorijevic, Charlie EM Strauss, Julia Koehler Leman, Kyunghyun Cho, and Richard Bonneau. Protein remote homology detection and structural alignment using deep learning. *Nature Biotechnology*, 42(6):975–985, 2024.

Klaus P Hoeflich and Mitsuhiko Ikura. Calmodulin in action: diversity in target recognition and activation mechanisms. *Cell*, 108(6):739–742, 2002.

Yinuo Jiang, Bing He, Zihan Wu, Fang Wang, Tianxu Lv, Yuran Jia, Yu Zhao, Chenchen Qin, Huajun Chen, Qiang Zhang, et al. Scalable homology detection with erast. *Nature Biotechnology*, pp. 1–9, 2026.

Bowen Jing, Anna Sappington, Mihir Bafna, Ravi Shah, Adrina Tang, Rohith Krishna, Adam Klivans, Daniel J Diaz, and Bonnie Berger. Generating functional and multistate proteins with a multimodal diffusion transformer. *bioRxiv*, 2025.

Bowen Jing, Mihir Bafna, Anisha Parsan, Heyuan Michael Ni, David Kwabi-Addo, Bryan Bryson, Adam Klivans, and Bonnie Berger. Switchcraft: A programmatic framework for designing state-switching proteins. *International Conference on Machine Learning*, 2026.

Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023.

Sidney Lyayuga Lisanza, Jacob Merle Gershon, Samuel WK Tipps, Jeremiah Nelson Sims, Lucas Arnoldt, Samuel J Hendel, Miriam K Simma, Ge Liu, Muna Yase, Hongwei Wu, et al. Multistate and functional protein design using rosettafold sequence space diffusion. *Nature Biotechnology*, 43(8):1288–1298, 2025.

Wei Liu, Ziye Wang, Ronghui You, Chenghan Xie, Hong Wei, Yi Xiong, Jianyi Yang, and Shanfeng Zhu. Plmsearch: Protein language model powers accurate and fast sequence search for remote homology. *Nature Communications*, 15(1):2775, 2024.

Patrick Löffler, Samuel Schmitz, Enrico Hupfeld, Reinhard Sterner, and Rainer Merkl. Rosetta: Msf: a modular framework for multi-state computational protein design. *PLoS computational biology*, 13(6):e1005600, 2017.

Milot Mirdita, Konstantin Sch"utze, Yoshitaka Moriwaki, Lim Heo, Sergey Ovchinnikov, and Martin Steinegger. Colabfold: making protein folding accessible to all. *Nature Methods*, 19:679–682, 2022.

Martin Pacesa, Lennart Nickel, Christian Schellhaas, et al. One-shot design of functional protein binders with bindcraft. *Nature*, 2025. doi: 10.1038/s41586-025-09429-6.

Max F Perutz. Stereochemistry of cooperative effects in haemoglobin: haem–haem interaction and the problem of allostery. *Nature*, 228(5273):726–734, 1970.

Michael Remmert, Andreas Biegert, Andreas Hauser, and Johannes Söding. Hhblits: lightning-fast iterative protein sequence searching by hmm-hmm alignment. *Nature Methods*, 9(2):173–175, 2012.

Marion F Sauer, Alexander M Sevy, James E Crowe Jr, and Jens Meiler. Multi-state design of flexible proteins predicts sequences optimal for conformational change. *PLoS Computational Biology*, 16 (2):e1007339, 2020.

Martin Steinegger and Johannes Söding. Mmseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nature Biotechnology*, 35(11):1026–1028, 2017.

Michel Van Kempen, Stephanie S Kim, Charlotte Tumescheit, Milot Mirdita, Jeongjae Lee, Cameron LM Gilchrist, Johannes Söding, and Martin Steinegger. Fast and accurate protein structure search with foldseek. *Nature Biotechnology*, 42(2):243–246, 2024.

Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eisenach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. De novo design of protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.

Hannah K Wayment-Steele, Adedolapo Ojoawo, Renee Otten, Julia M Apitz, Warintra Pitsawong, Marc Hömberger, Sergey Ovchinnikov, Lucy Colwell, and Dorothee Kern. Predicting multiple conformations via sequence clustering and alphafold2. *Nature*, 625(7996):832–839, 2024.

Jeremy Wohlwend, Gabriele Corso, Saro Passaro, Noah Getz, Mateo Reveiz, Ken Leidal, Wojtek Swiderski, Liam Atkinson, Tally Portnoi, Itamar Chinn, Jacob Silterra, Tommi Jaakkola, and Regina Barzilay. Boltz-1: Democratizing biomolecular interaction modeling. *bioRxiv*, 2024. doi: 10.1101/2024.11.19.624167.

Kai Yi, Kiarash Jamali, and Sjors HW Scheres. All-atom inverse protein folding through discrete flow matching. *arXiv preprint arXiv:2507.14156*, 2025.

A Additional Method Details

A.1 MSA Construction Details

We construct each motif MSA using complementary structure- and sequence-based search. Structure-based search serves as the primary source of homologs because structural similarity can identify motif-compatible proteins even when their sequences have diverged, and the resulting structural alignments provide a more reliable correspondence between residues occupying the same motif geometry. Sequence-based search is then used to increase coverage when structure databases contain few suitable hits or when structural alignments leave motif positions insufficiently represented.

For structure-based search, we submit the motif structure to the Foldseek online search server at `https://search.foldseek.com/search` (Van Kempen et al., 2024). Each returned structural alignment is projected onto the reference motif positions, producing one aligned sequence row over a common motif-column coordinate system.

For sequence-based search, we query the public ColabFold MSA API (Mirdita et al., 2022), which uses MMseqs2 (Steinegger & Söding, 2017) to search the UniRef sequence database. We use the corresponding full-length native protein sequence as the query and extract only the aligned residues mapped to the specified motif positions.

Finally, rows from both search strategies are converted to the same motif-column representation, filtered to remove low-quality sequences such as those containing unresolved "X" residues, and deduplicated.

A.2 Optimization Details

Following the design-optimization protocol of SwitchCraft (Jing et al., 2026), we optimize a shared matrix of sequence logits using a staged relaxation from continuous to discrete sequence representations. MevoCraft augments this procedure with the MSA-depth schedule in Equation (8) and the state-conditional MSA construction in Equation (6). The complete procedure is summarized in Algorithm 1.

The initial sequence logits $\boldsymbol{z}(0) \in \mathbb{R}^{20\times L}$ are sampled from a Gumbel distribution. Each column contains logits over the 20 canonical amino-acid identities. At step t, we construct soft and hard sequence representations as

$$\begin{aligned} \boldsymbol{z}_{\mathrm{soft}}(t) &= \mathrm{softmax}(\boldsymbol{z}(t)/\tau_t)\,, \\ \boldsymbol{z}_{\mathrm{hard}}(t) &= \mathrm{onehot}(\arg\max \boldsymbol{z}(t))\,, \end{aligned} \tag{9}$$

where τ_t is the temperature. A straight-through representation is defined by

$$\boldsymbol{z}_{\mathrm{st}}(t) = \mathrm{sg}(\boldsymbol{z}_{\mathrm{hard}}(t) - \boldsymbol{z}_{\mathrm{soft}}(t)) + \boldsymbol{z}_{\mathrm{soft}}(t), \tag{10}$$

where sg denotes stop-gradient. The representation supplied to the structure predictor interpolates between logits, soft probabilities, and the straight-through hard sequence:

$$\boldsymbol{z}_{\mathrm{in}}(t) = \beta_t \boldsymbol{z}_{\mathrm{st}}(t) + (1-\beta_t)\gamma_t \boldsymbol{z}_{\mathrm{soft}}(t) + (1-\beta_t)(1-\gamma_t)\boldsymbol{z}(t). \tag{11}$$

Fixed motif positions are replaced by one-hot encodings of $\mathbf{s}$ before each forward pass.

The pseudocode describes a single motif for clarity. For motif switching, the homolog set, motif-preserving state set, and sampled row indices are maintained independently for each motif, and the MSA assigned to a state contains the rows associated with the motif required in that state. Homolog rows are fixed, non-differentiable inputs; gradients propagate only through the query representation derived from $\boldsymbol{z}(t)$.

Optimization uses four stages. Stage 1 runs for 30 steps with $(\eta_t, \beta_t, \gamma_t, \tau_t) = (0.2, 0, 1, 0.5)$ and is followed by $\boldsymbol{z} \leftarrow \boldsymbol{z}_{\mathrm{in}}$. Stage 2 runs for 100 steps with $\eta_t = 0.1$, $\beta_t = 0$, $\tau_t = 0.5$, and γ_t increasing linearly from 0 to 1; it is followed by $\boldsymbol{z} \leftarrow 2\boldsymbol{z}$. Stage 3 runs for 100 steps with $\eta_t = 0.1$, $\beta_t = 0$, $\gamma_t = 1$, and τ_t decreasing from 0.5 to 0.005. Stage 4 runs for 10 steps with $(\eta_t, \beta_t, \gamma_t, \tau_t) = (0.1, 1, 1, 0.005)$.

When $0 < d(t) < K$, the injected homolog rows are sampled uniformly without replacement at each iteration; all rows are used when $d(t) = K$. States outside $\mathcal{S}_{\mathrm{motif}}$ receive only the query row.

Algorithm 1 MevoCraft design optimization

Require: target states $\mathcal{S}$ and contexts $\{\mathcal{C}_s\}$, losses $\{\mathcal{L}_n\}$, motif specification $(\mathbf{m}, \mathbf{s})$, embedded homolog rows $\widetilde{\mathcal{H}} = \{\tilde{\mathbf{h}}_i\}_{i=1}^{K}$, motif-preserving states $\mathcal{S}_{\mathrm{motif}}$, predictor f_θ;
Ensure: designed sequence $\mathbf{x}$;

Initialize $\boldsymbol{z}(0) \in \mathbb{R}^{20\times L}$ from a Gumbel distribution;
for $t = 0, \ldots, 239$ **do**
  Obtain $(\eta_t, \beta_t, \gamma_t, \tau_t)$ from the four-stage schedule;
  Compute $\boldsymbol{z}_{\mathrm{soft}}(t)$, $\boldsymbol{z}_{\mathrm{hard}}(t)$, $\boldsymbol{z}_{\mathrm{st}}(t)$, and $\boldsymbol{z}_{\mathrm{in}}(t)$;
  Replace motif columns of $\boldsymbol{z}_{\mathrm{in}}(t)$ with the one-hot encodings specified by $\mathbf{s}$;
  Set the MSA depth $d(t)$ according to Equation (8);
  if $0 < d(t) < K$ **then**
    Sample $I_t \subset \{1, \ldots, K\}$ uniformly without replacement, with $|I_t| = d(t)$;
  else if $d(t) = K$ **then**
    $I_t \leftarrow \{1, \ldots, K\}$;
  else
    $I_t \leftarrow \emptyset$;
  end if
  for all $s \in \mathcal{S}$ **do**
    if $s \in \mathcal{S}_{\mathrm{motif}}$ and $I_t \neq \emptyset$ **then**
      $M_s(t) \leftarrow [\boldsymbol{z}_{\mathrm{in}}(t); \{\tilde{\mathbf{h}}_i : i \in I_t\}]$;
    else
      $M_s(t) \leftarrow [\boldsymbol{z}_{\mathrm{in}}(t)]$;
    end if
    $\widehat{\mathcal{Y}}_s(t) \leftarrow f_\theta(\boldsymbol{z}_{\mathrm{in}}(t), \mathcal{C}_s, M_s(t))$;
  end for
  $\mathcal{L}(t) \leftarrow \sum_{n=1}^{N_{\mathrm{loss}}} \mathcal{L}_n(\{\widehat{\mathcal{Y}}_s(t)\}_{s\in\mathcal{S}_n})$;
  $\boldsymbol{z}(t+1) \leftarrow \boldsymbol{z}(t) - \eta_t \nabla_{\boldsymbol{z}(t)} \mathcal{L}(t)$;
  if $t = 29$ **then**
    Reparameterize $\boldsymbol{z}(t+1)$ using the Stage 1 input representation;
  else if $t = 129$ **then**
    $\boldsymbol{z}(t+1) \leftarrow 2\boldsymbol{z}(t+1)$;
  end if
end for
$\mathbf{x} \leftarrow \mathrm{argmax}_{\mathrm{AA}}\, \boldsymbol{z}(240)$, with motif positions fixed to $\mathbf{s}$;
return $\mathbf{x}$;

A.3 Evaluation Criteria

We adopt the success criteria of SwitchCraft (Jing et al., 2026). After optimization, each design is folded under the apo (unbound) and holo (bound) contexts. For a given design and state, motif RMSD is computed against the reference motif coordinates using Cα atoms at motif positions, and we report the mean and standard deviation over five independent structure samples. Let μ_{u} and σ_{u} denote the mean and standard deviation of motif RMSD in the unbound state, and μ_{b} and σ_{b} the corresponding quantities in the bound state. All thresholds below are in Ångstroms.

Positive allostery. A design is successful if the motif is inactive when unbound and active when bound,

$$\mu_{\mathrm{u}} > 1.0 \quad \text{and} \quad \mu_{\mathrm{b}} \leq 1.0, \tag{12}$$

and the predictions are sufficiently stable and state-discriminating,

$$\sigma_{\mathrm{u}} \leq 0.5, \quad \sigma_{\mathrm{b}} \leq 0.5, \quad \text{and} \quad |\mu_{\mathrm{u}} - \mu_{\mathrm{b}}| \geq 0.5. \tag{13}$$

Negative allostery. A design is successful if the motif is active when unbound and inactive when bound,

$$\mu_{\mathrm{u}} \leq 1.0 \quad \text{and} \quad \mu_{\mathrm{b}} > 1.0, \tag{14}$$

together with the same stability and discrimination constraints

$$\sigma_{\mathrm{u}} \leq 0.5, \quad \sigma_{\mathrm{b}} \leq 0.5, \quad \text{and} \quad |\mu_{\mathrm{u}} - \mu_{\mathrm{b}}| \geq 0.5. \tag{15}$$

Motif switching. For an ordered pair $A \to B$, let $(\mu_{\mathrm{u}}^{A}, \sigma_{\mathrm{u}}^{A})$ and $(\mu_{\mathrm{b}}^{A}, \sigma_{\mathrm{b}}^{A})$ be the unbound and bound motif-RMSD statistics for motif A, and likewise for motif B. Motif A must be active in the apo state and inactive in the holo state, while motif B must exhibit the opposite behavior. A design is successful if

$$\mu_{\mathrm{u}}^{A} \leq 1.0, \quad \mu_{\mathrm{u}}^{B} > 1.0, \quad \sigma_{\mathrm{u}}^{A} \leq 0.5, \quad \sigma_{\mathrm{u}}^{B} \leq 0.5, \tag{16}$$

$$\mu_{\mathrm{b}}^{B} \leq 1.0, \quad \mu_{\mathrm{b}}^{A} > 1.0, \quad \sigma_{\mathrm{b}}^{A} \leq 0.5, \quad \sigma_{\mathrm{b}}^{B} \leq 0.5. \tag{17}$$

That is, both motifs must satisfy reciprocal activity across the two states, and each of the four motif–state mean RMSDs must have standard deviation at most 0.5 Å. Unlike the allostery criteria, motif switching does not impose an additional unbound–bound gap threshold on either motif.

B Additional Results

B.1 Motif Lengths

Table S1 summarizes the lengths of the candidate motifs. We exclude `1ycr` and `1qjg` from our evaluation because each contains fewer than 10 residues.

Table S1: Motif lengths (number of residues) for motifs used in MevoCraft.

Motif	Length	Motif	Length	Motif	Length
1bcf	32	1prw	40	1qjg	3
1ycr	9	2kl8	59	3ixt	24
4jhw	24	4zyp	15	5ius	42
5tpn	19	5trv_long	21	5trv_med	21
5trv_short	21	5wn9	20	5yui	11
6e6r_long	13	6e6r_med	13	6e6r_short	13
6exz_long	15	6exz_med	15	6exz_short	15
7mrx_128	22	7mrx_60	22	7mrx_85	22

B.2 Positive and Negative Allostery

Figure S1 compares motif-level success rates for MevoCraft and SwitchCraft, aggregated across conditioning molecules, for positive and negative allostery.

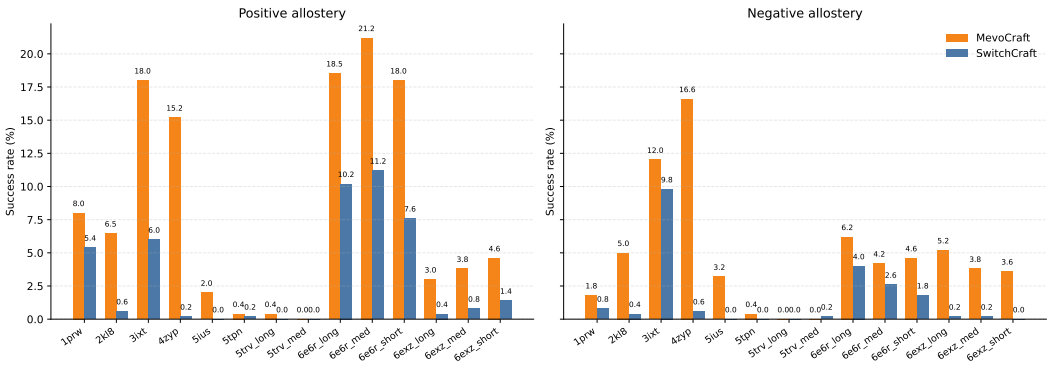
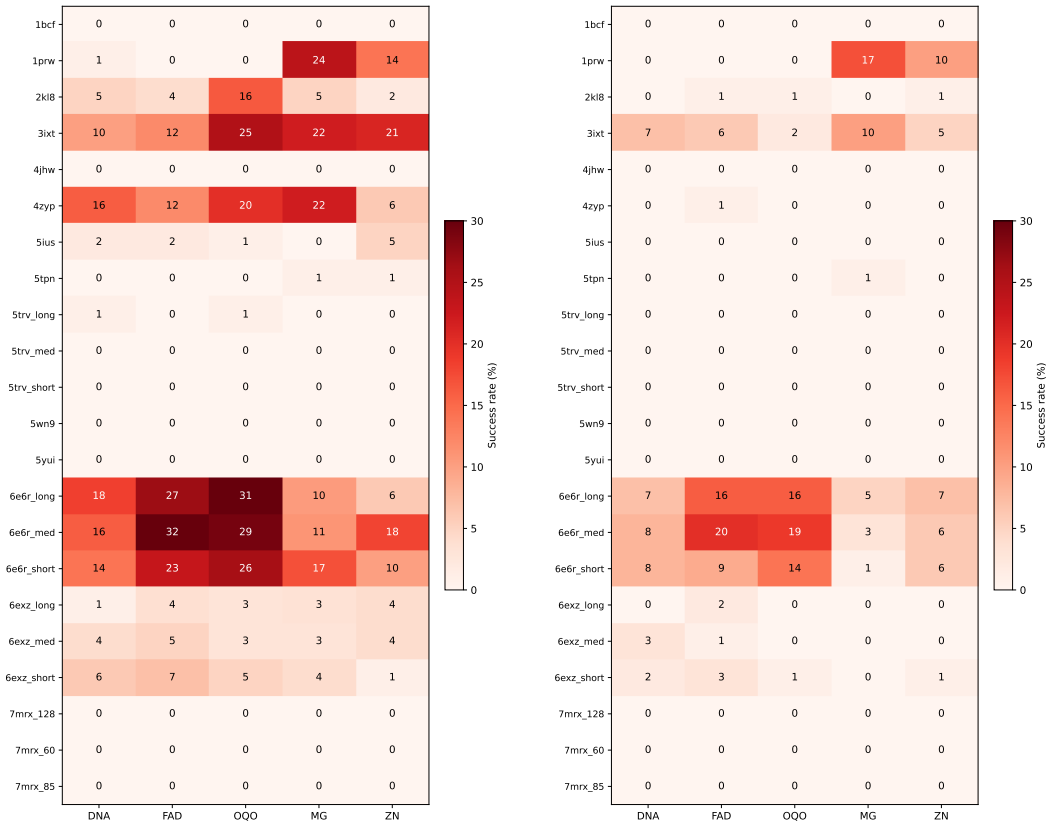


Figure S1: Motif-level success rates (%) for positive allostery (left) and negative allostery (right), aggregated across conditioning molecules. Orange bars represent MevoCraft and blue bars represent SwitchCraft. Values above the bars indicate the corresponding success rates.

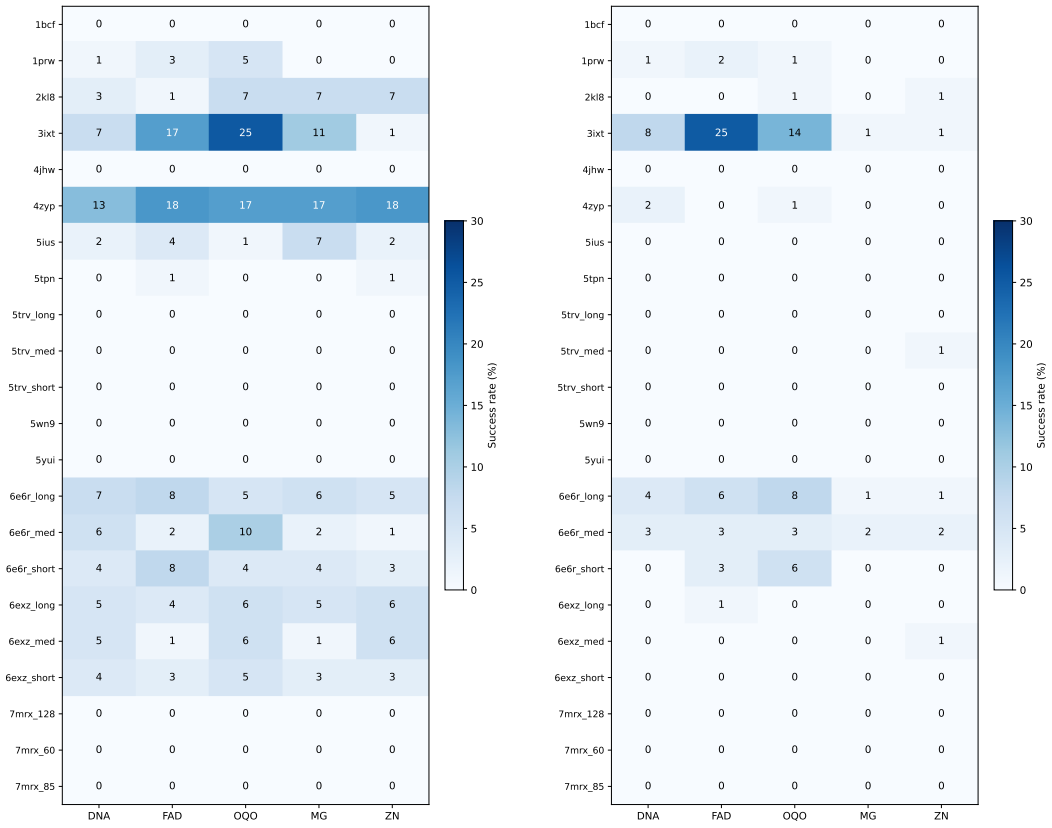
Figure S2 reports the positive-allostery success rates for MevoCraft and SwitchCraft across all 22 evaluated motifs and five conditioning molecules, including motifs omitted from the main-text comparison because both methods produce zero successful designs.



(a) MevoCraft (b) SwitchCraft

Figure S2: Positive-allostery success rates (%) for MevoCraft (left) and SwitchCraft (right). Rows denote motifs and columns denote conditioning molecules. Each cell reports the success rate over 100 designs.

Figure S3 reports the negative-allostery success rates for MevoCraft and SwitchCraft across all 22 evaluated motifs and five conditioning molecules, including motifs omitted from the main-text comparison because both methods produce zero successful designs.



(a) MevoCraft (b) SwitchCraft

Figure S3: Negative-allostery success rates (%) for MevoCraft (left) and SwitchCraft (right). Rows denote motifs and columns denote conditioning molecules. Each cell reports the success rate over 100 designs.

B.3 MOTIF SWITCHING

Figure S4 reports the successful design counts for each method underlying the differences shown in Figure 4.

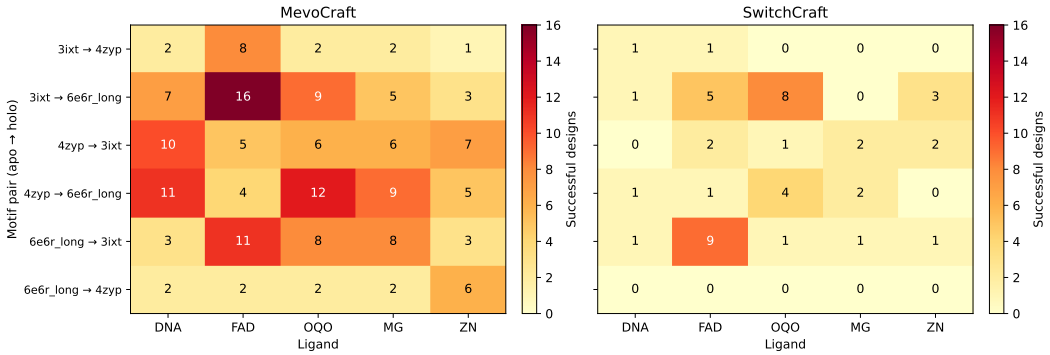


Figure S4: Successful design counts for motif switching with MevoCraft (left) and SwitchCraft (right). Each cell reports the number of successful designs for an ordered motif pair and conditioning molecule. Rows specify the target motifs in the apo and holo states, respectively, in the direction indicated by the arrow. Both panels use the same color scale.