





StructEvo: Structure-aware Reinforcement Learning for Protein Directed Evolution





Key challenges in protein directed evolution

- Protein directed evolution aims to improve protein function by identifying **optimal combinations of amino acid mutations**. It is widely applied in protein engineering, including enzyme engineering and antibody affinity maturation.
- Key challenges in protein directed evolution:
 - **Vast Search Space**
 - **Limited Experimental Budget**
 - **Sparse Beneficial Mutations**
 - **Strong Epistatic Effects**



How can we discover high-fitness protein mutants more efficiently under limited experimental budgets?



Why structure matters ?


- **Structure reveals spatial proximity:** Residues distant in sequence may directly interact in 3D space and exhibit synergistic effects.
- **Structure captures local microenvironments:** Is a residue located at a functional site, an interaction interface, or within a hydrophobic core?

Sequence-only search struggles to capture the spatial context of mutation sites and interactions among residues.

However, mutant structures are often unavailable:

- **Limited experimentally resolved** mutant structures database
- **Low mutation sensitivity** of existing structure prediction models
- **High computational cost** of predicting structures for every mutant

How can we effectively leverage structural knowledge to guide mutation search?



Schematic of GFP topology. Residues distant in sequence may be spatially adjacent and exhibit cooperative interactions.

➤ End-to-End Structural Change Awareness

Approximate mutation-induced changes in structural representations from sequence differences, avoiding explicit prediction of potentially unreliable 3D structures for individual mutants.

➤ Structure-Guided Mutation Policy

Integrate sequence and structural representations to capture local spatial environments and potential interactions among residues.


➤ From Fitness Prediction to Mutation Proposal

Formulate protein directed evolution as a sequential decision-making problem under a reinforcement learning framework, directly learning how to explore the mutation space.

Formulating Protein Directed Evolution as a Reinforcement Learning Problem

State

Protein Sequence
 $S_t = (x_t^1, x_t^2, \dots, x_t^L)$



MKTAYIAKQR...RQD

Policy Model
 $\pi(a_t | s_t)$

Action

$a_t = (a_t^{pos}, a_t^{type})$

MKTAYIAKQR...RQD

MKTAYIMKQR...RQD

Environment

Trajectory-level Sparse Reward


$$r_t = r(s_t, a_t) = \begin{cases} f_{proxy}(S_{t+1}), & S_{t+1} \text{ is a better mutant than } S_0 \\ 0, & \text{otherwise} \end{cases}$$

Proxy model

Fitness
Evaluation

MKTAYIMKQR...RQD

StructEvo: Structure-aware Reinforcement Learning for Protein Directed Evolution




- **Delta-Structure Fusion Encoder:** Approximate structural feature changes from sequence representation differences
- **Hierarchical Mutation Policy:** Decompose mutation decisions into position selection followed by amino acid selection
- **Geometric Constraint:** Enforce sequence-structure consistency and training robustness

Searching for High-Fitness Mutants Across Four Real-World Tasks

Target Protein	Protein Function	Fitness	Mutation Space	Oracle
GB1	Immunoglobulin-binding domain	IgG-Fc binding affinity	Four-site combinatorial	Experimental measurements
PhoQ	Sensor histidine kinase	Signal transduction activity	Four-site combinatorial	Experimental measurements
AAV	Adeno-associated virus capsid protein	Capsid fitness	Full-length (28 aa)	Learned oracle model
GFP	Green fluorescent protein	Fluorescence intensity	Full-length (237 aa)	Learned oracle model

Shared Characteristics Across All Four Tasks:

- **Sparse High-Fitness Mutants**
- **Complex Epistatic Effects**



Exploring Four-Site Combinatorial Mutation Spaces

GB1 Benchmark: Affinity optimization through combinatorial mutations in the immunoglobulin-binding domain.
StructEvo better captures synergistic interactions among multiple mutations.



Table 9: **Results on GB1 benchmarks across 3 rounds.** The best scores are **bolded** and the second best are underlined. All baseline results are derived from Wang et al. (2024).

Method	Round 1			Round 2			Round 3		
	Max↑	Mean↑	NDCG↑	Max↑	Mean↑	NDCG↑	Max↑	Mean↑	NDCG↑
MLDE	0.650	0.183	0.767	0.680	0.217	0.794	0.684	0.203	0.789
ftMLDE (EVmutation)	0.725	0.233	0.791	0.770	0.280	0.814	0.935	0.414	0.833
ftMLDE (Transformer)	0.761	0.239	0.792	0.814	0.298	0.819	0.932	0.416	0.813
EvoPlay	0.837	0.433	0.826	0.834	0.457	0.839	0.840	0.460	0.847
CLADE	0.785	0.309	0.801	0.802	0.303	0.803	0.835	0.458	0.857
CLADE 2.0	0.886	0.376	0.808	0.903	0.419	0.858	0.935	0.491	0.879
KnowRLM	<u>0.931</u>	<u>0.494</u>	<u>0.851</u>	<u>0.972</u>	<u>0.534</u>	<u>0.862</u>	<u>0.972</u>	<u>0.562</u>	<u>0.884</u>
StructEvo (ours)	0.942	0.514	0.870	0.973	0.540	0.902	0.990	0.587	0.932

PhoQ Benchmark: Combinatorial mutation optimization of signal transduction in a sensor histidine kinase.
StructEvo better models the relationship between multi-site mutations and protein function.




Table 10: **Results on PhoQ benchmarks across 3 rounds.** The best scores **bolded** and the second best underlined. All baseline results are derived from Wang et al. (2024).


Method	Round 1			Round 2			Round 3		
	Max↑	Mean↑	NDCG↑	Max↑	Mean↑	NDCG↑	Max↑	Mean↑	NDCG↑
MLDE	0.309	0.069	0.753	0.364	0.087	0.775	0.361	0.095	0.791
ftMLDE (EVmutation)	0.297	0.072	0.754	0.431	0.114	0.807	0.436	0.115	0.804
ftMLDE (Transformer)	0.346	0.073	0.756	0.414	0.108	0.802	0.422	0.117	0.815
EvoPlay	0.443	0.119	0.782	0.444	0.135	0.786	0.474	0.143	0.804
CLADE	0.319	0.070	0.759	0.441	0.089	0.762	0.467	0.089	0.777
CLADE 2.0	0.345	0.118	0.781	0.383	0.125	0.779	0.399	0.148	0.814
KnowRLM	<u>0.486</u>	<u>0.129</u>	<u>0.821</u>	<u>0.532</u>	<u>0.152</u>	<u>0.816</u>	0.658	<u>0.157</u>	<u>0.819</u>
StructEvo (ours)	0.530	0.158	0.844	0.581	0.169	0.848	<u>0.640</u>	0.186	0.856

Exploring Long-Range Mutations in Full-Length Benchmarks


AAV & GFP Long-Range Mutation Benchmarks: Guided by protein structural information, StructEvo better captures long-range mutational interactions and successfully discovers higher-fitness variants in vast sequence spaces of up to 20^{28} and 20^{237} possible sequences, respectively.

Table 2: **Performance on AAV and GFP benchmarks.** The results include average and standard deviation across 5 independent runs. *: results are derived from Kirjner et al. [30]. †: results are derived from Lee et al. [13]. -: not reported in the original paper.

Method	AAV medium				AAV hard			
	Mean↑	Max↑	Diversity	Novelty	Mean↑	Max↑	Diversity	Novelty
CMA-ES	0.05±0.00	0.40±0.01	20.9±0.5	17.2±0.5	0.05±0.00	0.32±0.03	20.4±0.7	18.6±0.5
BO	0.64±0.03	0.72±0.04	8.1±0.2	9.4±0.9	0.62±0.03	0.70±0.03	8.6±0.2	10.3±0.3
PEX	0.65±0.01	0.74±0.00	5.5±0.5	5.3±0.3	0.63±0.02	0.75±0.01	5.7±0.5	6.2±0.4
GGs*	0.51±0.01	-	4.0±0.2	5.4±0.5	0.60±0.02	-	4.5±0.5	7.0±0.0
AdaLead	0.74±0.04	0.80±0.05	5.2±0.9	7.6±0.4	0.76±0.03	0.81±0.03	4.0±0.4	8.9±0.6
VLGPO	0.69±0.01	0.76±0.01	3.1±0.2	5.7±0.1	0.71±0.01	0.78±0.01	3.3±0.2	6.9±0.1
LatProtRL†	0.71±0.02	-	5.4±0.8	6.0±0.2	0.66±0.01	-	6.0±0.8	7.0±0.6
GFN-δ-CS	0.77±0.02	0.81±0.02	4.4±1.1	5.8±1.5	0.76±0.03	0.80±0.03	4.0±0.7	6.7±1.1
StructEvo	0.82±0.01	0.86±0.01	4.5±0.7	8.6±0.2	0.83±0.03	0.87±0.04	4.5±0.3	8.4±0.5
Method	GFP medium				GFP hard			
	Mean↑	Max↑	Diversity	Novelty	Mean↑	Max↑	Diversity	Novelty
CMA-ES	-0.04±0.01	0.58±0.00	172.7±2.1	154.3±4.9	-0.03±0.00	0.16±0.03	173.8±2.0	218.6±7.6
BO	0.30±0.07	0.42±0.07	26.4±5.6	37.7±12.1	0.26±0.07	0.39±0.04	36.0±2.7	71.7±10.1
PEX	0.64±0.04	0.84±0.04	8.3±2.6	6.8±1.9	0.46±0.05	0.62±0.05	7.0±1.2	10.3±1.7
GGs*	0.76±0.01	-	3.7±0.2	5.0±0.1	0.74±0.00	-	3.6±0.0	6.9±0.1
AdaLead	0.93±0.01	0.99±0.02	5.2±0.9	9.5±1.2	0.75±0.04	0.81±0.04	5.4±0.6	13.5±1.1
VLGPO	0.93±0.03	1.02±0.01	2.8±0.2	5.3±0.0	0.79±0.02	0.95±0.00	1.9±0.0	6.2±0.0
LatProtRL†	0.93±0.00	-	4.6±0.4	5.5±0.1	0.85±0.01	-	4.8±0.3	7.0±0.4
GFN-δ-CS	1.06±0.04	1.09±0.03	7.0±0.4	9.5±1.9	0.86±0.06	0.90±0.05	4.3±1.6	8.2±1.5
StructEvo	1.11±0.02	1.17±0.02	6.3±0.8	11.9±0.8	1.00±0.04	1.05±0.05	4.9±0.3	11.4±0.6




Retrospective Validation: StructEvo Recovering an Experimentally Validated Triple-Mutation Epistasis Pattern



StructEvo identifies the **F64L/V163A/I171V triple-mutation pattern** associated with high GFP fluorescence.

Mutation Combinations	F64L	✓			✓	✓	✓	✓
	V163A		✓		✓		✓	✓
	I171V			✓		✓	✓	✓
Average Fluorescence Intensity	0.06	0.08	0.07	0.02	0.23	0.11	0.09	0.91



- **Biological Interpretation:** All three substitutions reduce side-chain bulk, collectively reshaping the structural environment around the chromophore. Their synergistic interactions contribute to a substantial increase in fluorescence intensity.
- **Incorporating structural information enables StructEvo to uncover long-range, higher-order epistatic patterns that sequence-only methods fail to identify.**