PROTEIN SEQUENCE DESIGN IN A LATENT SPACE VIA MODEL-BASED REINFORCEMENT LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Proteins are complex molecules responsible for different functions in the human body. Enhancing the functionality of a protein and/or cellular fitness can significantly impact various industries. However, their optimization remains challenging, and sequences generated by data-driven methods often fail in wet lab experiments. This study investigates the limitations of existing model-based sequence design methods and presents a novel optimization framework that can efficiently traverse the latent representation space instead of the protein sequence space. Our framework generates proteins with higher functionality and cellular fitness by modeling the sequence design task as a Markov decision process and applying model-based reinforcement learning. We discuss the results in a comprehensive evaluation of two distinct proteins, GPF and His3, along with the predicted structure of optimized sequences using deep learning-based structure prediction.

1 INTRODUCTION

Proteins mediate the fundamental processes of cellular fitness and life. Iterated mutations on various proteins and natural selection during the biological evolution diversify traits, eventually accumulating beneficial phenotypes. Similarly, in protein engineering and design, the directed evolution of proteins has proved to be an effective strategy for improving or altering the proteins' functions or cellular fitness for industrial, research, and therapeutic applications (Yang et al., 2019; Huang et al., 2016). However, the protein sequence space of possible combinations of 20 amino acids is too large to search exhaustively in the laboratory, even with high-throughput screening from the diversified library (Huang et al., 2016). In other words, directed evolution becomes trapped at local fitness maxima where library diversification is insufficient to cross *fitness valleys* and access neighboring fitness peaks. Moreover, functional sequences in this vast space of sequences are rare and overwhelmed by nonfunctional sequences.

Lately, to tackle the limitations, data-driven methods have been applied to protein sequence design. They used reinforcement learning (RL) (Angermueller et al., 2019), Bayesian optimization (Wu et al., 2017; Belanger et al., 2019; Terayama et al., 2021), and generative models (Jain et al., 2022; Kumar & Levine, 2020; Hoffman et al., 2022) in a model-based fashion, *i.e.*, using a protein functionality predictor trained on experimental data to model the local landscape. Despite the advances obtained by these methods, it is still challenging to generate optimized sequences that are experimentally validated. We suggest that the cause for this is two-fold. The first cause is that the optimization process is usually performed by generating candidate sequences through amino acid substitutions (Belanger et al., 2019) or additions (Angermueller et al., 2019). Given the vast search space, these methodologies are highly computationally inefficient and commonly lead to the exploration of parts of the space with a low chance of having functional proteins. To tackle this problem, in this paper we propose to optimize sequences in a latent representation space rather than in the protein sequence space. Actions, *e.g.,* small perturbations in the latent vector, taken in this representation space can intuitively be understood as *walking* through a local functionality/fitness landscape.

The second cause is related to the models used as an oracle for optimization. These models are trained on experimental data obtained for a specific function, covering only a small portion of the protein space, and their accuracy is consequently restricted to this small region. Later, we will demonstrate that even the most advanced model-based optimizations can assign high functionality values for randomly generated protein sequences. These random sequences, however, are unlikely to

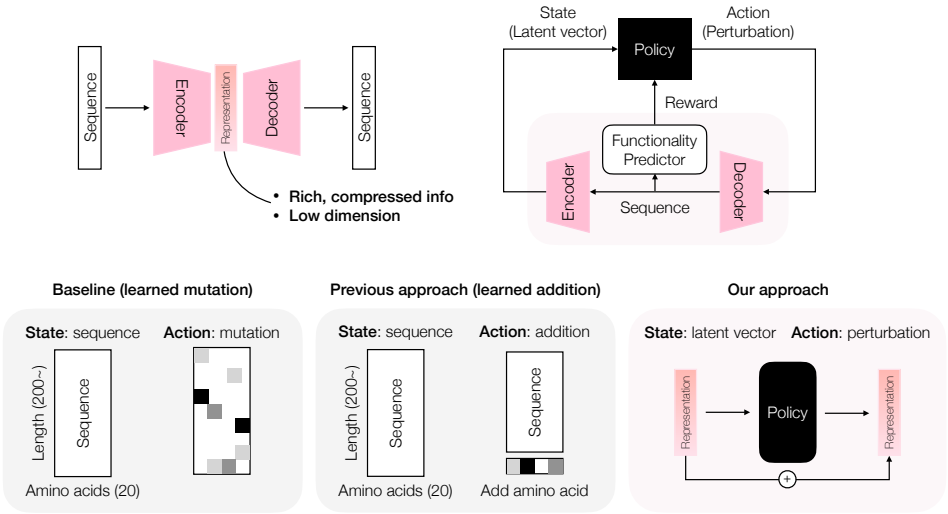

Figure 1: Overview of the Proposed Framework

be functional. To reduce false positives, we suggest augmenting the experimental data with random sequences (*i.e.*, negative examples) assigned to a low functionality value. Such augmentation also helps set boundaries around the experimental data distribution in which the oracle can be trusted.

We model protein sequence design as a Markov Decision Process (MDP) to optimize the latent representation for the first time. Our method trains the optimization policy using a model-based deep reinforcement learning (RL) framework (Fig. 1). At each timestep, the policy updates the latent representation by small perturbations to maximize protein functionality or cellular fitness, *i.e., walking* uphill through the local landscape until the end of an episode. We evaluate our method in two evaluation tasks, optimizing the functionality of the green fluorescence protein (GFP) and the cell fitness of imidazoleglycerol-phosphate dehydratase (His3). Our results show that the proposed framework design sequences with higher protein functionality and cellular fitness than existing methods. Ablation studies show that, based on the evaluation of various model options for state and action for the RL framework, the proposed latent representation update can successfully optimize the protein and search the vast design space. We provide visual evidence that the trained policy can traverse the local functionality landscape efficiently for GFP and His3. Our method is general and can also be applied to representations learned from protein structures, such as those presented in Zhang et al. (2022); Eguchi et al. (2022).

2 Related Works

Protein representation learning Representation learning methods aim to learn compact and expressive features describing data. Since a protein is composed of a sequence of distinct amino acids (N=20), it can be interpreted as a large word in which each character is one amino acid. Due to the similarity with natural language processing (NLP) tasks, methods such as BERT (Devlin et al., 2018) and GPT-2 (Radford et al., 2019) have been used to train character-level protein language models (Alley et al., 2019; Brandes et al., 2022; Lin et al., 2022; Ferruz et al., 2022; Rives et al., 2021). The protein language model trained with 250 million sequences by Rives et al. (2021) has shown to learn representations containing meaningful information about biological properties and reflecting protein structure. It was shown that the learned representation could be generalized across different applications achieving state-of-the-art results for supervised prediction of mutational effects. Representation learning has also been applied to learn protein structures (Zhang et al., 2022). An antibody-specific structure encoder was trained using variational autoencoders by Eguchi et al. (2022).

Protein functionality prediction Similar to protein representation learning, for functionality prediction, methods inspired by NLP continue to achieve high performance. For the TAPE bench-

mark (Rao et al., 2019), long short-term memory (LSTM) (Hochreiter & Schmidhuber, 1997) and transformers (Vaswani et al., 2017) are the top performing methods. Representing the protein as an image and using a convolutional neural network architecture (He et al., 2016) are also investigated in (Rao et al., 2019). Also proposed was the use of representations learned from a protein language model for zero-shot predictions (Alley et al., 2019; Meier et al., 2021). Utilizing protein structures as inputs, graph neural networks have been used to predict functionality in (Gligorijević et al., 2021). Finally, combining a sequence-based and structure-based representation for functionality prediction using language models and geometric vector perceptrons is proposed in (Wang et al., 2022).

Biological sequence design Various machine learning methods have been used to design biological sequences. DynaPPO is a model-based RL algorithm proposed specifically for biological sequence design (Angermueller et al., 2019), which generates sequence from left to right by adding amino acids. Bayesian optimization was applied for cell fitness maximization in (Brookes et al., 2019; Belanger et al., 2019; Terayama et al., 2021; Swersky et al., 2020). An evolutionary algorithm for efficient biological sequence optimization was proposed in (Sinai et al., 2020). Methods using generative models to search and sample optimized sequences are investigated in (Brookes et al., 2019; Kumar & Levine, 2020; Jain et al., 2022; Hoffman et al., 2022; Melnyk et al., 2021).

3 METHODOLOGY

This section describes the primary components of the model-based RL framework, depicted in Fig. 1. Details of network architectures are provided in Sec. 4 and Appendix A.

3.1 PROTEIN REPRESENTATION LEARNING

We use an encoder-decoder architecture, in which a pre-trained encoder represents protein sequences in the latent space, and a decoder uses the learned representation to recover the protein sequence.

Let $x \in \mathcal{X}$ be a (protein) sequence of length L, where L denotes the number of amino acids in the sequence and $\mathcal{X}$ is the sequence space. Each element x_i of the sequence is a discrete number representing one of 20 amino acids, and i is the i-th element from x. Given x, the encoder network produces a 2-dimensional vector representation $q \in \mathcal{Q}$ of dimensions (L, E), where E is the dimension size in the latent space, and $\mathcal{Q}$ is the representation space. A pre-trained protein language model is used to model the encoder network. We employ the ESM-2 model (Lin et al., 2022) trained on the Uniref50 dataset (Suzek et al., 2015). This model uses a BERT (Devlin et al., 2018) style encoder (See Appendix A.1 for details) and is trained in an unsupervised fashion with 650 million parameters. For our model, ESM-2 is used to map mutation effects onto the latent space.

The output of the encoder with dimensions (L, E) is preprocessed before being passed to the policy as input. In ESM-2, the CLS token is prepended to the sequence and functions as the start-of-sequence character. We use these token embeddings and reduce the dimension from (L, E) to $(E,)$. However, this could still be substantial, for example, a GFP protein of length 237 has a representation size of $L \times E = 237 \times 1,280$ when using ESM-2 with $E = 1,280$, which is a large action space. We hence use average pooling to further reduce the dimension to $(R,)$.

We recover the amino acid sequence from the reduced representation of size $(R,)$. Using a linear layer, the size of the representation is expanded to $(E,)$. To recover the L dimension, we concatenate the reduced representation with the wild-type representation of the pre-trained encoder of size (L, E), followed by a linear layer. The input to the decoder's ESM-2 head is the recovered representation of the original size (L, E), and the output is the recovered sequence.

3.2 PROTEIN FUNCTIONALITY PREDICTION

The functionality predictor network is trained as a downstream task of the pre-trained encoder network. Given an input sequence x, the predictor returns a value $k \in \mathbb{R}$, where k is the functionality metric to be maximized and $\mathbb{R}$ is the set of real numbers representing possible values for this metric. k is predicted from a representation q with dimensions (L, E), and the functionality predictor $f(q)$ is assumed to be trained independently for each task. To train the functionality predictor, negative examples are included when sampling the training batch to address the problem where the model

assigns high values to sequences outside the experimental data distribution. Negative examples are defined as random sequences with a zero functionality value. This was an important data augmentation method to reduce false positives. The augmentation can help set boundaries around the experimental data region that can be trusted.

3.3 Protein Sequence Design via Model-based Reinforcement Learning

The final component trains a policy using an off-policy RL algorithm that models the reward function based on the functionality predictor.

System Modeling Consider a policy π that interacts with a fully observable environment during each episode. An episode is defined by the number of timesteps in which the policy takes actions to optimize the target sequence in the latent space. We sample an initial state $s_0 \in \mathcal{Q}$ from the representation space (state space) at the beginning of each episode. At each timestep t, the agent observes a state $s_t \in \mathcal{Q}$ and selects an action $a_t \in \mathcal{A}$ according to a stochastic policy $\pi : \mathcal{Q} \rightarrow \mathcal{A}$, where $\mathcal{A}$ is the action space. The action a_t is defined as a perturbation in which a continuous value $a_{t_j} \in [-\epsilon, \epsilon]$ is chosen for the j-th dimension, and ϵ is a hyperparameter that can be tuned based on the representation distribution of the experimental data. The hyperparameter ϵ controls how conservative the policy is when traversing the functionality or fitness landscape. Both s_t and a_t have E dimensions. The next state s_{t+1} is sampled from a transition probability function $p(s_{t+1}|s_t, a_t)$.

We model the environment using a Markov Decision Process (MDP) and assume that the next state s_{t+1} is conditioned only on the current state s_t and action a_t. This function calculates the element-wise sum of s_t and a_t so that $s_{t+1} = s_t + a_t$. As a result of performing the action a_t, the agent receives the reward r_t. The reward is a function $r_t = r(s_t, a_t)$, where r is the reward function $r : \mathcal{Q} \rightarrow \mathbb{R}$, and $\mathbb{R}$ is the set of real numbers representing reward values. This data is later used to train the policy using an off-policy reinforcement learning algorithm. The agent interacts with the environment until reaching a terminal state, which is one of the followings: (i) when the last timestep T is reached; (ii) when $f(s_t)$ for the current state is greater than a hyperparameter threshold; and (iii) when s_t is in a part of the representation space $\mathcal{Q}$ where $f(s_t)$ cannot be trusted. Condition (iii) states how far s_t is from the experimental data. Each timestep's data is stored in a replay buffer as a 5-tuple $(s_t, a_t, s_{t+1}, r_t, d_t)$ where d_t identifies the terminal state.

Off-Policy Reinforcement Learning We further optimize the policy using Soft Actor Critic (SAC) (Haarnoja et al., 2018), an off-policy entropy-regularized reinforcement learning (RL) algorithm for continuous action spaces. Off-policy is a type of RL learning that learns to optimize independently of the agent's actions. SAC is trained to maximize future rewards and entropy to promote randomness in action space exploration. An optimal policy π^* following this objective is given as

$$\pi^* = arg \max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t (r(s_t, a_t) + \alpha H(\pi(\cdot|s_t))) \right], \tag{1}$$

where $\gamma \in [0, 1]$ is a discount factor, α is a temperature hyperparameter and H is entropy. Details regarding SAC optimization and how actions are sampled by the policy π are given in Appendix A.3.

4 Results

4.1 Experiment Setup

Datasets We chose proteins of different lengths and functions for robustness in evaluation: the green fluorescent protein (GFP) that can be found naturally in jellyfish and imidazoleglycerol-phosphate dehydratase (His3) that is a key enzyme in our human body. The dataset proposed in (Sarkisyan et al., 2016) is used to train the GFP encoder-decoder and its functionality predictor. The dataset contains 54,025 mutant sequences, with log-fluorescence intensity associated with each sequence. The length L of protein sequences is 237. For the His3 protein, we used the dataset in (Pokusaeva et al., 2019) that consists of mutant sequences of its evolutionarily-relevant segments and is associated with its growth rate. We processed the data to 50,000 sequences with the length of $L = 30$. This dataset is used to train the His3 encoder-decoder and fitness predictor. Datasets were split into train and test sets.

Implementation Details The encoder and decoder use a pre-trained ESM-2 model with 650 million trainable parameters. The latent representation space is set to $R = 8$, which sets the size of state and action vectors. The magnitude of the perturbation ϵ applied to each element of the action vector is set to 0.1. The episode length T is set to 20 for GFP and 10 for His3. We set three experimental rewards: (1) a *dense* reward, defined as $r_t = f(s_t)$, (2) an *absolute* reward, defined as a binary value of $r_t = 1$, if $f(s_t) > r_{th}$ and 0 if otherwise (values for r_{th} are listed in Table 4), and (3) a *binary* reward, defined as a binary value of $r_t = 1$, if $f(s_t) > f(s_{t-1})$ and 0 if otherwise. To set the initial state s_0 of GFP, we sampled sequences with low (1.4 to 1.7) and high (3.2 to 3.4) functionality values from the test set. For His3, we sampled sequences with fitness values between 0.5 and 0.8. Appendix A describes other details regarding selecting initial states, network architectures, training hyperparameters, and optimization of our method and baseline methods.

Evaluation Metrics We use three evaluation metrics as in (Hoffman et al., 2022): performance, novelty, and diversity. The performance evaluation metric is calculated as the mean log-fluorescence intensity from the top K generated sequences. Let the generated sequences be contained in the following set $\mathcal{G}^* = \{g_1^*, \cdots, g_K^*\}$, performance is defined as $\frac{1}{K}\sum_i f(g_i^*)$. The novelty evaluation metric is defined to evaluate if the policy is generating sequences similar to the ones contained in the experimental data. Defining $\mathcal{D}$ as the experimental data set containing the wild-type protein sequence, novelty is given as $\frac{1}{K}\sum_i 1([g_i^* \notin \mathcal{D}])$. The diversity evaluation metric is defined as the mean of the number of amino acids that are different among the optimized sequences and is given as follows

$$\frac{1}{K(K-1)} \sum_{g_i^* \in \mathcal{G}^*} \sum_{g_j^* \in \mathcal{G}^* - \{g_i^*\}} d(g_i^*, g_j^*), \tag{2}$$

where d is defined as the number of different amino acids of two sequences. When testing the protein functionality of GFP, we include the presence of the chromophore region (residues SYG in the wild-type protein) in the optimized sequence, as these residues are related to the ability to emit fluorescence. We generate 100 sequences for each method, and the ten highest-performing sequences are utilized. Evaluating and ranking designed sequences based on performance is problematic in parts of the protein space where the oracle is inaccurate. This issue directly impacts the optimization methods' ability to generate de novo (novel) proteins outside the trust region, and it will be discussed in greater detail in Sec. 5.

4.2 Experiment

GFP Sequence Design The results obtained for GFP sequence design are described in Table 1. We compare with four optimization methods: CbAS (adaptive sampling) (Brookes et al., 2019), BO (bayesian optimization) (Swersky et al., 2020), GFlowNet (generative model) (Jain et al., 2022), and DynaPPO (reinforcement learning) (Angermueller et al., 2019). The baseline refers to the average functionality values of the initial states used in the optimization process. We also include random mutations of initial states. Our framework outperforms the baseline. Only the proposed method and CbAsoptimize GFP effectively, whereas BO, GFlowNet, and DynaPPO achieve low performance. Even though our method limits the action to a small step in the latent space, it gets more diversity than CbAS, which also limits the search space when optimizing.

His3 Sequence Design Table 2 displays the results for His3, where our framework achieves the best performance. DynaPPO and BO were unable to optimize His3 effectively. Given that the length of the His3 protein is only 30, it is interesting to observe that a single random mutation is an effective strategy. Note, however, that even though the proposed method achieves higher overall performance, its novelty is lower than all other methods. This indicates that the decoder is recovering similar sequences within the representation space. The proposed method achieves higher diversity when compared to CbAS.

4.3 Ablation studies

State and Action Modeling Table 3 shows how state and action modeling influence sequence design in GFP. For sequence modeling, we examine two types: a representation in which each amino

Model	Performance	Novelty	Diversity	Chromophore
Ours	**3.531** ± 0.06	100%	6.311	100%
Random mutation (N=1)	3.410 ± 0.094	80%	7.716	100%
Baseline	3.370 ± 0.013	-	4.858	100%
CbAS (Brookes et al., 2019)	3.328 ± 0.044	80%	1.956	100%
Random mutations (N=5)	2.354 ± 0.522	100%	12.368	90%
Random perturbation	1.973 ± 0.832	100%	14.616	100%
BO (Swersky et al., 2020)	1.231 ± 0.034	100%	6.867	100%
GFlowNet (Jain et al., 2022)	0.017 ± 0.000	100%	12.533	0%
DynaPPO (Angermueller et al., 2019)	0.014 ± 0.001	100%	224.133	0%

Table 1: Results obtained for GFP sequence design ($K = 10$)

Model	Performance	Novelty	Diversity
Ours	**0.961** ± 0.050	60%	3.521
CbAS (Brookes et al., 2019)	0.889 ± 0.092	90%	2.356
Random mutation (N=1)	0.856 ± 0.070	80%	7.716
Baseline	0.756 ± 0.013	-	6.942
Random mutations (N=5)	0.518 ± 0.184	100%	12.368
DynaPPO (Angermueller et al., 2019)	-0.067 ± 0.053	100%	27.4674
BO (Swersky et al., 2020)	-0.089 ± 0.029	100%	4.756

Table 2: Results obtained for His3 sequence design ($K = 10$)

acid is represented by one-hot encoding and a latent vector that is the encoder's output. We investigate three types of action modeling: (i) a mutation in the amino acid sequence, (ii) the autoregressive addition of amino acids proposed in (Angermueller et al., 2019), and (iii) the proposed perturbation in the latent vector. Only the latent vector as the state and the perturbation as the action can simultaneously optimize GFP and His3 tasks. The sequence alone cannot easily convey structural information about a protein's functional sites, making it unsuitable for guiding the optimization process. Also, adding amino acids is ineffective for traversing the vast and sparse protein space.

State	Action	Performance	
		GFP	His3
Latent vector	Perturbation on latent vector	3.531 ± 0.06	0.901 ± 0.039
Baseline		3.370 ± 0.013	0.756 ± 0.013
Sequence	Sequence	0.014 ± 0.001	0.000 ± 0.037
Sequence	Amino acid addition	0.014 ± 0.001	-0.067 ± 0.053
Latent vector	Sequence	0.012 ± 0.001	-0.172 ± 0.044

Table 3: Comparison on state and action modeling. For rows with the action defined as a mutation in the protein sequence, we train the policy using Proximal Policy Optimization (PPO) (Schulman et al., 2017) to handle the multiple discrete action space setting. More details regarding the state and action modeling for this ablation are presented in Appendix A.3.

Initial State and Reward Modeling Table 4 reports the effect of the initial state and reward. Starting from a state of low or high functionality, our framework optimizes performance relative to the baseline. The best performance is obtained when starting from a state of high functionality and training with an absolute reward. Our framework can generate novel sequences even when beginning with a high functionality value. Nonetheless, beginning from a state with low functionality, the trained policy does not produce designs with high functionality. This could mean that there may be a gap between regions with low functionality and regions with high functionality in the representation space that requires additional time steps to be explored in each episode.

Initial state	Reward	Performance	Novelty	Diversity
Low	Dense	1.856 ± 0.631	70%	8.032
	Absolute ($r_{th} = 1.8$)	1.855 ± 0.601	50%	5.653
	Binary	1.546 ± 0.271	100%	6.347
	Baseline	1.492 ± 0.008	-	5.426
High	Dense	3.448 ± 0.094	100%	5.700
	Absolute ($r_{th} = 3.3$)	3.531 ± 0.06	100%	6.311
	Binary	3.452 ± 0.081	95%	5.279
	Baseline	3.370 ± 0.013	-	4.858

Table 4: Ablation studies investigating the sampling of initial states and the reward modeling for the GFP dataset. The policy trained with absolute reward and starting from high initial states is used for comparisons in Tables 1 and 3.

Representation Analysis We analyzed the influence of the optimization method on the latent vector as an input. Our method, which employs reinforcement learning for taking actions that update a latent vector, is compared to random perturbation and BO, both of which take a latent vector as input. First, BO performs better when using the latent vector input (2.266 ± 0.899) than when using the sequence as input (1.231 ± 0.034). It implies that the latent vector provides rich information in a low-dimensional space that is easier to optimize than the sequence space. Even with the same representation, our method outperforms BO while achieving competitive novelty.

Model	Performance	Novelty	Diversity	Chromophore
Ours	**3.531** ± 0.06	100%	6.311	100%
Baseline	3.370 ± 0.013	-	4.858	100%
BO on latent space	2.266 ± 0.899	90%	6.600	100%
Random perturbation	1.973 ± 0.832	100%	14.616	100%

Table 5: Representation analysis comparing our method and BO when using latent representation as an input.

Training with Negative Examples We trained the GFP functionality predictor with the original training dataset and the training dataset augmented with 40,000 negative examples (i.e., 82% of the size of the original training dataset). The model architecture and training are detailed in Appendix A.2. We tested both predictors with 10,000 negative examples that were not seen during training. Fig. 2 shows that the functionality predictor trained without negative examples incorrectly predicts a high value for non-functional sequences (mean=4.002), whereas the one trained with negative examples accurately assigns zero functionality (mean=0.0) to random non-functional sequences.

5 Discussion

How the trained policy traverses the functionality landscape We qualitatively evaluate our trained policy by analyzing its ability to traverse the local landscape. Figs. 3 (a) and (b) visualize the optimization trajectory during one episode of GFP, indicating continued improvement of the fluorescence level. This demonstrates that the model can successfully *climb up* the local landscape throughout the episode. The policy can also traverse local optima thanks to the MDP formulation (See Appendix A.7). For His3, the cell fitness landscape of the entire dataset is shown in Fig. 3(c), which interestingly depicts a valley region that corresponds to part of the space leading to low fitness. Also, we can observe that the optimized states are scattered throughout the landscape, implying that our policy can design diverse proteins. Fig. 3(d) presents optimization steps taken by the trained policy, which now shows that large (optimistic) perturbations taken by the policy often lead to failure in optimization.

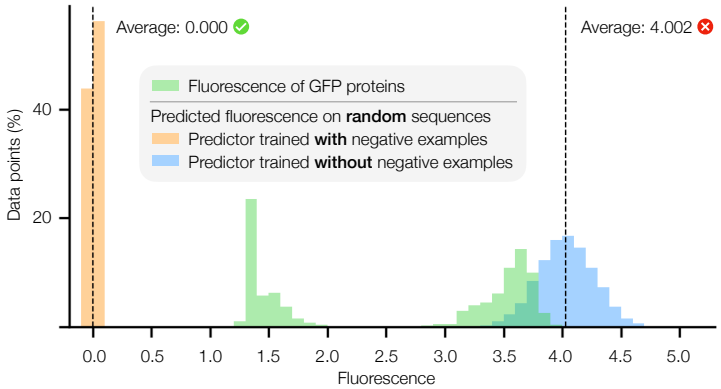

Figure 2: The predicted fluorescence of random sequences tested using a functionality predictor trained **with** or **without** negative examples.

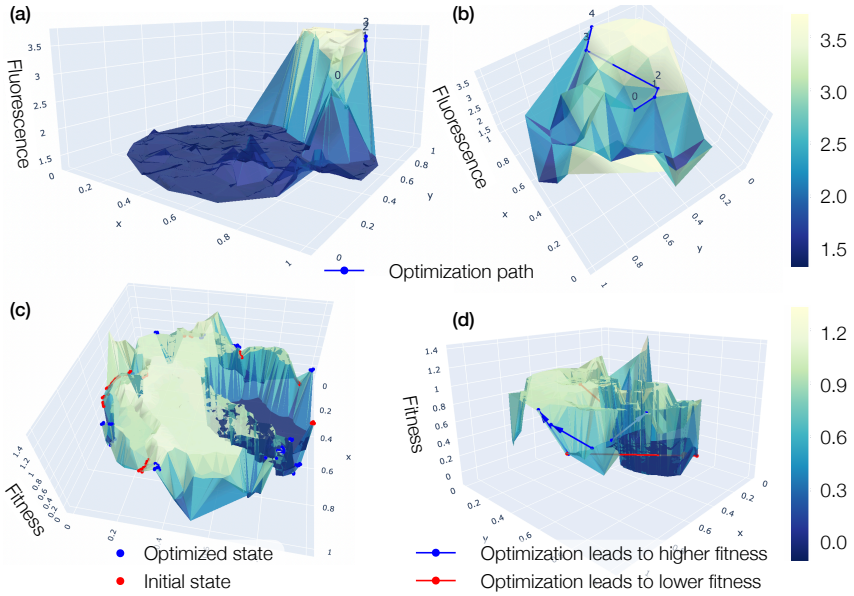

Figure 3: Visualization of the climbing process through a local functionality and fitness landscape. **(a) and (b)** Resulting optimization carried out by the trained policy during a single episode for the GFP task. The z-axis represents the intensity of log fluorescence. In (b) it is seen the policy traversing local optima due to the RL-based formulation of the proposed method. **(c)** Fitness landscape for the His3 dataset. We annotated initial states (red) and optimized states (blue). **(d)** Optimization performed by the trained policy for the His3 task. The representation space is reduced to two using t-SNE (Van der Maaten & Hinton, 2008). Appendix A.6 describes the methodology used to generate these landscapes. An HTML-formatted interactive version of these graphs can be found in the supplementary materials.

AlphaFold2 predictions of optimized sequences Fig. 4 visualizes the structure of the optimized sequence of GFP based on AlphaFold2 (Jumper et al., 2021). The sequences produced by our method and CbAS maintain the critical chromophore region that is known for emitting fluorescence in GFP and the beta sheets that secure the chromophore region. In contrast, DynaPPO-optimized sequences failed to preserve these essential structures of GFP. The different structures obtained by two reinforcement learning methods highlight the significance of state and action modeling in the design of biological sequences. Learning from a single-hot sequence encoding makes it difficult for an algorithm to identify crucial structural information. On the other hand, the latent vector trained by extracting information from a language model trained with millions of protein sequences can efficiently learn mutational effects and reflect protein structures, as shown in (Rives et al., 2021).

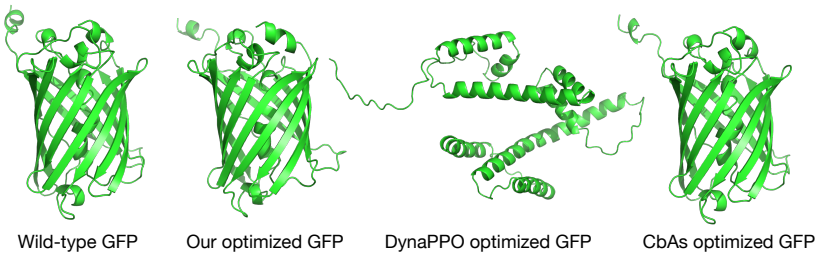

Figure 4: Structures predicted by AlphaFold2 for sequences optimized by the proposed method, DynaPPO, and CbAs for GFP.

How to design *de novo* functional proteins outside the experimental data distribution In protein sequence design, a policy should also be able to design *de novo* functional proteins outside the experimental data distribution. This problem is still an open question. A Markov chain Monte Carlo-based method named Hallucination (Anishchenko et al., 2021) was able to effectively design *de novo* proteins optimized to resemble a general (background) protein distribution. Combining background distribution loss and predictions from a functionality predictor could be an exciting step in the design of *de novo* proteins.

Nevertheless, if the functionality predictor is unreliable outside of the data distribution, the resulting de novo designs will likely fail. Consequently, data-driven methods should be able to make accurate predictions and identify which portions (subgraphs) of a protein structure are responsible for the target function. Using hallucination and modern in-painting (Lugmayr et al., 2022) techniques, these subgraphs may be used to generate *de novo* proteins (Wang et al., 2021). In-painting methods based on diffusion models have shown promising results in the design of proteins in (Anand & Achim, 2022). Combining proteins generated by hallucination and in-painting methods with our optimization framework is a promising approach.

6 Conclusion

This paper addressed the problem of optimizing protein functionality and cellular fitness using data-driven methods. We investigated the limitations of the latest model-based biological sequence design methods and proposed a novel optimization framework that can efficiently traverse the latent representation space, as opposed to optimizing the protein sequence by mutations in its amino acid sequence. Our framework trained a policy that continually updates the latent representation throughout an episode. We modeled the problem as an MDP to maximize future rewards, which allowed the agent to explore the landscape while learning an optimal policy. This special mechanism was critical to efficiently traversing the local functionality and fitness landscape. The proposed framework outperformed other baseline methods in performance and is competitive in terms of novelty. We have also demonstrated that a trained policy can navigate a local landscape effectively and traverse local optima. In the future, we intend to filter the generated sequences and conduct wet lab experiments to validate the proposed framework.

REFERENCES

Ethan C Alley, Grigory Khimulya, Surojit Biswas, Mohammed AlQuraishi, and George M Church. Unified rational protein engineering with sequence-based deep representation learning. *Nature methods*, 16(12):1315–1322, 2019.

Namrata Anand and Tudor Achim. Protein structure and sequence generation with equivariant denoising diffusion probabilistic models. *arXiv preprint arXiv:2205.15019*, 2022.

Christof Angermueller, David Dohan, David Belanger, Ramya Deshpande, Kevin Murphy, and Lucy Colwell. Model-based reinforcement learning for biological sequence design. In *International conference on learning representations*, 2019.

Ivan Anishchenko, Samuel J Pellock, Tamuka M Chidyausiku, Theresa A Ramelot, Sergey Ovchinnikov, Jingzhou Hao, Khushboo Bafna, Christoffer Norn, Alex Kang, Asim K Bera, et al. De novo protein design by deep network hallucination. *Nature*, 600(7889):547–552, 2021.

David Belanger, Suhani Vora, Zelda Mariet, Ramya Deshpande, David Dohan, Christof Angermueller, Kevin Murphy, Olivier Chapelle, and Lucy Colwell. Biological sequences design using batched bayesian optimization. 2019.

Nadav Brandes, Dan Ofer, Yam Peleg, Nadav Rappoport, and Michal Linial. Proteinbert: A universal deep-learning model of protein sequence and function. *Bioinformatics*, 38(8):2102–2110, 2022.

David Brookes, Hahnbeom Park, and Jennifer Listgarten. Conditioning by adaptive sampling for robust design. In *International conference on machine learning*, pp. 773–782. PMLR, 2019.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

Raphael R Eguchi, Christian A Choe, and Po-Ssu Huang. Ig-vae: Generative modeling of protein structure by direct 3d coordinate generation. *PLoS computational biology*, 18(6):e1010271, 2022.

Noelia Ferruz, Steffen Schmidt, and Birte Höcker. Protgpt2 is a deep unsupervised language model for protein design. *Nature communications*, 13(1):1–10, 2022.

Karl Pearson F.R.S. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901. doi: 10.1080/14786440109462720.

Vladimir Gligorijević, P Douglas Renfrew, Tomasz Kosciolek, Julia Koehler Leman, Daniel Berenberg, Tommi Vatanen, Chris Chandler, Bryn C Taylor, Ian M Fisk, Hera Vlamakis, et al. Structure-based protein function prediction using graph convolutional networks. *Nature communications*, 12(1):1–14, 2021.

Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.

Samuel C Hoffman, Vijil Chenthamarakshan, Kahini Wadhawan, Pin-Yu Chen, and Payel Das. Optimizing molecules using efficient queries from property evaluations. *Nature Machine Intelligence*, 4(1):21–31, 2022.

Po-Ssu Huang, Scott E Boyken, and David Baker. The coming of age of de novo protein design. *Nature*, 537(7620):320–327, 2016.

Moksh Jain, Emmanuel Bengio, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Bonaventure FP Dossou, Chanakya Ajit Ekbote, Jie Fu, Tianyu Zhang, Michael Kilgour, Dinghuai Zhang, et al. Biological sequence design with gflownets. In *International Conference on Machine Learning*, pp. 9786–9801. PMLR, 2022.

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Aviral Kumar and Sergey Levine. Model inversion networks for model-based optimization. *Advances in Neural Information Processing Systems*, 33:5126–5137, 2020.

Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Allan dos Santos Costa, Maryam Fazel-Zarandi, Tom Sercu, Sal Candido, et al. Language models of protein sequences at the scale of evolution enable accurate structure prediction. *bioRxiv*, 2022.

Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11461–11471, 2022.

Joshua Meier, Roshan Rao, Robert Verkuil, Jason Liu, Tom Sercu, and Alex Rives. Language models enable zero-shot prediction of the effects of mutations on protein function. *Advances in Neural Information Processing Systems*, 34:29287–29303, 2021.

Igor Melnyk, Payel Das, Vijil Chenthamarakshan, and Aurelie Lozano. Benchmarking deep generative models for diverse antibody sequence design. *arXiv preprint arXiv:2111.06801*, 2021.

Victoria O Pokusaeva, Dinara R Usmanova, Ekaterina V Putintseva, Lorena Espinar, Karen S Sarkisyan, Alexander S Mishin, Natalya S Bogatyreva, Dmitry N Ivankov, Arseniy V Akopyan, Sergey Ya Avvakumov, et al. An experimental assay of the interactions of amino acids from orthologous sequences shaping a complex fitness landscape. *PLoS genetics*, 15(4):e1008079, 2019.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

Roshan Rao, Nicholas Bhattacharya, Neil Thomas, Yan Duan, Peter Chen, John Canny, Pieter Abbeel, and Yun Song. Evaluating protein transfer learning with tape. *Advances in neural information processing systems*, 32, 2019.

Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Jason Liu, Demi Guo, Myle Ott, C Lawrence Zitnick, Jerry Ma, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15):e2016239118, 2021.

Karen S Sarkisyan, Dmitry A Bolotin, Margarita V Meer, Dinara R Usmanova, Alexander S Mishin, George V Sharonov, Dmitry N Ivankov, Nina G Bozhanova, Mikhail S Baranov, Onuralp Soylemez, et al. Local fitness landscape of the green fluorescent protein. *Nature*, 533(7603):397–401, 2016.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Sam Sinai, Richard Wang, Alexander Whatley, Stewart Slocum, Elina Locane, and Eric D Kelsic. Adalead: A simple and robust adaptive greedy search algorithm for sequence design. *arXiv preprint arXiv:2010.02141*, 2020.

Baris E Suzek, Yuqi Wang, Hongzhan Huang, Peter B McGarvey, Cathy H Wu, and UniProt Consortium. Uniref clusters: a comprehensive and scalable alternative for improving sequence similarity searches. *Bioinformatics*, 31(6):926–932, 2015.

Kevin Swersky, Yulia Rubanova, David Dohan, and Kevin Murphy. Amortized bayesian optimization over discrete spaces. In *Conference on Uncertainty in Artificial Intelligence*, pp. 769–778. PMLR, 2020.

Kei Terayama, Masato Sumita, Ryo Tamura, and Koji Tsuda. Black-box optimization for automated discovery. *Accounts of Chemical Research*, 54(6):1334–1346, 2021.

Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Jue Wang, Sidney Lisanza, David Juergens, Doug Tischer, Ivan Anishchenko, Minkyung Baek, Joseph L Watson, Jung Ho Chun, Lukas F Milles, Justas Dauparas, et al. Deep learning methods for designing proteins scaffolding functional sites. *bioRxiv*, 2021.

Zichen Wang, Steven A Combs, Ryan Brand, Miguel Romero Calvo, Panpan Xu, George Price, Nataliya Golovach, Emmanuel O Salawu, Colby J Wise, Sri Priya Ponnapalli, et al. Lm-gvp: an extensible sequence and structure informed deep learning framework for protein property prediction. *Scientific reports*, 12(1):1–12, 2022.

Jian Wu, Matthias Poloczek, Andrew G Wilson, and Peter Frazier. Bayesian optimization with gradients. *Advances in neural information processing systems*, 30, 2017.

Kevin K Yang, Zachary Wu, and Frances H Arnold. Machine-learning-guided directed evolution for protein engineering. *Nature methods*, 16(8):687–694, 2019.

Zuobai Zhang, Minghao Xu, Arian Jamasb, Vijil Chenthamarakshan, Aurelie Lozano, Payel Das, and Jian Tang. Protein representation learning by geometric structure pretraining. *arXiv preprint arXiv:2203.06125*, 2022.

A APPENDIX

A.1 SEQUENCE ENCODER-DECODER

The sequence encoder comprises a token embedding layer and 33 transformer layers. Each transformer layer consists of multi-head attention followed by a layer normalization layer and two fully connected layers with GeLU (Hendrycks & Gimpel, 2016) non-linear activation. The number of attention heads is set to 20. We used the pre-trained weights provided in (Lin et al., 2022). The encoder weights are not updated during the fine-tuning process. We jointly fine-tuned the encoder and decoder to the train set of each dataset for 16 epochs using the Adam optimizer (Kingma & Ba, 2014). The initial learning rate of Adam is set to 8e-5, with weight decay set to 1e-5. We defined the loss function as the mean of cross entropy loss on mutated positions and the mean cross-entropy loss on non-mutated positions to ensure that the decoder also focuses on recovering mutated positions. Since we provide the wild-type representation during sequence recovery and mutants in the GFP dataset have a maximum of 15 mutations, such a loss function is important to prevent the decoder from predicting every sequence as the wild-type. The learning rate is reduced by $0.8\times$ every epoch. Table 6 shows the performance of a sequence decoder by the chosen embedding size of R. Note that the information on mutated positions such as masks is not provided to the model. Therefore, the recovery of non-mutated positions is non-trivial even though the wild-type representation is given.

Dataset	Embedding size	Top-1 Accuracy	
		Mutated positions	Non-mutated positions
GFP	8	0.463	0.9958
	32	0.4762	0.9903
His3	8	0.6439	0.8407
	32	0.8194	0.9022

Table 6: Performance of the sequence decoder for GFP and His3.

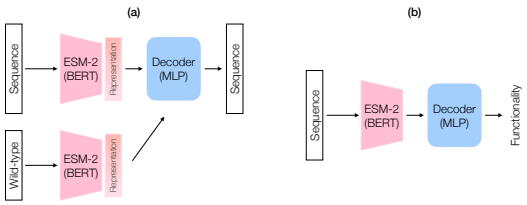

Figure 5: Model Architecture

A.2 FUNCTIONALITY PREDICTOR

Functionality and fitness predictors for GFP and His3 are also based on the pre-trained ESM-2 protein language model. We replaced the sequence decoder with three fully-connected layers with GeLU (Hendrycks & Gimpel, 2016) non-linear activation that predicts protein functionality and fitness from the representation of size (L, E). We fine-tuned the predictor on the train set of each dataset for 16 epochs using the Adam (Kingma & Ba, 2014) optimizer with the initial learning rate set to 8e-5 and weight decay set to 1e-5. The encoder weights are not updated during the fine-tuning process. We used the mean squared error between empirically obtained log-fluorescence intensity and predicted fluorescence as the loss function. In Table 7 we report the test Spearman's ρ and mean squared error (MSE) of the functionality predictor used for model-based optimization, *i.e.*, oracle.

A.2.1 TRAINING CBAS AND DYNAPPO WITH AND WITHOUT NEGATIVE SAMPLES

We conducted an ablation study by training two baseline optimization methods, CbAS and DynaPPO, using both functionality predictors. The optimized sequences are then visualized using predicted structures obtained by AlphaFold2. As shown in Fig. 6, CbAS managed to propose a valid fluorescent protein when guided by an oracle trained with negative examples but failed when

Task	Test set		Negative test set
	Spearman's ρ	MSE	MSE
GFP	0.8426	0.1436	2.930e-6
His3	0.6635	0.0080	0.0164

Table 7: Performance of the functionality predictor for GFP and His3

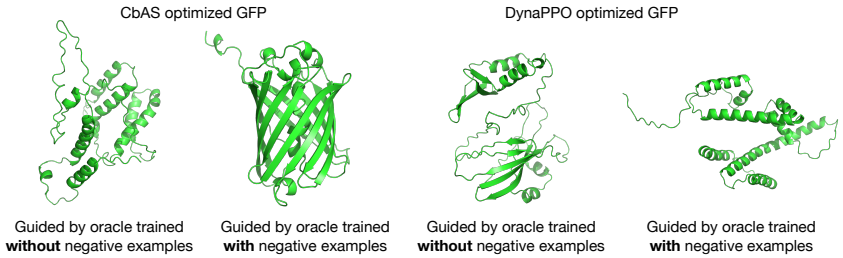

Figure 6: The structure of optimized sequences without and with negative examples.

guided by an oracle trained without negative examples. Even though the structure predicted for the CbAS model for this case seems unstable, the oracle trained without negative examples assigned a very high log-fluorescence intensity value (4.014) for the sequence. Therefore, we argue that biological sequence design should be guided and evaluated by an oracle trained with an experimental dataset augmented with negative examples.

A.3 Soft Actor Critic (SAC)

Recapitulating, an optimal policy π^* following the SAC objective is given as

$$\pi^* = arg\max_{\pi} \mathbb{E}_{\pi}\left[\sum_{t=0}^{\infty} \gamma^t (r(s_t, a_t) + \alpha H(\pi(\cdot|s_t)))\right], \tag{3}$$

where $\gamma \in [0, 1]$ is a discount factor, α is a temperature hyperparameter and H is entropy. Following a policy π the Q function or action-value function is defined as

$$Q^{\pi}(s, a) = \mathbb{E}_{\pi}\left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) + \alpha \sum_{t=1}^{\infty} \gamma^t H(\pi(\cdot|s_t))\right], \tag{4}$$

where the entropy is added from the first timestep. The Bellman equation for Q^{π} at a timestep t is given by

$$Q^{\pi}(s_t, a_t) = \mathbb{E}_{\pi}\left[r(s_t, a_t) + \gamma(Q^{\pi}(s_{t+1}, \tilde{a}_{t+1}) + \alpha H(\pi(\cdot|s_{t+1}))\right], \tag{5}$$

in which he state s_{t+1} is used from the replay buffer and $\tilde{a}_{t+1}$ is sampled from the current policy. During training, SAC models three networks: a policy π_θ and two Q-functions $Q_{\phi 1}$ and $Q_{\phi 2}$. The two Q functions are due to the clipped double-Q trick, more details can be found in Haarnoja et al. (2018). The loss functions to train $Q_{\phi 1}$ and $Q_{\phi 2}$ in SAC are defined as

$$L(\phi_i, \mathcal{D}) = \mathbb{E}_{(s_t, a_t, r_t, s_{t+1}, d_t)\sim\mathcal{D}}\left[(Q_{\phi_i}(s_t, a_t) - y(r_t, s_{t+1}, d_t))^2\right], \tag{6}$$

where the target $y(r_t, s_{t+1}, d_t)$, substituting the entropy function, is given by

$$y(r_t, s_{t+1}, d_t) = r_t + \gamma(1 - d_t)\left(\min_{j=1,2} Q_{\phi_{targ,j}}(s_{t+1}, \tilde{a}_{t+1}) - \alpha \log \pi_\theta(\tilde{a}_{t+1}|s_{t+1})\right), \tag{7}$$

where $\tilde{a}_{t+1} \sim \pi_\theta(\cdot|s_{t+1})$. The policy π_θ is trained to maximize the value function of state s_t. The value function V^{π} is as follows

$$V^{\pi}(s_t) = \mathbb{E}_{a_t\sim\pi}\left[Q^{\pi}(s_t, a_t) - \alpha \log \pi(a_t|s_t)\right]. \tag{8}$$

The policy is reparameterized and then optimized following Haarnoja et al. (2018). Using a squashing Gaussian function, sampling the action is finally defined as

$$\tilde{a}_{t_\theta}(s_t, \xi) = tanh(\mu_\theta(s_t) + \sigma_\theta(s_t) \odot \xi), \qquad \xi \sim \mathcal{N}(0, I). \tag{9}$$

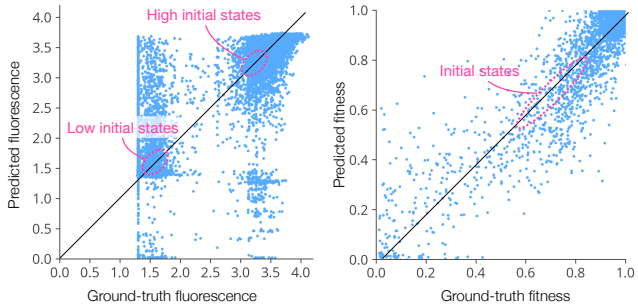

Figure 7: Ground truth and predicted fluorescence of GFP and fitness of His3.

A.4 Implementation Details

A.4.1 Policy

Soft Actor-Critic (SAC) We used three fully connected layers with a hidden size set to 256 for both the actor and critic networks of SAC. The Adam optimizer Kingma & Ba (2014) is used for both actor and critic networks. The entropy temperature α is automatically tuned during training following Haarnoja et al. (2018).

Proximal Policy Optimization (PPO) For the ablation studies in Table 3 using the mutation in the protein sequence as output, we use Proximal Policy Optimization (PPO) Schulman et al. (2017) to handle the discrete action space setting. The PPO architecture is similar to the one used for the actor of SAC, using three fully connected layers with a hidden size set to 256. We define the action as a multiple of discrete actions for the third and fifth row in Table 3.

A.4.2 Random Perturbation Baseline

We added a random value between $[-1, 1]$ to the initial state representation to test the effect of random perturbations on performance. Then, we decoded the sequence using the decoder presented in Sec. 3.1 and predicted the functionality using the functionality predictor presented in Sec. 3.2. Note that the baseline and random performance reported in the Tables 1 and 2 are also calculated on the top K candidates.

A.4.3 Previous Methods

For the comparison with previous baseline methods, we use the FLEXS implementation (Sinai et al., 2020) of DynaPPO, CbAS, and BO. All the models are trained for ten rounds. During training, we observed the performance saturation after the first few rounds. The sequences proposed in the last round are used to evaluate the baseline models. For DynaPPO, its ensemble model is trained for ten rounds, and the environment batch size is set to 10. For CbAS, a CNN architecture with a hidden size equal to 100 and a number of channels equal to 32 is used. Also, the generator block of CbAS is trained for ten epochs.

A.5 Initial States of the Optimization Process

We selected initial states for the optimization process in our framework according to two criteria: (i) sequences with room to improve, (ii) sequences sampled from regions that we can trust the oracle. Since the distribution for the ground truth functionality of GFP is bimodal, we sampled two sets of initial states from each mode, low initial states, and high initial states. For His3, only one set of initial states is chosen. As shown in Fig. 7, we sample the sequences from pre-determined regions that do not include the most functional mutants. Additionally, the sequences are sampled from examples in which the difference between their ground truth and predicted functionality or fitness is less than a threshold set to 0.1.

Figure	Task	Number of Representations	Number of Sequences	Dimension reduction
Fig. 3 (a)	GFP	1.55M	1403	t-SNE
Fig. 3 (b)	GFP	1.17M	46	t-SNE
Fig. 3 (c)(d)	His3	6.43K	6434	t-SNE
Fig. 8	GFP	1.17M	46	PCA

Table 8: Details of the complete functionality landscape.

A.6 METHODOLOGY TO OBTAIN THE COMPLETE FUNCTIONALITY LANDSCAPE

The methodology to plot the Fig. 3 and Fig. 8 is detailed next. First, a range based on the representations obtained during the episode being plotted is defined. Then we decode the distinct sequences from the large number of representations using the decoder presented in Sec. 3.1. Table 8 shows the number of representations and sequences used to plot the complete functionality and fitness landscape. After this step, the reward is decoded for each sequence using the functionality predictor presented in Sec. 3.2. To reduce the dimensionality, t-SNE (Van der Maaten & Hinton, 2008) is applied in Fig. 3 and Principal Component Analysis (PCA) (F.R.S., 1901) is applied in Fig 8 with the two principal components kept to create the local landscape.

A.7 POLICY ABILITY TO TRAVERSE LOCAL OPTIMA

We qualitatively evaluate our trained policy for the GFP by analyzing its ability to traverse local optima of the functionality landscape. The optimization process during one episode is shown in Fig. 8. Until the end of the episode, at timestep $t = 10$, the policy can traverse the landscape and maximize functionality efficiently. At timestep $t = 3$, it can escape local optima that do not trigger the conditions for the end of an episode. During timesteps $t = 7, 8, 9$, it can be seen that the policy is still maximizing the functionality metric until the end of the episode. Additional details and the methodology used to obtain the functionality landscape in Fig. 8 are explained in A.6.

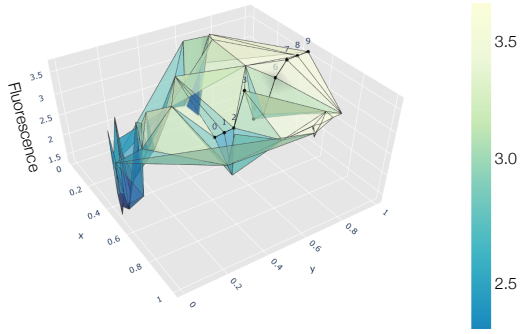

Figure 8: Optimization process performed by the trained policy during one episode for the GFP task. The x-axis and y-axis are the two principal components of the representation space calculated using Principal Component Analysis (F.R.S., 1901). The z-axis is the log-fluorescence intensity.