

Importance Weighted Expectation-Maximization for Protein Sequence Design

Zhenqiao Song¹ Lei Li¹


Abstract

Designing protein sequences with desired biological function is crucial in biology and chemistry. Recent machine learning methods use a surrogate sequence-function model to replace the expensive wet-lab validation. How can we efficiently generate diverse and novel protein sequences with high fitness? In this paper, we propose IsEM-Pro, an approach to generate protein sequences towards a given fitness criterion. At its core, IsEM-Pro is a latent generative model, augmented by combinatorial structure features from a separately learned Markov random fields (MRFs). We develop an Monte Carlo Expectation-Maximization method (MCEM) to learn the model. During inference, sampling from its latent space enhances diversity while its MRFs features guide the exploration in high fitness regions. Experiments on eight protein sequence design tasks show that our IsEM-Pro outperforms the previous best methods by at least 55% on average fitness score and generates more diverse and novel protein sequences. The code is available at <https://github.com/JocelynSong/IsEM-Pro.git>

1. Introduction

Protein engineering aims to discover protein variants with desired biological function, such as fluorescence intensity (Biswas et al., 2021), enzyme activity (Fox et al., 2007), and therapeutic efficacy (Lagassé et al., 2017). Protein sequences embody their function through spontaneous folding of amino-acid sequences into three dimensional structures (Go, 1983; Chothia, 1984; Starr & Thornton, 2017). The mapping from protein sequence to functional property forms a protein fitness landscape that characterizes the protein functional levels, such as the capability to catalyze reaction

¹Department of Computer Science, University of California, Santa Barbara, California, the United States. Correspondence to: Zhenqiao Song <zhenqiao@ucsb.edu>, Lei Li <leili@cs.ucsb.edu>.



(a) Fujiyama landscape. (b) Badlands landscape.

Figure 1. Protein Fitness Landscape: distribution of a functional property for proteins). Protein may exhibit single-peaked fitness landscape (Fujiyama landscape (a)) or multi-peaked landscape (Badlands landscape (b)) (Kauffman & Weinberger, 1989). In Fujiyama landscape, any method could perform well. However, for the rougher Badlands landscape, previous methods get trapped in a worse local optima while our proposed IsEM-Pro can climb much closer to the global optima through the iterative sampling in the latent space.

or bind a specific ligand (Romero & Arnold, 2009; Ren et al., 2022). Traditional approaches to design a specific protein with desired fitness objective involve obtaining protein variants by random mutagenesis (Labrou, 2010) or recombination in laboratory experiments (Ma et al., 2003). These variants are screened and selected in wet-lab experiments (Arnold, 1998) as illustrated in Figure 2.

However, these approaches requires iterative cycles of random mutagenesis and wet-lab validation, which are both money-consuming and time-intensive. Recent machine learning methods attempt to build a surrogate model of protein fitness landscape to accelerate expensive wet-lab screening (Luo et al., 2021; Meier et al., 2021). How can we efficiently discover satisfactory proteins over the exponentially large discrete space? Ideal protein molecules should be novel, diverse and exhibit high fitness. On one hand, designing novel and diverse protein sequences can uncover new functions and also lead to functional diversification (Singh et al., 2016). On the other hand, in addition to the evolutionary pressure for a protein to conserve specific positions of its sequence for function, the diversification of protein sequences can avoid undesired inter-domain association such as misfolding (Wright et al., 2005).

In this paper, we propose an Importance sampling based




Figure 2. Workflow of traditional protein sequence design. We aim to accelerate this process by directly generating desirable sequences.

Expectation-Maximization (EM) method to efficiently design novel, diverse and desirable **Protein** sequences (IsEM-Pro). Specifically, we introduce a latent variable in the generative model to capture the inter-dependencies in protein sequences. Sampling in latent space leads to more diverse candidates and can escape from locally optimal fitness regions. Instead of using standard variational inference models such as variational auto-encoder (VAE) (Kingma & Welling, 2014), we leverage importance sampling inside the EM algorithm to learn the latent generative model. As illustrated in Figure 1, our approach can navigate through multiple local optimum, and yield better overall performance. We further incorporate combinatorial structure of amino acids in protein sequences using Markov random fields (MRFs). It guides the model towards higher fitness landscape, leading to faster uphill path to desired proteins.

We carry out extensive experiments on eight protein sequence design tasks and compare the proposed method with previous strong baselines. The contribution are as follows:

- We propose a structure-enhanced latent generative model for protein sequence design.
- We develop an efficient method to learn the proposed generative model, based on importance sampling inside the EM algorithm.
- Experiments on eight protein datasets with different objectives demonstrate that our IsEM-Pro generates protein sequences with at least 55% higher average fitness score and higher diversity and novelty than previous best methods. Further analyse show that the protein sequences designed by our model can fold stably, giving empirical evidence that our proposed IsEM-Pro has the ability to generate real proteins.

2. Background

2.1. Protein Sequence Design upon Wild-Type

The protein sequence design problem is to search for a sequence with desired property in the sequence space \mathcal{V}^L , where \mathcal{V} denotes the vocabulary of amino acids and L denotes the desired sequence length. The target is to find a protein sequence with highest fitness given by a protein fitness function $f : \mathcal{V}^L \rightarrow \mathbb{R}$, which can be measured through wet-lab experiments. Wild-type refers to protein occurring

in nature. Evolutionary search based methods are widely used (Bloom & Arnold, 2009; Arnold, 2018; Angermüller et al., 2020; Ren et al., 2022). They use wild-type sequence as starting point during iterative search. In this paper, we do not focus on the modification upon the wild-type sequence, but aim to efficiently generate novel and diverse sequences with improved protein functional properties.

2.2. Monte Carlo Expectation-Maximization

We will develop our method based on Monte Carlo expectation-maximization (MCEM) (Bishop, 2006). A latent generative model assumes data x (e.g. a protein sequence) is generated from a latent variable z . To learn this latent model, the optimization procedure for maximizing the log marginal likelihood is to alternate between expectation step (E-step) and maximization step (M-step). EM directly targets the log marginal likelihood of an observation x by involving a variational distribution $q_\phi(z)$:

$$\log p_\theta(x) = E_{q_\phi(z)}[\log p_\theta(x, z) - \log q_\phi(z)] + D_{KL}(q_\phi(z) || p_\theta(z|x)) \quad (1)$$

where $p_\theta(z|x)$ is the true posterior distribution and $p_\theta(x, z) = p_\theta(x|z)p_\theta(z)$ is the joint distribution, composed of the conditional likelihood $p_\theta(x|z)$ and the prior $p_\theta(z)$. In MCEM, E-step samples a set of z from $q_\phi(z)$ to estimate the expectation using Monte Carlo method, and then M-step fits the model parameters θ by maximizing the Monte Carlo estimation (Wei & Tanner, 1990). It can be proved that this process will never decrease the log marginal likelihood (Bishop, 2006).

3. Proposed Method: IsEM-Pro

In this section, we describe our method in detail. We will first present the probabilistic model and its learning algorithm. To make the learning more efficient, we describe how to uncover and use the potential constraints conveyed in the protein sequences.

3.1. Problem Formulation

Our goal is to search over a space of discrete protein sequences \mathcal{V}^L – \mathcal{V} consists of 20 amino acids and L is the sequence length – for sequence $x \in \mathcal{V}^L$ that maximizes a

given fitness function $f : \mathcal{V}^L \rightarrow \mathbb{R}$. Let the fitness value $y = f(x)$, given a predefined threshold λ , we define a conditional likelihood function:

$$P(\mathcal{S}|x) = \begin{cases} 1, & f(x) \geq \lambda, \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

where \mathcal{S} represents the event that the fitness of x is ideal ($y \geq \lambda$). We also assume a class of generative models $P_\theta(x)$ that can be trained to model the raw protein sequences and it will be kept fixed afterwards. Since the search space is exponentially large ($O(20^L)$), random search would be time-intensive. Following Brookes et al. (2019), we formulate the protein design problem as generating satisfactory sequences from the posterior distribution $P_\theta(x|\mathcal{S})$:

$$P_\theta(x|\mathcal{S}) = \frac{P_\theta(x)P(\mathcal{S}|x)}{P_\theta(\mathcal{S})} \quad (3)$$

where $P_\theta(\mathcal{S}) = \int_x P_\theta(x)P(\mathcal{S}|x)dx$ is a normalization constant which does not rely on x . Protein sequences generated from $P_\theta(x|\mathcal{S})$ are not only more likely to be real proteins, but also have higher functional scores (i.e., fitness). The higher the λ is, the higher fitness the discovered protein sequences have.

3.2. Probabilistic Model

Directly generating satisfactory sequences from the posterior distribution $P_\theta(x|\mathcal{S})$ is highly efficient compared with randomly search over the exponentially discrete space. However, realizing this idea is difficult as computing $P_\theta(\mathcal{S}) = \int_x P(\mathcal{S}|x)P_\theta(x)dx$ needs an integration over all possible x , which is intractable. Instead, we propose to learn a variational distribution $Q_\phi(x)$ with learnable parameter ϕ to approximate satisfactory proteins $P_\theta(x|\mathcal{S})$. Following Brookes et al. (2019), to find the optimal ϕ of the proposal distribution, we choose to minimize the KL divergence between the posterior distribution $P_\theta(x|\mathcal{S})$ and the variational distribution $Q_\phi(x)$:

$$\begin{aligned} \phi^* &= \underset{\phi}{\operatorname{argmax}} -D_{KL}(P_\theta(x|\mathcal{S})||Q_\phi(x)) \\ &= \underset{\phi}{\operatorname{argmax}} E_{P_\theta(x|\mathcal{S})} \log Q_\phi(x) + \mathcal{H}(\mathcal{P}_\theta) \end{aligned} \quad (4)$$

where $\mathcal{H}(\mathcal{P}_\theta) = -E_{P_\theta(x|\mathcal{S})} \log P_\theta(x|\mathcal{S})$ is the entropy of $P_\theta(x|\mathcal{S})$ and can be dropped because it does not matter ϕ .

Diversity is a key consideration in our protein design procedure, which not only satisfies the diverse nature of species, but also can reduce undesired inter-domain misfolding (Wright et al., 2005). In order to promote the diversity of the designed protein sequences, we introduce a latent variable z into our model to capture the high-order dependencies among amino acids in protein sequences. We assume the joint approximate distribution $Q_\phi(x, z) = Q_0(z)Q_\phi(x|z)$.

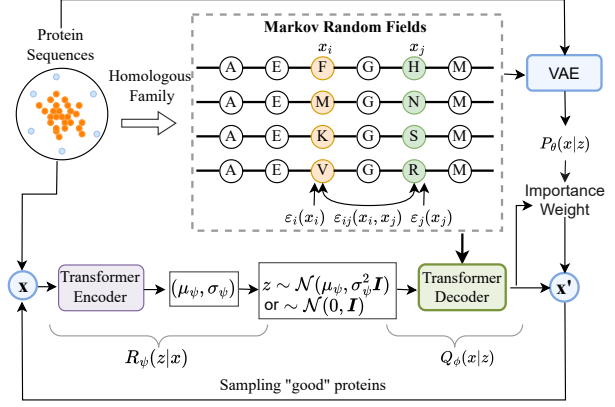


Figure 3. Overall architecture of the proposed IsEM-Pro. The upper half illustrates the Markov random fields which learns the combinatorial structure of amino acids in protein sequences from the same family. The bottom half shows the combinatorial structure feature augmented probabilistic model. Red lines show the calculation of importance weight.

Thus our final goal is to maximize the expected log-likelihood $\log Q_\phi(x)$ of the satisfactory proteins with respect to the ideal posterior distribution $P_\theta(x|\mathcal{S})$. By using the EM objective from Eq.(1), we derive

$$\begin{aligned} \mathcal{L} &= E_{P_\theta(x|\mathcal{S})} \log Q_\phi(x) \\ &= E_{P_\theta(x|\mathcal{S})} \{ E_{R_\psi(z|x)} [\log Q_\phi(x, z) - \log R_\psi(z|x)] \\ &\quad + D_{KL}(R_\psi(z|x)||Q_\phi(z|x)) \} \\ &= E_{P_\theta(x|\mathcal{S})} \mathcal{F}(R_\psi(z|x), \phi) \end{aligned} \quad (5)$$

where $R_\psi(z|x)$ is another variational distribution with its parameter ψ to approximate the intractable posterior $Q_\phi(z|x)$, and $\mathcal{F}(R_\psi(z|x), \phi) = E_{R_\psi(z|x)} [\log Q_\phi(x, z) - \log R_\psi(z|x)] + D_{KL}(R_\psi(z|x)||Q_\phi(z|x))$. Here $Q_0(z)$ is implemented as a standard normal distribution, and $Q_\phi(x|z)$ is a Transformer decoder (Vaswani et al., 2017) with z as the first embedding input, which will be augmented by the combinatorial structure features (Sec. 3.4). We implement $R_\psi(z|x)$ as a normal distribution $\mathcal{N}(\mu_\psi, \sigma_\psi^2 \mathbf{I})$ with $\mu_\psi, \sigma_\psi = \text{Transformer}_\psi(x)$. The overall architecture is illustrated in Figure 3.

3.3. Importance Weighted EM

To maximize the objective defined in Eq.(5), we plan to learn the proposal distribution $Q_\phi(x, z)$ through Monte Carlo EM, because the sampling procedure and iterative optimization process can lead to a better estimate (Figure 1 (b)), resulting in novel and diverse proteins with higher fitness.

Since Eq.(5) involves the expectation over the intractable distribution $P_\theta(x|\mathcal{S})$ and the KL divergence of the intractable posterior distribution $Q_\phi(z|x)$, we use importance sampling to approximate its variational lower bound. We assume the original protein sequence is also generated from

a same latent variable z , i.e. $P_\theta(x, z) = P_\theta(x|z)P_0(z)$, where $P_0(z)$ is a standard normal distribution and $P_\theta(x|z)$ is a Transformer with z as the initial input embedding. We also assume the latent variable z in $P_\theta(x, z)$ and $Q_\phi(x, z)$ are defined on the same latent space. In practice, we learn $P_\theta(x|z)$ as the decoder in a pre-trained variational auto-encoder on raw protein sequences. We assume \mathcal{S} is only conditioned on x , leading to $P_\theta(x, z|\mathcal{S}) = \frac{P_\theta(x, z)P(\mathcal{S}|x)}{P_\theta(\mathcal{S})}$. Then the final training objective can be derived as:

$$\begin{aligned} \mathcal{L} &= E_{P_\theta(x|\mathcal{S})} \mathcal{F}(R_\psi(z'|x), \phi) = E_{P_\theta(x, z|\mathcal{S})} \mathcal{F}(R_\psi(z'|x), \phi) \\ &= E_{Q_\phi(x, z)} \frac{P_\theta(x, z|\mathcal{S})}{Q_\phi(x, z)} \mathcal{F}(R_\psi(z'|x), \phi) \\ &\geq \frac{1}{\mathcal{C}} E_{Q_\phi(x, z)} \frac{P_\theta(x|z)}{Q_\phi(x|z)} P(\mathcal{S}|x) \left(E_{R_\psi(z'|x)} \log Q_\phi(x|z') \right. \\ &\quad \left. - D_{KL}(R_\psi(z'|x) || Q_0(z')) \right) = \tilde{\mathcal{L}} \end{aligned} \quad (6)$$

Here we use the importance sampling fundamental identity to derive the equation (Robert & Casella, 2004). $\mathcal{C} = P_\theta(\mathcal{S})$ is a constant which does not rely on ϕ and ψ . The last inequality adopts the evidence lower bound (ELBO) (Bishop, 2006). We use importance sampling based EM to approximate the above objective (Robert & Casella, 2004) with joint samples $(x_n, z_n) \sim Q_\phi(x, z)$. Specifically, at each iteration, we perform,

E-step:

$$\begin{aligned} \text{Sample } x_n &\sim P_{data}, \quad z_n \sim R_\psi(z|x_n) \\ \text{and } z_n &\sim Q_0(z), \quad x_n \sim Q_\phi(x|z_n), \quad z'_n \sim R_\psi(z|x_n) \\ \tilde{\mathcal{L}}_t &= \sum_{n=1}^N w'(x_n, z_n) \left(\log Q_\phi(x_n|z'_n) \right. \\ &\quad \left. - D_{KL}(R_\psi(z'|x) || Q_0(z')) \right) \end{aligned} \quad (7)$$

where $w(x_n, z_n) = \frac{P_\theta(x_n|z_n)}{Q_\phi(x_n|z_n)} P(\mathcal{S}|x_n)$ is the unnormalized importance weight, $w'(x_n, z_n) = \frac{w(x_n, z_n)}{\sum_{n=1}^N w(x_n, z_n)}$ is the normalized one, and N is the sample size. Since $R_\psi(z|x)$ is assumed to be a normal distribution with its mean μ_ψ and variance $\sigma_\psi^2 \mathbf{I}$ calculated from a Transformer encoder $\mu_\psi, \sigma_\psi = \text{Transformer}_\psi(x_n)$, the KL term can be calculated in closed form (Eq. (17)).

M-step:

$$\begin{aligned} \psi^{(t+1)}, \phi^{(t+1)} &= \underset{\psi, \phi}{\operatorname{argmax}} \tilde{\mathcal{L}}_t \\ \psi^{(t+1)} &= \psi^{(t)} + r * \nabla_\psi \tilde{\mathcal{L}}_t, \quad \phi^{(t+1)} = \phi^{(t)} + r * \nabla_\phi \tilde{\mathcal{L}}_t \end{aligned} \quad (8)$$

In E-step, we use two techniques to generate samples for z and x – using the real protein sequences x with generated z from $R_\psi(z|x)$ and using z from a standard normal distribution with generated proteins from $Q_\phi(x|z)$. In M-step, we optimize the KL regularized data generation log likelihood with both high-fitness real and synthetic proteins. The

procedure can be viewed as self-training with augmented synthetic protein data, which differs from prior approaches such as CbAS (Brookes et al., 2019).

3.4. Guiding Model Climbing through Combinatorial Structure

As shown in previous work, the combinatorial structure of amino acids in protein sequences can be learned from a generative graphical model Markov random fields (MRFs) fitted on the sequences from the same family (Hopf et al., 2017; Luo et al., 2021). These structure constraints are the results of the evolutionary process under natural selection and may reveal clues on which amino-acid combinations are more favorable than others. Thus we incorporate these features into our model to guide it towards higher fitness landscape to faster find desired protein sequences.

Given a protein sequence $x = (x_1, x_2, \dots, x_M)$ with M amino acids, the generative model generates it with likelihood $P_\epsilon(x) = \frac{\exp(\Upsilon(x))}{Z}$ where $Z = \int_x \exp(\Upsilon(x)) dx$ is a normalization constant and $\Upsilon(x)$ is the corresponding energy function, which is defined as the sum of all pairwise constraints and single-site constraints as follows:

$$\Upsilon(x) = \sum_{i=1}^M \varepsilon_i(x_i) + \sum_{i=1}^M \sum_{j=1, j \neq i}^M \varepsilon_{ij}(x_i, x_j) \quad (9)$$

where $\varepsilon_i(x_i)$ denotes the single-site constraint of x_i at position i and $\varepsilon_{ij}(x_i, x_j)$ denotes the pairwise constraint of x_i and x_j at position i, j . The above graphical model is illustrated in the upper half of Figure 3.

We train the model on protein sequences from the same family following CCMpred (Seemayer et al., 2014) using a pseudo-likelihood $\hat{P}_\epsilon(x)$ (provided in Appendix E) combined with L_2 regularization to make the learning of $P_\epsilon(x)$ easier. But different from them, we additionally add L_1 regularization to the training objective to make the graph sparse, of which the regularization coefficients are set to the same values as the L_2 regularization:

$$\begin{aligned} \max_x L_\epsilon &= \sum_x \log \hat{P}_\epsilon(x) - L_1(\epsilon) - L_2(\epsilon) \\ L_1(\epsilon) &= \alpha_{\text{single}} \sum_{i=1}^M \|\varepsilon_i\|_1 + \alpha_{\text{pair}} \sum_{i,j=1, i \neq j}^M \|\varepsilon_{ij}\|_1 \\ L_2(\epsilon) &= \alpha_{\text{single}} \sum_{i=1}^M \|\varepsilon_i\|_2^2 + \alpha_{\text{pair}} \sum_{i,j=1, i \neq j}^M \|\varepsilon_{ij}\|_2^2 \end{aligned} \quad (10)$$

where x denotes protein sequences from the same family, $\varepsilon_i = [\varepsilon_i(a_1), \varepsilon_i(a_2), \dots, \varepsilon_i(a_{20})]$ is the vector of the single-site constraints of the 20 amino acids at position i , and $\varepsilon_{ij} = [\varepsilon_{ij}(a_1, a_2), \varepsilon_{ij}(a_1, a_3), \dots, \varepsilon_{ij}(a_M, a_{M-1})]$ is the vector of all possible pairwise constraints at position i, j .

After training the MRFs, we can encode a protein sequence x with the learned constraints. Specifically, we first encode the i -th amino acid by concatenating its corresponding single-site constraint as well as the possible pairwise ones:

$$\begin{aligned}\varepsilon_i(x_i) &= [\varepsilon_i(x_i), \varepsilon_{i1}(x_i, a_{1\cdot}), \dots, \varepsilon_{iM}(x_i, a_{M\cdot})] \\ \varepsilon_{ij}(x_i, a_{j\cdot}) &= [\varepsilon_{ij}(x_i, a_1), \varepsilon_{ij}(x_i, a_2), \dots, \varepsilon_{ij}(x_i, a_{20})]\end{aligned}\quad (11)$$

where $\varepsilon_{ij}(x_i, a_{j\cdot})$ gathers the 20 amino acids for any position $j \neq i$. Then we map $\varepsilon_i(x_i)$ to the amino-acid embedding space with trainable parameter W_ε , and add the mapped vector to the original amino-acid embedding $e(x_i)$ to get the final feature vector as our model input:

$$\begin{aligned}\hat{e}(x_i) &= e(x_i) + W_\varepsilon * \varepsilon_i(x_i) \\ H_0 &= z', \quad H_i = \hat{e}(x_{i-1}) \text{ for } 1 \leq i < M\end{aligned}\quad (12)$$

where H_i is the input embedding of the Transformer decoder, i.e., the first input is set to the sampled latent vector z' and the input for other position i is set to the combinatorial structure augmented feature vector $\hat{e}(x_{i-1})$.

Combining Eq. (7) and (12), the learning process becomes: **E-step**:

$$\begin{aligned}\tilde{\mathcal{L}}_t &= \sum_{n=1}^N w'(x_n, z_n) \left(\log Q_\phi(x_n | z'_n; \varepsilon) \right. \\ &\quad \left. - \left(\frac{1}{2} \|\sigma_\psi(x_n)\|_2^2 + \frac{1}{2} \|\mu_\psi(x_n)\|_2^2 - \sum_{i=1}^d \log \sigma_{\psi,i}(x_n) \right) \right) \\ w'(x_n, z_n) &= \frac{w(x_n, z_n)}{\sum_{n=1}^N w(x_n, z_n)} \\ w(x_n, z_n) &= \frac{P_\theta(x_n | z_n; \varepsilon)}{Q_{\phi^{(t)}}(x_n | z_n; \varepsilon)} P(\mathcal{S} | x_n)\end{aligned}\quad (13)$$

where $Q_\phi(x_n | z'_n; \varepsilon)$ is computed from $\text{Transformer}(H_0, H_{1:M})$. In practice, we also learn $P_\theta(x, z; \varepsilon) = P_0(z)P_\theta(x | z; \varepsilon)$ as a combinatorial structure feature enhanced latent generative model, where $P_\theta(x | z; \varepsilon)$ is a combinatorial structure feature enhanced Transformer with a learnable mapping matrix W_η to transform the combinatorial structure features into the embedding space. d is the dimensionality of $\sigma_\psi(x_n)$ and $\sigma_{\psi,i}(x_n)$ is its i -th dimension.

M-step:

$$\psi^{(t+1)} = \psi^{(t)} + r * \nabla_\psi \tilde{\mathcal{L}}_t, \quad \phi^{(t+1)} = \phi^{(t)} + r * \nabla_\phi \tilde{\mathcal{L}}_t \quad (14)$$

The overall learning algorithm is given in Appendix B.4.

4. Experiments

In this section, we conduct extensive experiments to validate the effectiveness of our proposed IsEM-Pro on protein sequence design task.

4.1. Implementation Details

We first train a VAE model on raw protein sequences using a 6-layer Transformer as the encoder and a 2-layer Transformer as the decoder. We add the combinatorial structure features to the Transformer decoder input as described in Sec. 3.4 and its weight W_η is learned jointly. The raw protein probability $P_\theta(x | z; \varepsilon)$ is then defined by the combinatorial structure feature augmented decoder in this VAE. The protein combinatorial structure constraints ε are learned on the training sequences for each dataset, rather than using real multiple sequence alignments (MSAs), to ensure a fair comparison. The approximate posterior probability $\tilde{P}_\theta(z | x)$ is defined by its encoder. Its embedding and feed-forward network sizes are 320 and 1280. The encoder is initialized with the pre-trained ESM-2 weights (Lin et al., 2023)¹. The latent variable z 's dimension is also 320. We use another 6-layer Transformer encoder followed by a linear mapping to learn μ_ψ and σ_ψ for $R_\psi(z | x) = \mathcal{N}(\mu_\psi, \sigma_\psi^2 \mathbf{I})$ and another 2-layer Transformer decoder to learn $Q_\phi(x | z; \varepsilon)$. We initialize parameters $\psi^{(0)}$ and $\phi^{(0)}$ with θ .

The number of iterations in the importance sampling-based MCEM is set to 10. The mini-batch size and learning rate are set to 4,096 tokens and 1e-5 respectively. The model is trained with 1 NVIDIA RTX A6000 GPU card. We apply Adam algorithm (Kingma & Ba, 2015) as the optimizer with a linear warm-up over the first 4,000 steps and linear decay for later steps. We randomly split each dataset into training/validation sets with the ratio of 9:1. We run all the experiments for five times and report the average scores. More experimental settings are given in Appendix B.1. Following (Kamisetty et al., 2013), we set $\alpha_{\text{single}} = 1$ and $\alpha_{\text{pair}} = 0.2 * (M - 1)$ with M equals to the protein sequence length.

In inference, we design protein sequences by taking the wild-type as encoder input and the latent vector is sampled from prior distribution $Q_0(z) = N(0, \mathbf{I})$. The sequences are decoded using sampling strategy with top-5. The candidate number is set to K=128 following the setting of Jain et al. (2022) on the GFP dataset.

4.2. Datasets

Following Ren et al. (2022), we evaluate our method on the following eight protein engineering benchmarks:

- (1) **Green Fluorescent Protein (avGFP)**: The goal is to design sequences with higher log-fluorescence intensity values. We collect data following Sarkisyan et al. (2016).
- (2) **Adeno-Associated Viruses (AAV)**: The target is to generate amino-acid segment (position 561 – 588) for higher gene therapeutic efficiency. We collect data following Bryant et al. (2021).
- (3) **TEM-1 β -Lactamase (TEM)**: The goal is

¹https://dl.fbaipublicfiles.com/faresm/models/esm2_t6.8M_UR50D.pt

to design high thermodynamic-stable sequences. We merge the data from Firnberg et al. (2014). (4) **Ubiquitination Factor Ube4b (E4B)**: The objective is to design sequences with higher enzyme activity. We gather data following Starita et al. (2013). (5) **Aliphatic Amide Hydrolase (AMIE)**: The goal is to produce amidase sequences with higher enzyme activity. We merge data following Wrenbeck et al. (2017). (6) **Levoglucosan Kinase (LGK)**: The target is to optimize LGK protein sequences with improved enzyme activity. We collect data following Klesmith et al. (2015). (7) **Poly(A)-binding Protein (Pab1)**: The goal is to design sequences with higher binding fitness to multiple adenosine monophosphates. We gather data following Melamed et al. (2013). (8) **SUMO E2 Conjugase (UBE2I)**: We aim to find human SUMO E2 conjugase with higher growth rescue rate. Data are obtained following Weile et al. (2017). The detailed data statistics, including protein sequence length, data size and data source are provided in Appendix A.

4.3. Baseline Models

We compare our method against the following representative baselines: (1) **CMA-ES** (Hansen & Ostermeier, 2001) is a famous evolutionary search algorithm. (2) **FBGAN** proposed by Gupta & Zou (2019) is a novel feedback-loop architecture with generative model GAN. (3) **DbAS** (Brookes & Listgarten, 2018) is a probabilistic modeling framework and uses adaptive sampling algorithm. (4) **CbAS** (Brookes et al., 2019) improves on DbAS by conditioning on the desired properties. (5) **PEX** proposed by Ren et al. (2022) is a model-guided sequence design algorithm using proximal exploration. (6) **GFlowNet-AL** (Jain et al., 2022) applies GFlowNet to design biological sequences. We use the implementations of CMA-ES, DbAS and CbAS provided in Trabucco et al. (2022) and for other baselines, we apply their released codes. To better analyze the influence of different components in our model, we also conduct ablation tests as follows: (1) **IsEM-Pro-w/o-ESM** removes ESM-2 as encoder initialization. (2) **IsEM-Pro-w/o-ISEM** removes iterative optimization process. (3) **IsEM-Pro-w/o-MRFs** removes MRFs features and iterative optimization process. (4) **IsEM-Pro-w/o-LV** removes latent variable, MRFs features and iterative optimization process. (5) **ESM-Search** samples sequences from the softmax distribution obtained by finetuning ESM-2 on the protein datasets and taking the wild-type as input.

4.4. Evaluation Metrics

We use three automatic metrics to evaluate the performance of the designed sequences: (1) **MFS**: Maximum fitness score. The oracle model adopted to evaluate MFS is described in Appendix B.1; (2) **Diversity** proposed by Jain et al. (2022) is used to evaluate how different the designed candidates are from each other; (3) **Novelty** proposed by

Jain et al. (2022) is used to evaluate how different the proposed candidates are from the sequences in training data.

4.5. Main Results

Table 1, 2 and 3 respectively report the maximum fitness scores, diversity scores and novelty scores of all models.

IsEM-Pro achieves the highest fitness scores on all protein families and outperforms the previous best method GFlowNet-AL by 55% on average (Table 1). The reasons are two-folds. On one hand, the importance sampling based MCEM can help our model to navigate to a better region instead of getting trapped in a worse local optima. On the other hand, the combinatorial structure features help to recognize the preferred mutation patterns which have higher success rate under the nature selection pressure, potentially leading to sequences with higher fitness scores.

IsEM-Pro achieves the highest average diversity score over the eight tasks (Table 2). Our model gains the highest diversity on 4 out of 8 tasks while GFlowNet-AL gains 2 and CMA-ES gains 1. It indicates that though involving combinatorial structure constraints can give the guidance for preferred protein patterns, it might also limit the sequence design to these patterns to some extent. The involved latent variable can capture complex inter-dependencies among amino acids, which benefits for more diverse protein design.

IsEM-Pro can design more novel protein sequences on all datasets (Table 3). Our model achieves higher novelty scores on all datasets due to the reason that more new samples are involved during the importance sampling based iterative optimization process, which is beneficial for more novel protein design.

4.6. Ablation Study

Bottom halves of Table 1, 2 and 3 report the results of ablation tests. IsEM-Pro-w/o-MRFs improves the average diversity and novelty scores by as much as 33x compared with IsEM-Pro-w/o-LV, which demonstrates that introducing a latent variable can significantly help to generate diverse proteins. IsEM-Pro-w/o-MRFs achieves higher maximum fitness scores than IsEM-Pro-w/o-ESM on all datasets, validating that adopting a pretrained protein language model as the encoder helps to design more satisfactory protein sequences. However, directly finetuning ESM-2 to sample candidates (ESM-Search) drops 1.3 points on average fitness score compared with taking ESM-2 as an encoder (IsEM-Pro-w/o-MRFs), demonstrating that ESM-2 is not suitable for direct sequence design. Incorporating combinatorial structure features can further improve the fitness of the designed proteins (IsEM-Pro-w/o-ISEM V.S. IsEM-Pro-w/o-MRFs), based on which learning the proposal distribution by importance sampling based MCEM can better

Importance Weighted Expectation-Maximization for Protein Sequence Design

Models	avGFP	AAV	TEM	E4B	AMIE	LGK	Pab1	UBE2I	Average
CMA-ES	4.492	-3.417	0.375	-0.768	-8.224	-0.077	0.164	2.461	-0.624
FBGAN	1.251	-4.227	0.006	0.369	-2.410	-1.206	0.029	0.208	-0.747
DbAS	3.548	4.327	0.003	-1.286	-2.658	-1.148	1.524	3.088	0.924
CbAS	3.550	4.336	0.106	-1.000	-1.306	-0.362	1.842	3.263	1.303
PEX	3.764	3.265	0.121	5.019	-0.474	0.007	1.153	1.995	1.856
GFlowNet-AL	5.062	1.205	1.552	3.155	0.059	0.027	2.168	3.576	2.101
ESM-Search	2.610	-5.099	0.148	-1.860	-2.351	-0.029	1.406	3.244	-0.241
IsEM-Pro	6.185	4.813	1.850	5.737	0.062	0.035	2.923	4.536	3.267
- w/o ESM	1.214	-4.313	0.005	-1.352	-6.376	-0.225	0.072	1.843	-1.141
- w/o ISEM	4.708	1.130	0.708	0.046	-2.335	-0.077	1.913	0.475	1.342
- w/o MRFs	4.376	1.008	0.952	0.045	-1.771	-0.012	1.652	2.418	1.083
- w/o LV	4.274	2.251	0.078	-1.612	-2.266	-0.931	0.041	-0.262	0.196

Table 1. Maximum fitness scores (MFS) of all methods on eight datasets. Higher values indicate better functional properties in the dataset. Our proposed IsEM-Pro achieves the highest fitness scores on all datasets.

Models	avGFP	AAV	TEM	E4B	AMIE	LGK	Pab1	UBE2I	Average
CMA-ES	225.12	23.50	261.60	86.81	283.90	317.08	61.16	140.92	175.01
FBGAN	0.64	8.31	0.46	3.87	33.87	17.35	3.07	3.00	8.82
DbAS	3.04	3.00	3.67	5.94	1.32	2.30	4.05	11.80	4.33
CbAS	1.31	3.01	7.03	7.09	6.01	6.15	9.86	22.73	8.23
PEX	6.83	4.35	10.26	5.22	7.56	13.24	5.33	10.32	7.88
GFlowNet-AL	224.78	25.57	266.43	43.62	219.84	212.25	37.13	49.79	134.92
ESM-Search	3.79	3.58	11.56	3.82	3.83	3.78	5.71	6.59	5.33
IsEM-Pro	218.62	22.92	202.09	91.35	293.30	405.99	68.27	122.66	178.15
- w/o ESM	204.21	13.87	194.78	7.90	276.88	362.98	3.30	119.87	147.96
- w/o ISEM	122.15	22.91	70.12	86.17	145.74	169.17	60.22	13.05	153.09
- w/o MRFs	217.35	17.02	225.88	84.64	268.07	381.66	66.26	143.29	175.52
- w/o LV	22.70	5.26	10.12	5.24	8.65	20.28	0.67	2.98	9.48

Table 2. Diversity scores of all models on eight datasets. Higher values indicate more diverse protein sequences. Our IsEM-Pro achieves the highest average diversity scores over the eight protein datasets.

Models	avGFP	AAV	TEM	E4B	AMIE	LGK	Pab1	UBE2I	Average
CMA-ES	221.55	22.73	269.25	93.78	256.07	415.63	59.35	128.10	183.30
FBGAN	0.05	2.76	0.08	0.63	57.87	39.36	0.75	0.80	12.43
DbAS	1.01	3.01	1.47	1.09	1.12	1.63	1.64	2.05	1.66
CbAS	4.02	3.03	2.06	1.90	1.33	1.09	2.71	2.95	1.92
PEX	3.59	1.88	8.57	4.08	4.50	10.53	3.63	10.24	5.87
GFlowNet-AL	221.95	22.83	266.99	86.78	316.79	412.58	61.33	143.28	191.56
ESM-Search	1.50	1.83	7.32	1.21	0.92	0.91	5.46	6.14	3.16
IsEM-Pro	226.31	23.81	270.27	96.57	332.23	420.93	70.09	153.27	199.18
- w/o ESM	198.08	9.25	176.20	3.79	264.36	340.62	1.49	110.53	138.04
- w/o ISEM	195.85	16.36	244.05	85.81	306.22	382.12	53.01	99.51	180.40
- w/o MRFs	207.59	9.53	221.49	90.81	244.14	327.14	58.44	129.75	161.11
- w/o LV	16.67	0.89	4.39	0.48	3.98	9.64	0.21	1.76	4.75

Table 3. Novelty scores of all models on eight datasets. Higher values indicate more novel protein sequences. Our IsEM-Pro achieves the highest novelty scores on all the datasets.

promote more desirable, diverse and novel protein generation. We also provide the model performance with a fully-connected decoder (non-autoregressive decoder) instead of the current autoregressive one in Appendix F.2. It shows the non-autoregressive decoder can not generate good candidates for protein with longer sequences such as LGK (439 amino acids).

5. Analysis

5.1. Approximate KL Divergence

To validate how close the proposal distribution $Q_\phi(x)$ is to the posterior distribution $P_\theta(x|S)$, we calculate the KL divergence through Monte Carlo approximation. Here we calculate the KL divergence between $Q_\phi(x)$ and $P_\theta(x|S)$ as

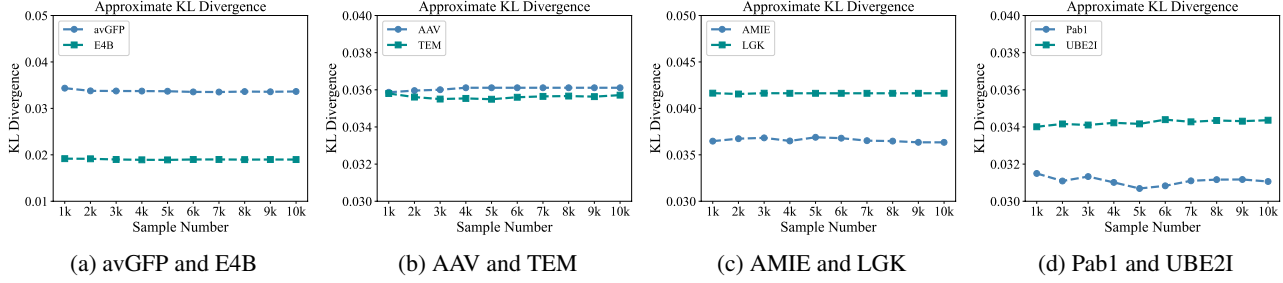


Figure 4. Approximate KL divergence on eight protein datasets. It shows the variance of KL divergence is very small over different sample size for all datasets, giving empirical evidence that when we sample from the ultimate $Q_\phi(x)$, it has minor difference compared with sampling from the posterior distribution $P_\theta(x|S)$.



Figure 5. 3-D visualization of our designed green fluorescent protein, validating that IsEM-Pro can generate realistic fluorescent protein.

Methods	MFS	Diversity	Novelty
Latent-Add	4.102	218.26	209.85
Latent-Memory	4.040	219.00	211.38
IsEM-Pro	6.185	218.62	226.31

Table 4. Results of different schemes of introducing a latent variable with a pretrained encoder evaluated on avGFP dataset. It shows that our method, which takes the latent representation as the first token input of decoder, achieves a higher fitness and novelty scores though with a mild decrease on diversity.

we have proven in Lemma C.1 (provided in Appendix C.2) that the sampling difference between these two distributions can be bounded under this divergence. We leverage an unbiased and low-variance estimator (proof shown in Appendix C.1) to approximate the KL divergence as follows:

$$D_{KL}(Q_\phi(x)||P_\theta(x|S)) = E_{Q_\phi(x)}[r(x) - 1 - \log r(x)] \quad (15)$$

where $r(x) = \frac{P_\theta(x|S)}{Q_\phi(x)}$. The approximate KL divergence on eight protein datasets over 1k-10k samples are illustrated in Figure 4. From the figure, we can see that the variance of KL divergence is very small over different sample size for all datasets. Besides, the KL divergence finally arrives at a small value, such as 0.018 for E4B and 0.033 for avGFP. It gives empirical evidence that when we sample from the ultimate $Q_\phi(x)$, it has minor difference compared with sampling from the posterior distribution $P_\theta(x|S)$.

5.2. Effect of VAE Implementation Method

Next, we study the effect of different implementation schemes of involving a latent variable with a pretrained

encoder. Some works have tried adding the latent representation to the original embedding layer (Latent-Add) or using it as an additional memory (Latent-Memory) when adopting a pretrained language model as encoder (Li et al., 2020). We also implement our model with these two schemes, and evaluate the model performance on avGFP dataset. Table 4 shows that our method, which takes the latent representation as the first token input of decoder, achieves a higher fitness and novelty scores though with a mild decrease on diversity. We also provide the conditional VAE (CVAE) performance in Appendix F.1, which shows decreased performance. It demonstrates that iterative sampling inside the EM algorithm with explicit fitness constraints can lead to better proteins.

5.3. Case Study

To gain an insight on how well the designed proteins are, we analyze the generated avGFP sequence with highest fitness in detail using Phyre2 tool (Kelley et al., 2015). Figure 5 (a) illustrates the generated variant can fold stably. According to the software, the most similar protein is Cytochrome b562 integral fusion with enhanced green fluorescent protein (EGFP) (Edwards et al., 2008). There are 227 residues (96% of the candidate sequence) have been modeled with 100.0% confidence by using this protein as template. Details are given in Appendix D. Figure 5(b) visualizes the superposition of the top-5 most similar templates to our sequence in the protein data bank, which are all fluorescent proteins and show highly consistent structure in most regions, validating that our model can design a real fluorescent protein.

6. Related Work

Machine Learning for Protein Fitness Landscape Prediction. Machine learning has been increasingly used for modeling protein fitness landscape, which is crucial for protein engineering. Some work leverage co-evolution information from multiple sequence alignments to predict fitness scores (Kamisetty et al., 2013; Luo et al., 2021). Melamed et al. (2013) propose to construct a deep latent generative model to capture higher-order mutations. Meier et al. (2021) propose to use pretrained protein language models to enable zero-shot prediction. The learned protein landscape models can be used to replace the expensive wet-lab validation to screen enormous designed sequences (Rao et al., 2019; Ren et al., 2022).

Methods for Protein Sequence Design. Protein sequence design has been studied with a wide variety of methods, including traditional directed evolution (Arnold, 1998; Dalby, 2011; Packer & Liu, 2015; Arnold, 2018) and machine learning methods. The mainly used machine learning algorithms include reinforcement learning (Angermüller et al., 2020; Jain et al., 2022), Bayesian optimization (Belanger et al., 2019; Moss et al., 2020; Terayama et al., 2021), search using adaptive evolution methods (Hansen, 2006; Swersky et al., 2020), likelihood-free inference (Zhang et al., 2022), deep generative models (Brookes & Listgarten, 2018; Madani et al., 2020; Kumar & Levine, 2020; Das et al., 2021; Hoffman et al., 2022; Melnyk et al., 2021; Ren et al., 2022) and latent deep generative model (Brookes et al., 2019). Our approach has the same objective using Bayes rule and KL divergence (Eq. (3) (4)) as CbAS (Brookes et al., 2019). Different from CbAS, our approach includes two latent models – one representing the raw protein sequences and the other representing “good” proteins. The solution is derived under the Monte Carlo EM framework. Secondly, our approach is essentially self-training since we use both real proteins and (importance) sampled high-quality ones in each iteration to train the generation model. Thirdly, we augment the generative model with combinatorial structure features learned from MRF. This comprehensive framework not only enables the generative model to explore optimal regions in either the Fujiyama landscape or the Badlands landscape (Kauffman & Weinberger, 1989), but also significantly enhances protein diversity and novelty.

7. Discussion

We opt for MCEM over the standard VAE to learn the latent generative model due to our belief that the standard VAE has certain limitations. The major limitation of the standard VAE is that it maximizes the likelihood of observed data, say protein sequence, but higher likelihood does not necessarily associate with higher fitness of a protein. Our IsEM-Pro tackles this by explicitly modelling fitness constraints and

amino acid correlation in a protein sequence. As shown in Table 1, 2 and 3, the fitness scores of standard VAE (IsEM-Pro-w/o-MRFs) decrease a lot compared with our IsEM-Pro and the diversity and novelty scores also slightly decrease. Besides, as Dieng & Paisley (2019) analyzed, VAE amortizes the cost of inference by using a recognition network to parameterize the variational family, which introduces an amortization gap and leads to approximate posteriors of reduced expressivity due to the problem known as posterior collapse. Instead, EM directly maximizes the log marginal likelihood of the data and each iteration in EM is guaranteed to increase the log marginal likelihood from the previous iteration (Bishop, 2006). Therefore, using EM in the context of deep generative models could lead to better performance. Additionally, standard VAE attempts to perform variational inference through a direct, discriminative mapping from data observations to approximate posterior parameters. Though generative models can adapt to accommodate sub-optimal approximate posteriors, it’s likely limited to direct inference mapping, leading to being trapped in a worse local optima (Marino et al., 2018). Instead, EM guarantees that the log marginal likelihood of the data keeps increasing in each iteration, helping our model climb much closer to the global optima.

One limitation of this work is, although our model has demonstrated promising results, the designed protein sequences have not undergone wet-lab testing and there might be some uncertainty correlated with the oracle model. The results reported in our paper are averaged over five runs, which can accommodate some variance on this metric. Furthermore, all fitness data used in our paper are obtained from wet-lab experiments, ensuring that the fitness values of the training datasets are realistic. To reduce the cost of wet-lab validation and select good protein candidates, we evaluate the designed sequences using the oracle model trained on real protein sequences following (Ren et al., 2022; Jain et al., 2022), which may associate with some uncertainties. While there is currently no perfect automatic metric available, we believe that new methods should be encouraged in the field of computational biology. With time, we are confident that the field will continue to develop and mature.

8. Conclusion

This paper proposes IsEM-Pro, a latent generative model for protein sequence design, which incorporates additional combinatorial structure features learned by MRFs. We use importance weighted EM to learn the model, which can not only enhance design diversity and novelty, but also lead to protein sequences with higher fitness. Experimental results on eight protein sequence design tasks show that our method outperforms several strong baselines on all metrics.

Acknowledgements

The authors would like to thank the anonymous reviewers for their valuable comments. We are also grateful to Siqui Ouyang, Yujian Liu, Jingjing Xu, Danqing Wang, Antonis Antoniadis, and Jennifer Listgarten for their great suggestions. This work is partially supported by UCSB Faculty Research Award.

References

- Angermüller, C., Dohan, D., Belanger, D., Deshpande, R., Murphy, K., and Colwell, L. Model-based reinforcement learning for biological sequence design. In *Proc. of ICLR*. OpenReview.net, 2020.
- Arnold, F. H. Design by directed evolution. *Accounts of chemical research*, 31(3):125–131, 1998.
- Arnold, F. H. Directed evolution: bringing new chemistry to life. *Angewandte Chemie International Edition*, 57(16):4143–4148, 2018.
- Belanger, D., Vora, S., Mariet, Z., Deshpande, R., Dohan, D., Angermüller, C., Murphy, K., Chapelle, O., and Colwell, L. Biological sequences design using batched bayesian optimization. 2019.
- Bishop, C. M. *Pattern recognition and machine learning*. Springer, 2006.
- Biswas, S., Khimulya, G., Alley, E. C., Esvelt, K. M., and Church, G. M. Low-n protein engineering with data-efficient deep learning. *Nature methods*, 18(4):389–396, 2021.
- Bloom, J. D. and Arnold, F. H. In the light of directed evolution: pathways of adaptive protein evolution. *Proceedings of the National Academy of Sciences*, 106(supplement_1):9995–10000, 2009.
- Brookes, D. H. and Listgarten, J. Design by adaptive sampling. *ArXiv preprint*, abs/1810.03714, 2018.
- Brookes, D. H., Park, H., and Listgarten, J. Conditioning by adaptive sampling for robust design. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proc. of ICML*, volume 97 of *Proceedings of Machine Learning Research*, pp. 773–782. PMLR, 2019.
- Bryant, D. H., Bashir, A., Sinai, S., Jain, N. K., Ogden, P. J., Riley, P. F., Church, G. M., Colwell, L. J., and Kelsic, E. D. Deep diversification of an aav capsid protein by machine learning. *Nature Biotechnology*, 39(6):691–696, 2021.
- Chothia, C. Principles that determine the structure of proteins. *Annual review of biochemistry*, 53(1):537–572, 1984.
- Csiszár, I. and Körner, J. *Information theory: coding theorems for discrete memoryless systems*. Cambridge University Press, 2011.
- Dalby, P. A. Strategy and success for the directed evolution of enzymes. *Current opinion in structural biology*, 21(4):473–480, 2011.
- Das, P., Sercu, T., Wadhawan, K., Padhi, I., Gehrmann, S., Cipcigan, F., Chenthamarakshan, V., Strobel, H., Dos Santos, C., Chen, P.-Y., et al. Accelerated antimicrobial discovery via deep generative models and molecular dynamics simulations. *Nature Biomedical Engineering*, 5(6):613–623, 2021.
- Dieng, A. B. and Paisley, J. Reweighted expectation maximization. *ArXiv preprint*, abs/1906.05850, 2019.
- Edwards, W. R., Busse, K., Allemann, R. K., and Jones, D. D. Linking the functions of unrelated proteins using a novel directed evolution domain insertion method. *Nucleic acids research*, 36(13):e78–e78, 2008.
- Firnberg, E., Labonte, J. W., Gray, J. J., and Ostermeier, M. A comprehensive, high-resolution map of a gene’s fitness landscape. *Molecular biology and evolution*, 31(6):1581–1592, 2014.
- Fox, R. J., Davis, S. C., Mundorff, E. C., Newman, L. M., Gavrilovic, V., Ma, S. K., Chung, L. M., Ching, C., Tam, S., Muley, S., et al. Improving catalytic function by prosar-driven enzyme evolution. *Nature biotechnology*, 25(3):338–344, 2007.
- Go, N. Theoretical studies of protein folding. *Annual review of biophysics and bioengineering*, 12(1):183–210, 1983.
- Gupta, A. and Zou, J. Feedback gan for dna optimizes protein functions. *Nature Machine Intelligence*, 1(2):105–111, 2019.
- Hansen, N. The cma evolution strategy: a comparing review. *Towards a new evolutionary computation*, pp. 75–102, 2006.
- Hansen, N. and Ostermeier, A. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2):159–195, 2001.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-vae: Learning basic visual concepts with a constrained variational framework. In *Proc. of ICLR*. OpenReview.net, 2017.
- Hoffman, S. C., Chenthamarakshan, V., Wadhawan, K., Chen, P.-Y., and Das, P. Optimizing molecules using efficient queries from property evaluations. *Nature Machine Intelligence*, 4(1):21–31, 2022.

- Hopf, T. A., Ingraham, J. B., Poelwijk, F. J., Schärfe, C. P., Springer, M., Sander, C., and Marks, D. S. Mutation effects predicted from sequence co-variation. *Nature biotechnology*, 35(2):128–135, 2017.
- Jain, M., Bengio, E., Hernández-García, A., Rector-Brooks, J., Dossou, B. F. P., Ekbote, C. A., Fu, J., Zhang, T., Kilgour, M., Zhang, D., Simine, L., Das, P., and Bengio, Y. Biological sequence design with gflownets. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvári, C., Niu, G., and Sabato, S. (eds.), *Proc. of ICML*, volume 162 of *Proceedings of Machine Learning Research*, pp. 9786–9801. PMLR, 2022.
- Kamisetty, H., Ovchinnikov, S., and Baker, D. Assessing the utility of coevolution-based residue–residue contact predictions in a sequence-and structure-rich era. *Proceedings of the National Academy of Sciences*, 110(39): 15674–15679, 2013.
- Kauffman, S. A. and Weinberger, E. D. The nk model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of theoretical biology*, 141(2):211–245, 1989.
- Kelley, L. A., Mezulis, S., Yates, C. M., Wass, M. N., and Sternberg, M. J. The phyre2 web portal for protein modeling, prediction and analysis. *Nature protocols*, 10(6): 845–858, 2015.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *Proc. of ICLR*, 2015.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In Bengio, Y. and LeCun, Y. (eds.), *Proc. of ICLR*, 2014.
- Klesmith, J. R., Bacik, J.-P., Michalczyk, R., and Whitehead, T. A. Comprehensive sequence-flux mapping of a levoglucosan utilization pathway in e. coli. *ACS synthetic biology*, 4(11):1235–1243, 2015.
- Kumar, A. and Levine, S. Model inversion networks for model-based optimization. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Proc. of NeurIPS*, 2020.
- Labrou, N. E. Random mutagenesis methods for in vitro directed enzyme evolution. *Current Protein and Peptide Science*, 11(1):91–100, 2010.
- Lagassé, H. D., Alexaki, A., Simhadri, V. L., Katagiri, N. H., Jankowski, W., Sauna, Z. E., and Kimchi-Sarfaty, C. Recent advances in (therapeutic protein) drug development. *F1000Research*, 6, 2017.
- Li, C., Gao, X., Li, Y., Peng, B., Li, X., Zhang, Y., and Gao, J. Optimus: Organizing sentences via pre-trained modeling of a latent space. In Webber, B., Cohn, T., He, Y., and Liu, Y. (eds.), *Proc. of EMNLP*, pp. 4678–4699. Association for Computational Linguistics, 2020.
- Lin, Z., Akin, H., Rao, R., Hie, B., Zhu, Z., Lu, W., Smetanin, N., Verkuil, R., Kabeli, O., Shmueli, Y., et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637): 1123–1130, 2023.
- Luo, Y., Jiang, G., Yu, T., Liu, Y., Vo, L., Ding, H., Su, Y., Qian, W. W., Zhao, H., and Peng, J. Ecnet is an evolutionary context-integrated deep learning framework for protein engineering. *Nature communications*, 12(1): 5743, 2021.
- Ma, J. K., Drake, P. M., and Christou, P. The production of recombinant pharmaceutical proteins in plants. *Nature reviews genetics*, 4(10):794–805, 2003.
- Madani, A., McCann, B., Naik, N., Keskar, N. S., Anand, N., Eguchi, R. R., Huang, P.-S., and Socher, R. Progen: Language modeling for protein generation. *ArXiv preprint*, abs/2004.03497, 2020.
- Marino, J., Yue, Y., and Mandt, S. Iterative amortized inference. In Dy, J. G. and Krause, A. (eds.), *Proc. of ICML*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3400–3409. PMLR, 2018.
- Meier, J., Rao, R., Verkuil, R., Liu, J., Sercu, T., and Rives, A. Language models enable zero-shot prediction of the effects of mutations on protein function. In Ranzato, M., Beygelzimer, A., Dauphin, Y. N., Liang, P., and Vaughan, J. W. (eds.), *Proc. of NeurIPS*, pp. 29287–29303, 2021.
- Melamed, D., Young, D. L., Gamble, C. E., Miller, C. R., and Fields, S. Deep mutational scanning of an rrm domain of the saccharomyces cerevisiae poly (a)-binding protein. *Rna*, 19(11):1537–1551, 2013.
- Melnyk, I., Das, P., Vijil, V., and Lozano, A. Benchmarking deep generative models for diverse antibody sequence design. *Proc. of NeurIPS*, 2021.
- Moss, H. B., Leslie, D. S., Beck, D., Gonzalez, J., and Rayson, P. BOSS: bayesian optimization over string spaces. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Proc. of NeurIPS*, 2020.
- Packer, M. S. and Liu, D. R. Methods for the directed evolution of proteins. *Nature Reviews Genetics*, 16(7): 379–394, 2015.

- Qian, L., Zhou, H., Bao, Y., Wang, M., Qiu, L., Zhang, W., Yu, Y., and Li, L. Glancing transformer for non-autoregressive neural machine translation. In Zong, C., Xia, F., Li, W., and Navigli, R. (eds.), *Proc. of ACL*, pp. 1993–2003. Association for Computational Linguistics, 2021.
- Rao, R., Bhattacharya, N., Thomas, N., Duan, Y., Chen, P., Canny, J. F., Abbeel, P., and Song, Y. S. Evaluating protein transfer learning with TAPE. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), *Proc. of NeurIPS*, pp. 9686–9698, 2019.
- Ren, Z., Li, J., Ding, F., Zhou, Y., Ma, J., and Peng, J. Proximal exploration for model-guided protein sequence design. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvári, C., Niu, G., and Sabato, S. (eds.), *Proc. of ICML*, volume 162 of *Proceedings of Machine Learning Research*, pp. 18520–18536. PMLR, 2022.
- Riesselman, A. J., Ingraham, J. B., and Marks, D. S. Deep generative models of genetic variation capture the effects of mutations. *Nature methods*, 15(10):816–822, 2018.
- Rives, A., Meier, J., Sercu, T., Goyal, S., Lin, Z., Liu, J., Guo, D., Ott, M., Zitnick, C. L., Ma, J., 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.
- Robert, C. and Casella, G. *Monte Carlo statistical methods*. Springer, 2004.
- Romero, P. A. and Arnold, F. H. Exploring protein fitness landscapes by directed evolution. *Nature reviews Molecular cell biology*, 10(12):866–876, 2009.
- Sarkisyan, K. S., Bolotin, D. A., Meer, M. V., Usmanova, D. R., Mishin, A. S., Sharonov, G. V., Ivankov, D. N., Bozhanova, N. G., Baranov, M. S., Soylemez, O., et al. Local fitness landscape of the green fluorescent protein. *Nature*, 533(7603):397–401, 2016.
- Seemayer, S., Gruber, M., and Söding, J. Ccnpred—fast and precise prediction of protein residue–residue contacts from correlated mutations. *Bioinformatics*, 30(21):3128–3130, 2014.
- Singh, A., Pandey, A., Srivastava, A. K., Tran, L.-S. P., and Pandey, G. K. Plant protein phosphatases 2c: from genomic diversity to functional multiplicity and importance in stress management. *Critical Reviews in Biotechnology*, 36(6):1023–1035, 2016.
- Starita, L. M., Pruneda, J. N., Lo, R. S., Fowler, D. M., Kim, H. J., Hiatt, J. B., Shendure, J., Brzovic, P. S., Fields, S., and Klevit, R. E. Activity-enhancing mutations in an e3 ubiquitin ligase identified by high-throughput mutagenesis. *Proceedings of the National Academy of Sciences*, 110(14):E1263–E1272, 2013.
- Starr, T. N. and Thornton, J. W. Exploring protein sequence–function landscapes. *Nature biotechnology*, 35(2):125–126, 2017.
- Swersky, K., Rubanova, Y., Dohan, D., and Murphy, K. Amortized bayesian optimization over discrete spaces. In Adams, R. P. and Gogate, V. (eds.), *Proceedings of the Thirty-Sixth Conference on Uncertainty in Artificial Intelligence, UAI 2020, virtual online, August 3-6, 2020*, volume 124 of *Proceedings of Machine Learning Research*, pp. 769–778. AUAI Press, 2020.
- Terayama, K., Sumita, M., Tamura, R., and Tsuda, K. Black-box optimization for automated discovery. *Accounts of Chemical Research*, 54(6):1334–1346, 2021.
- Trabucco, B., Geng, X., Kumar, A., and Levine, S. Designbench: Benchmarks for data-driven offline model-based optimization. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvári, C., Niu, G., and Sabato, S. (eds.), *Proc. of ICML*, volume 162 of *Proceedings of Machine Learning Research*, pp. 21658–21676. PMLR, 2022.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), *Proc. of NeurIPS*, pp. 5998–6008, 2017.
- Wei, G. C. and Tanner, M. A. A monte carlo implementation of the em algorithm and the poor man’s data augmentation algorithms. *Journal of the American statistical Association*, 85(411):699–704, 1990.
- Weile, J., Sun, S., Cote, A. G., Knapp, J., Verby, M., Mellor, J. C., Wu, Y., Pons, C., Wong, C., van Lieshout, N., et al. A framework for exhaustively mapping functional missense variants. *Molecular systems biology*, 13(12): 957, 2017.
- Wrenbeck, E. E., Azouz, L. R., and Whitehead, T. A. Single-mutation fitness landscapes for an enzyme on multiple substrates reveal specificity is globally encoded. *Nature communications*, 8(1):1–10, 2017.
- Wright, C. F., Teichmann, S. A., Clarke, J., and Dobson, C. M. The importance of sequence diversity in the aggregation and evolution of proteins. *Nature*, 438(7069): 878–881, 2005.

Zhang, D., Fu, J., Bengio, Y., and Courville, A. C. Unifying likelihood-free inference with black-box optimization and beyond. In *Proc. of ICLR*. OpenReview.net, 2022.

A. Data Statistics

We provide the detailed data statistics in the following table, including protein sequence length, data size and data source. We have checked and cleaned the data and make sure the data do not contain personally identifiable information or offensive content.

Protein	Length	Size	Data Source
avGFP	237	49,855	https://figshare.com/articles/dataset/Local_fitness_landscape_of_the_green_fluorescent_protein
AAV	28	296,914	https://github.com/churchlab/Deep_diversification_AAV
TEM	286	17,238	https://github.com/facebookresearch/esm/tree/main/examples/data
E4B	102	91,033	https://figshare.com/articles/dataset
AMIE	341	6,631	https://figshare.com/articles/dataset/Normalized_fitness_values_for_AmiE_selections/3505901/2
LGK	439	8,069	https://figshare.com/articles/dataset
Pab1	75	36,522	https://figshare.com/articles/dataset
UBE2I	159	5,355	http://dalai.mshri.on.ca/~jweile/projects/dmsData/

Table 5. Detailed statistics of the eight protein datasets.

B. More Implementation Details

B.1. Additional Experimental Settings

We apply the annealing schedule for the KL term during $P_\theta(x, z)$ training process following β -VAE (Higgins et al., 2017) to prevent posterior collapse. Specifically, the KL term coefficient starts from 0 and is gradually increased to 1.0 as training goes on. At each iteration in importance sampling based EM learning process, the number of samples from current $Q_\phi(x, z)$ is set to 10% of the training data size.

Following (Ren et al., 2022), We construct the oracle model $f(x)$ by adopting the features produced by ESM-1b (Rives et al., 2021) with dimension 1280 and finetuning an Attention1D decoder to predict the fitness values. Since Brookes et al. (2019) state that the results are insensitive when λ is set in the range [50, 100]-th percentile of the fitness scores in the training set, we set λ to 50-th percentile of the fitness values in the training data to accommodate more diversity.

B.2. KL Divergence for Two Gaussian Distribution

The kl divergence for two Gaussian distribution is defined as follows:

$$\begin{aligned}
 p(x) &= \frac{1}{2\pi^{0.5n}|\Sigma|^{0.5}} \exp\left(-\frac{1}{2}(x-\mu)^T \Sigma^{-1}(x-\mu)\right) \\
 q(x) &= \frac{1}{2\pi^{0.5n}|L|^{0.5}} \exp\left(-\frac{1}{2}(x-m)^T L^{-1}(x-m)\right) \\
 D_{KL}(p||q) &= \frac{1}{2} \left\{ \log \frac{|L|}{|\Sigma|} + \text{Tr}(L^{-1}\Sigma) + (\mu-m)^T L^{-1}(\mu-m) - n \right\}
 \end{aligned} \tag{16}$$

When $p(x) = \mathcal{N}(\mu, \sigma^2 \mathbf{I})$ and $q(x) = \mathcal{N}(0, \mathbf{I})$, the KL divergence can be computed in closed form as follows:

$$D_{KL}(p||q) = \frac{1}{2} \|\sigma\|_2^2 + \frac{1}{2} \|\mu\|_2^2 - \sum_{i=1}^d \log \sigma_i - \frac{d}{2} \tag{17}$$

where d is the dimensionality of σ .

B.3. Full Derivation Details about EM Algorithm

We provide the derivation details about EM algorithm in Eq (7) and Eq (8) as follows:

E-step:

$$\begin{aligned}
 & \text{Sample } x_n \sim P_{data}, z_n \sim R_\psi(z|x_n) \text{ and } z_n \sim Q_0(z), x_n \sim Q_\phi(x|z_n), z'_n \sim R_\psi(z|x_n) \\
 \mathcal{L}_t &= \frac{\sum_{n=1}^N w(x_n, z_n) \mathcal{F}(R_\psi(z|x), \phi)}{\sum_{n=1}^N w(x_n, z_n)} \\
 &= \sum_{n=1}^N w'(x_n, z_n) \left\{ E_{R_\psi(z|x)} [\log Q_\phi(x, z) - \log R_\psi(z|x)] + D_{KL}(R_\psi(z|x) || Q_\phi(z|x)) \right\} \\
 &\geq \sum_{n=1}^N w'(x_n, z_n) \left\{ E_{R_\psi(z|x)} [\log Q_\phi(x, z) - \log R_\psi(z|x)] \right\} \\
 &= \sum_{n=1}^N w'(x_n, z_n) \left\{ E_{R_\psi(z|x)} [\log Q_\phi(x|z) + \log Q_0(z) - \log R_\psi(z|x)] \right\} \\
 &= \sum_{n=1}^N w'(x_n, z_n) \left\{ E_{R_\psi(z|x_n)} [\log Q_\phi(x_n|z) + \log Q_0(z) - \log R_\psi(z|x_n)] \right\} \\
 &= \sum_{n=1}^N w'(x_n, z_n) \left\{ \log Q_\phi(x_n|z'_n) - D_{KL}(R_\psi(z|x_n) || Q_0(z)) \right\} \\
 &= \sum_{n=1}^N w'(x_n, z_n) \left\{ \log Q_\phi(x_n|z'_n) - \left(\frac{1}{2} \|\sigma_\psi(x_n)\|_2^2 + \frac{1}{2} \|\mu_\psi(x_n)\|_2^2 - \sum_{i=1}^d \log \sigma_{\psi,i}(x_n) \right) \right\} \\
 &= \tilde{\mathcal{L}}_t
 \end{aligned} \tag{18}$$

where $w(x_n, z_n) = \frac{P_\theta(x_n|z_n)}{Q_{\phi^{(t)}}(x_n|z_n)} P(\mathcal{S}|x_n)$ is the unnormalized importance weight, $w'(x_n, z_n) = \frac{w(x_n, z_n)}{\sum_{n=1}^N w(x_n, z_n)}$ is the normalized one, and N is the sample size. We use two techniques to generate samples for z and x – using the real protein sequences x with generated z from $R_\psi(z|x)$ and using z from a standard normal distribution with generated proteins from $Q_\phi(x|z)$. Since $R_\psi(z|x)$ is assumed to be a normal distribution with its mean μ_ψ and variance $\sigma_\psi^2 \mathbf{I}$ calculated from a Transformer encoder $\mu_\psi, \sigma_\psi = \text{Transformer}_\psi(x_n)$, the KL term can be calculated in close form (Eq. (17)). d is the dimensionality of $\sigma_\psi(x_n)$ and $\sigma_{\psi,i}(x_n)$ is the i -th dimension of $\sigma_\psi(x_n)$. $\tilde{\mathcal{L}}_t$ is the lower bound of \mathcal{L}_t .

M-step:

$$\begin{aligned}
 \psi^{(t+1)}, \phi^{(t+1)} &= \underset{\psi, \phi}{\operatorname{argmax}} \tilde{\mathcal{L}}_t \Rightarrow \psi^{(t+1)} = \psi^{(t)} + r * \nabla_\psi \tilde{\mathcal{L}}_t, \phi^{(t+1)} = \phi^{(t)} + r * \nabla_\phi \tilde{\mathcal{L}}_t \\
 \nabla_\psi \tilde{\mathcal{L}}_t &= \sum_{n=1}^N w'(x_n, z_n) \left\{ -\nabla_\psi \left(\frac{1}{2} \|\sigma_\psi(x_n)\|_2^2 + \frac{1}{2} \|\mu_\psi(x_n)\|_2^2 - \sum_{i=1}^d \log \sigma_{\psi,i}(x_n) \right) \right\} \\
 \nabla_\phi \tilde{\mathcal{L}}_t &= \sum_{n=1}^N w'(x_n, z_n) \left\{ \nabla_\phi \log Q_\phi(x_n|z'_n) \right\}
 \end{aligned} \tag{20}$$

r is the learning rate. $\nabla_\psi \mathcal{L}_t$ and $\nabla_\phi \mathcal{L}_t$ are the gradients of ψ and ϕ at the t -th iteration, respectively.

B.4. IsEM-Pro Algorithm

Algorithm 1 IsEM-Pro Model Learning

Input: ϵ : separately learned combinatorial structure features through MRFs

$P_\theta(x, z; \epsilon)$: a pre-trained protein sequence probability model (VAE's decoder augmented with ϵ)

T: number of iteration for importance sampling based EM learning

$\mathcal{X}_{\text{data}}$: real protein sequences

N_{data} : number of real protein sequences

r : learning rate

\mathcal{B} : batch size

d : dimensionality of the latent variable

Output: Final proposal model $Q_{\phi^{(T)}}(x|z; \epsilon)$ and variational distribution $R_{\psi^{(T)}}(z|x)$

```

1: set  $Q_{\phi^{(0)}}(x|z; \epsilon) = P_\theta(x|z; \epsilon)$ ,  $R_{\psi^{(0)}}(z|x) = \tilde{P}_\theta(z|x)$ 
2: for t=0 to T-1 do
3:    $D_1 = \emptyset, D_2 = \emptyset$ 
4:   for each  $x_n \sim \mathcal{X}_{\text{data}}$  do
5:     sample  $z_n \sim R_{\psi^{(t)}}(z|x_n) = \mathcal{N}(\mu_\psi, \sigma_\psi^2 \mathbf{I})$ 
6:      $D_1 = D_1 \cup \{(x_n, z_n)\}$ 
7:   end for
8:   for i=1 to  $(N_{\text{data}} * 0.1)$  do
9:     sample  $z_i \sim \mathcal{N}(0, \mathbf{I})$ 
10:    sample  $x_i \sim Q_{\phi^{(t)}}(x|z_i; \epsilon)$ 
11:     $D_2 = D_2 \cup \{(x_i, z_i)\}$ 
12:  end for
13:   $D = D_1 \cup D_2, N = |D|$ 
14:  for minibatch  $\{(x_n, z_n)\}_{n=1}^{\mathcal{B}} \subset D$  do
15:     $w(x_n, z_n) = \frac{P_\theta(x_n|z_n; \epsilon)}{Q_{\phi^{(t)}}(x_n|z_n; \epsilon)} P(S|x_n), w'(x_n, z_n) = \frac{w(x_n, z_n)}{\sum_{n=1}^{\mathcal{B}} w(x_n, z_n)}$ 
16:    sample  $z'_n \sim R_{\psi^{(t)}}(z|x_n) = \mathcal{N}(\mu_\psi, \sigma_\psi^2 \mathbf{I})$ 
17:     $\tilde{\mathcal{L}}_t = \sum_{n=1}^{\mathcal{B}} w'(x_n, z_n) \left\{ \log Q_{\phi^{(t)}}(x_n|z'_n; \epsilon) - \left( \frac{1}{2} \|\sigma_{\psi^{(t)}}(x_n)\|_2^2 + \frac{1}{2} \|\mu_{\psi^{(t)}}(x_n)\|_2^2 - \sum_{i=1}^d \log \sigma_{\psi^{(t)}, i}(x_n) \right) \right\}$ 
18:    update  $\psi^{(t)} = \psi^{(t)} + r * \nabla_\psi \tilde{\mathcal{L}}_t$ 
19:    update  $\phi^{(t)} = \phi^{(t)} + r * \nabla_\phi \tilde{\mathcal{L}}_t$ 
20:  end for
21:  update  $\psi^{(t+1)} = \psi^{(t)}$ 
22:  update  $\phi^{(t+1)} = \phi^{(t)}$ 
23: end for

```

C. Approximate KL Divergence

C.1. Proof of the Unbiased and Low-Variance Estimator

Letting $r(x) = \frac{P_\theta(x|S)}{Q_\phi(x)}$, we have:

$$E_{Q_\phi(x)}[(r(x) - 1) - \log r(x)] = E_{Q_\phi(x)}\left[\log \frac{Q_\phi(x)}{P_\theta(x|S)}\right] = D_{KL}(Q_\phi(x) || P_\theta(x|S)) \quad (21)$$

Therefore, this estimator for KL divergence is unbiased.

Let $y = r(x)$ and $f(y) = (y - 1) - \log y$, since $f(y)$ is a convex function and it achieves the minimum value when $y = 1$, we have:

$$(y - 1) - \log y \geq f(1) = 0 \quad (22)$$

Thus, $(r(x) - 1) - \log r(x)$ is always larger than or equals to 0. Instead, in the original KL divergence, $\log \frac{Q_\phi(x)}{P_\theta(x|S)} = -\log r(x)$ would be negative for half of the samples. Therefore, $E_{Q_\phi(x)}[(r(x) - 1) - \log r(x)]$ has lower variance compared to the original one.

C.2. Theoretical Understanding

We can prove that under acceptable KL divergence, the samples from the proposal distribution $Q_\phi(x)$ can be bounded within a reasonable sampling error with samples from the posterior distribution $P_\theta(x|\mathcal{S})$.

Lemma C.1. *If the KL divergence between two distributions P and Q is less than a small positive value δ , then the sampling probability difference between P and Q will be bounded by $\sqrt{2\delta}$ for each sample.*

Proof. Let $\delta(P_\theta(x|\mathcal{S}), Q_\phi(x))$ be the total variation distance between $P_\theta(x|\mathcal{S})$ and $Q_\phi(x)$. We have:

$$\frac{1}{2} \sum_x |P_\theta(x|\mathcal{S}) - Q_\phi(x)| = \delta(P_\theta(x|\mathcal{S}), Q_\phi(x)) \leq \sqrt{\frac{1}{2} D_{KL}(P_\theta(x|\mathcal{S}) || Q_\phi(x))} < \sqrt{\frac{1}{2} \delta} \quad (23)$$

The first inequality is due to Pinsker’s inequality (Csiszár & Körner, 2011), which is tight if and only if $P_\theta(x|\mathcal{S}) = Q_\phi(x)$. Then there is no difference between sampling from $P_\theta(x|\mathcal{S})$ and $Q_\phi(x)$.

From the above analysis, we can get:

$$\sum_x |P_\theta(x|\mathcal{S}) - Q_\phi(x)| < \sqrt{2\delta} \quad (24)$$

When δ approaches 0, the sampling difference between $P_\theta(x|\mathcal{S})$ and $Q_\phi(x)$ would be very minor.

D. Case Study

Figure 6 illustrates the complete sequence and secondary structure analyse of our designed protein of avGFP compared with Cytochrome b562 integral fusion with enhanced green fluorescent protein (EGFP). From the figure, we can see that there are much overlap between our designed protein sequence and the chain B of Cytochrome b562 integral fusion with EGFP. It gives empirical evidence that the green fluorescent protein generated by our model is highly likely to be a real protein compared with the proteins we already know. But whether the designed sequences can accelerate wet-lab experiments still need more exploration as we can not 100% trust it.

E. Pseudo-Likelihood for Combinatorial Structure Learning

We train the Markov random fields using a pseudo-likelihood as CCMpred (Seemayer et al., 2014) additionally combined with the L1 regularization and L2 regularization. The pseudo-likelihood is given in the following equation:

$$\begin{aligned} \hat{P}_\epsilon(x) &= \log \prod_{i=1}^M P_\epsilon(x_i | x_1, x_2, \dots, x_{i-1}, x_{i+1}, \dots, x_M, \epsilon) \\ &= \sum_{i=1}^M \log \frac{\exp(\epsilon_i(x_i) + \sum_{j=1, j \neq i}^M \epsilon_{ij}(x_i, x_j))}{\sum_{c \in \mathcal{V}} \exp(\epsilon_i(c) + \sum_{j=1, j \neq i}^M \epsilon_{ij}(c, x_j))} \\ &= \sum_{i=1}^M \{ \epsilon_i(x_i) + \sum_{j=1, j \neq i}^M \epsilon_{ij}(x_i, x_j) - \log Z_i \} \\ Z_i &= \sum_{c \in \mathcal{V}} \exp(\epsilon_i(c) + \sum_{j=1, j \neq i}^M \epsilon_{ij}(c, x_j)) \end{aligned} \quad (25)$$

where \mathcal{V} denotes the vocabulary of 20 amino acids.

F. Additional Experiments

F.1. Conditional VAE

We provide the results of conditional VAE (CVAE) in Table 6. We implement the CVAE which is built by adding the fitness condition on deepsequence (Riesselman et al., 2018). The table shows CVAE exhibits decreased performance. It demonstrates that iterative sampling inside EM algorithm with explicit fitness constraints can lead to better proteins.

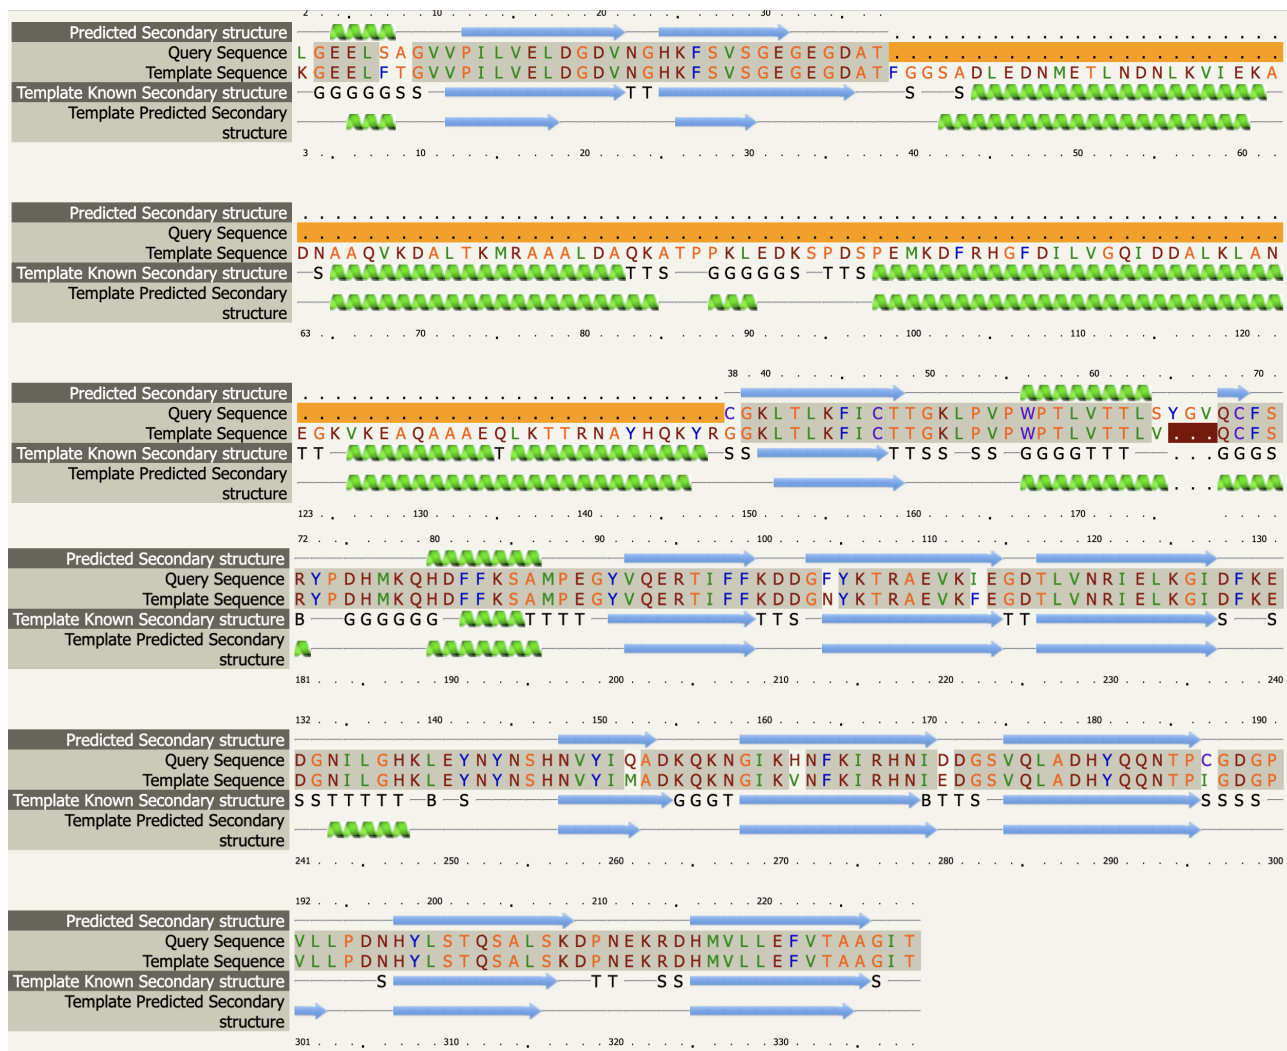


Figure 6. Complete sequence and secondary structure comparison between the designed sequence of avGFP and chain B of Cytochrome b562 integral fusion with enhanced green fluorescent protein. Green and blue parts respectively represent the Alpha helix and Beta strand.

F.2. Non-Autoregressive Decoder

We present a performance comparison between our IsEM-Pro and a fully-connected non-autoregressive decoder in Table 7. The non-autoregressive model, a VAE, includes a 6-layer Transformer encoder followed by a linear mapping to learn the mean and variance of the latent variable, and a two-layer non-autoregressive Transformer decoder. The embedding and feed-forward network dimensions are set to 320 and 1280, respectively. The model uses sampled latent vector as decoder input for all positions and is trained using the glancing strategy (Qian et al., 2021). The results show that the non-autoregressive model struggles to generate high-quality candidates for proteins with longer sequences, such as LGK, which has 439 amino acids.

Dataset	avGFP	AAV	TEM	E4B	AMIE	LGK	Pab1	Average	
CVAE	3.570	4.128	0.083	−0.819	−1.564	−0.796	1.638	3.492	1.217
IsEM-Pro	6.185	4.813	1.850	5.737	0.062	0.035	2.923	4.536	3.267

Table 6. Fitness of conditional VAE compared with our IsEM-Pro.

Dataset	avGFP	AAV	TEM	E4B	AMIE	LGK	Pab1	Average	
Non-autoregressive Decoder	1.312	3.137	1.015	3.096	−3.297	−3.125	1.548	2.694	0.797
IsEM-Pro	6.185	4.813	1.850	5.737	0.062	0.035	2.923	4.536	3.267

Table 7. Fitness of a non-autoregressive decoder based VAE model compared with our IsEM-Pro.