
MSAGPT: Neural Prompting Protein Structure Prediction via MSA Generative Pre-Training

Bo Chen^{1*†}, Zhilei Bei^{1*†}, Xingyi Cheng², Pan Li², Jie Tang¹, Le Song^{2,3}

¹Tsinghua University ²BioMap Research ³MBZUAI

cb21@mails.tsinghua.edu.cn

<https://github.com/THUDM/MSAGPT>

Abstract

Multiple Sequence Alignment (MSA) plays a pivotal role in unveiling the evolutionary trajectories of protein families. The accuracy of protein structure predictions is often compromised for protein sequences that lack sufficient homologous information to construct high-quality MSA. Although various methods have been proposed to generate virtual MSA under these conditions, they fall short in comprehensively capturing the intricate co-evolutionary patterns within MSA or require guidance from external oracle models. Here we introduce MSAGPT, a novel approach to prompt protein structure predictions via MSA generative pre-training in the low-MSA regime. MSAGPT employs a simple yet effective 2D evolutionary positional encoding scheme to model the complex evolutionary patterns. Endowed by this, its flexible 1D MSA decoding framework facilitates zero- or few-shot learning. Moreover, we demonstrate that leveraging the feedback from AlphaFold2 can further enhance the model's capacity via Rejective Fine-tuning (RFT) and Reinforcement Learning from AF2 Feedback (RLAF). Extensive experiments confirm the efficacy of MSAGPT in generating faithful virtual MSA to enhance the structure prediction accuracy (up to +8.5% TM-Score on few-shot scenarios). The transfer learning capabilities also highlight its great potential for facilitating other protein tasks.

1 Introduction

The advent of deep learning has significantly propelled progress across various scientific domains, exemplified by breakthroughs such as AlphaFold series [1, 2] for accurate biomolecular interaction predictions, AlphaGeometry [3] for intricate geometry and mathematical reasoning—to name a few. Among these, AlphaFold2 (AF2) represents a landmark within structural biology, achieving an *in silico* precision of approximately 90% atomic accuracy that rivals wet lab experiments on protein structure predictions (PSP). The remarkable success of AF2 can be attributed to its innovative end-to-end use of co-evolutionary information supported by Multiple Sequence Alignment (MSA). MSA aggregates homologous sequences from vast databases, providing a comprehensive overview of evolutionary trajectories, which is critical for accurately predicting protein structures [1, 2, 4]. An illustrative example in Figure 1(a) showcases that the correlations analysis among amino acids (AAs) sites could reveal contacts or conservative regions in the folding structure. Unfortunately, not all proteins possess a rich set of homologous counterparts. Statistical investigations reveal that approximately 20% of metagenomic proteins [5] and around 11% of proteins from eukaryotic and viral origins [6] are classified as "orphan" proteins. This presents a significant challenge for MSA-search algorithms in constructing high-quality MSA, consequently impeding the performance of PSP models [2].

^{*}BC and ZB contributed equally.

[†]Work done while interned at BioMap.




Figure 1: (a) The illustration of MSA and (b) performance comparisons between MSAGPT and advanced baselines on three natural MSA-scarce benchmark.

Drawing on the impressive capabilities of large language models endowed either by the autoencoding [7] or the autoregressive language modeling regime [8, 9], protein language models (PLMs) have been developed to unveil the evolutionary patterns and sequence characteristics intrinsic to protein structures. Specifically, generative PLMs [10, 11, 12], trained on vast protein databases [13, 14, 15, 16] have achieved unparalleled success in generating novel proteins with desired structural properties. These achievements underscore the efficacy of language models in identifying evolutionary patterns within individual protein sequences. Inspired by this, subsequent works [17, 18] attempt to further integrate MSA as the input or by directly generating virtual yet informative MSA [19, 20, 21] to provide additional evolutionary insights. These approaches usually adopt customized attentions that merely allow attention aggregated among specific directions, such as axial attention [22], for separately analyzing the row- and column-wise co-evolutionary patterns in MSA. However, these attention mechanisms usually have low efficiency in capturing the evolutionary information in MSA, or even fail to adequately capture intricate co-evolutionary dynamics. Taking Figure 1(a) as an example, it is imperative to concurrently analyze the pairwise or high-order relationships of amino acid sites across all homologs to deduce the structural constraints influencing the folding structures, which may not be achieved by customized attention. The limited capacity may result in compromised performance on the task that highly resorts to co-evolutionary information.

Built upon the insights mentioned above, we introduce MSAGPT, a simple yet effective framework that prompts protein structure prediction via MSA generative pre-training. This method facilitates *de novo* MSA generation, aiding in protein structure prediction in scenarios with limited MSA available. MSAGPT is characterized by its unique features:

- **2D Evolutionary Positional Encoding.** We employ an innovative dual-axis positional encoding scheme that captures column- and row-wise co-evolutionary information concurrently. This method provides a comprehensive understanding of complex evolutionary relationships with high efficacy, enhancing the model’s generative capabilities.
- **1D Zero-/Few-Shot MSA Decoding.** With 2D positional encoding, MSAGPT re-formalizes MSA generation as a one-dimensional sequence generation task, optimized by the simple next-token-prediction objective. This enables MSAGPT to conduct zero- or few-shot MSA generation under a flexible in-context learning framework.
- **Learning from AlphaFold2 Feedback.** MSAGPT further utilizes feedback from AlphaFold2 to reduce hallucinations during MSA generation. This approach ensures the generation of reliable and informative MSA, thus enhancing protein structure prediction.

Extensive experiments conducted on three benchmarks, CAMEO [23], CASP, and PDB [14], demonstrate the superior capacity of MSAGPT in generating high-quality MSA. Notably, MSAGPT outperforms existing MSA generation models on both zero- and few-shot scenarios. Impressively, AF2 with virtual MSA generated by MSAGPT significantly improves the structure prediction accuracy than that with natural MSA on cases with limited homologous information. Moreover, the subsequent Rejective Fine-tuning (RFT) and learning from AF2 feedback (RLAF) stage further enhance the model’s ability to generate informative and faithful MSA, outperforming the original MSAGPT by a large margin, as shown in Figure 1(b). Additionally, we demonstrate that virtual MSA can also benefit other tasks.




Figure 2: **The overall framework of prompting protein structure predictions via MSA generation.** **Left:** The challenge faced by conventional search algorithms on protein with scarce homologous sequences, resulting in suboptimal alignments. **Middle-to-Right:** MSAGPT generates informative and high-quality MSA for such challenging queries, presenting a promising approach to overcoming these limitations. [M] denotes the sequence separator. [S], [E] are the special tokens to represent the start or end of MSA generation.

We expect MSAGPT to become integral in supplementing protein-related tasks requiring critical evolutionary information from MSA. The model is available at <https://github.com/THUDM/MSAGPT>.

2 Related work

Protein Structure Prediction. Proteins are fundamental to the various biological processes that sustain, grow, and protect living organisms. Groundbreaking deep learning approaches [1, 2, 4] have been developed to predict the folding structures based on their sequences. These methods have achieved structure prediction accuracy to conventional wet-lab experiments. The success largely relies on the utilization of MSA, which are retrieved through search algorithms [24, 25, 26, 27] against vast databases [13, 14, 15, 16]. However, challenges arise with “orphan” protein sequences, which lack sufficient homologous sequences for accurate structure prediction. Single-sequence PSP methods [11, 28, 29, 30] are designed to infer folding structures directly from the query protein sequences. Despite these advancements, the accuracy of predictions from single-sequence methodologies generally falls short of those derived from MSA-based algorithms.

Protein Language Models. Protein Language Models (PLMs), such as ESM [28, 31], ProGen [10, 32], etc [12, 33, 34] have emerged as a groundbreaking development in computational biology. PLMs are trained on single sequences, towards understanding protein structural features or enabling the generation of diverse and realistic protein sequences. MSA Transformer [17] further incorporates MSA as the input, achieving better performance than these single sequence models, underscoring the importance of utilizing the evolutionary information from MSA [35, 36, 37]. To enhance MSA generation, MSA-Augmentor [20], PoET [19] employ the seqs2seqs pre-training, which adopts the sequential axial attention mechanism to capture the evolutionary data across and within the input sequences, both horizontally and vertically. EvoGen [21], serving as the meta generative model, aims at producing virtual MSA for enhancing protein structure predictions. However, it highly resorts to external structural prediction models to optimize its performance.

Aligning with Human Preferences. Aligning language models with human preferences has been shown to be effective in improving generation quality [8, 38, 39, 40]. Existing methods typically employ supervised fine-tuning using human-annotated datasets or reinforcement learning from human feedback pipelines [38, 39]. Inspired by these, we utilize the feedback from AlphaFold2 to further enhance the generation capability of the pre-trained model, which helps mitigate hallucinations and enables the model to generate accurate and reliable MSA.

3 Preliminary

Definition 1 Multiple Sequence Alignment (MSA). Given the query protein sequence $Q \in \mathbb{A}^{1 \times L}$, where \mathbb{A} denotes the set of alphabetic symbols used to represent the 20 basic amino acids and L

represents the number of amino acids per sequence, the MSA $M \in \mathbb{A}^{N \times L}$ of Q is comprised of N homologous protein sequences, which can be obtained either by searching over protein databases or generating with MSA generation methods.

Problem 1 Prompting Protein Structure Prediction by MSA Generation. Given Q with initial MSA $M_{init} \in \mathbb{A}^{n \times L}$ as the prompt, where $n = 0$ indicates the zero-shot MSA generation and $n > 0$ signifies the few-shot MSA generation, we target at learning a function f to generate virtual MSA $M_{gen} \in \mathbb{A}^{m \times L}$ based on Q and M_{init} , such that the structure prediction accuracy based on the augmented MSA $M_{aug} \in \mathbb{A}^{(n+m) \times L}$ significantly surpasses that based on the initial MSA M_{init} ,

$$M_{aug} = f(Q, M_{init}),$$

$$\mathbb{I}_{acc}(Q, M_{aug}) > \mathbb{I}_{acc}(Q, M_{init})$$

where the \mathbb{I}_{acc} is prediction accuracy comparing the prediction result of AF2 and the ground truth.

In this paper, we mainly focus on improving the structure prediction accuracy in the low-MSA regime, i.e., the cases that lack a sufficient number of homologous sequences.

4 Methodology

Given a query sequence and its retrieved natural MSA, we aim to comprehensively understand the co-evolutionary patterns in MSA, such that we can generate informative virtual MSA for prompting protein structure prediction in the low-MSA regime. Conceptually, the co-evolutionary information is analogous to the covariance matrix in mathematics, depicting the correlations among amino acids by comparing pairwise or high-order correlations among amino acid sites in MSA, as depicted in Figure 1(a). To achieve this goal, MSAGPT contains two key adoptions, distinguishing it from existing MSA-based PLMs that rely on customized attentions [2, 17, 20, 19]: **2D Evolutionary Positional Encoding.** Introduces an adaptive dual-axis positional encoding scheme that captures column- and row-wise co-evolutionary information concurrently. And **1D Zero-/Few-Shot MSA Decoding.** Re-formalizes MSA generation as a one-dimensional sequence generation task based on the proposed 2D positional encoding scheme, which enables MSAGPT to conduct zero- or few-shot context learning MSA generation framework. The overall framework is illustrated in Figure 2.

4.1 2D Evolutionary Positional Encoding

Vanilla transformers typically use 1D positional embeddings to incorporate absolute and relative positional information of tokens. However, when dealing with MSA, which represents stacked homologs, the structure is different. Each row of MSA corresponds to a distinct protein sequence, while each column represents the evolutionary trajectories of a specific amino acids (AAs) site. To effectively capture the evolutionary patterns, recent approaches [2, 17, 20] have employed decoupled axial attentions, which are designed to capture explicit co-evolutionary information along the rows (protein sequences) and columns (AAs sites). However, these methods often suffer from low efficiency in capturing the information dynamics or fail to capture the evolutionary information adequately.

In light of this, we introduce a novel two-dimensional evolutionary positional encoding scheme, illustrated in Figure 2. Given an MSA $M \in \mathbb{A}^{N \times L}$, we define a 2D positional id matrix $\mathbf{P} \in \mathbb{N}^{2 \times N \times L}$, where the first positional id $\mathbf{P}_0 \in \mathbb{N}^{1 \times N \times L}$ indicates the column position, i.e., $P_0[i, \cdot] = \{0, 1, \dots, L\}$, and the second positional id \mathbf{P}_1 indicates the row position, i.e., $P_1[j, \cdot] = \{0, 1, \dots, N\}$. The two positional ids are projected into two vectors added to the input token embeddings. We utilize the Rotary Positional Encoding (RoPE) [41] technique, specifically adapting its two-dimensional variant* to suit our 2D positional encoding requirements.

Comparison with Axial Attentions. Considering the 2D positional id (P_0, P_1) , the self-attention among AAs (α, β) meets the following unit patterns, as illustrated in Figure 3:

- $P_0^\alpha = P_0^\beta \ \& \ P_1^\alpha \neq P_1^\beta$. Indicates α and β reside in the same site across different protein sequences, such as the AA pair (A, K) and (P, G), enabling column-wise self-attention that highlights evolutionary patterns conserved across sequences.

*<https://kexue.fm/archives/8397>

- $P_0^\alpha \neq P_0^\beta$ & $P_1^\alpha = P_1^\beta$. Suggests α and β are positioned in the same protein sequence but at different sites, such as the AA pair (A, P) and (K, G), facilitating row-wise self-attention that captures sequence-specific features.

- $P_0^\alpha \neq P_0^\beta$ & $P_1^\alpha \neq P_1^\beta$. Denotes α and β lack explicit correlation, such as the AA pair (A, G) and (P, K), may be serving as anchor nodes for complex co-evolutionary information diffusion.

Conceptually, the 2D positional encoding encapsulates the explicit row- and column-wise self-attention patterns with high efficacy. Moreover, it allows unrestricted information diffusion, that is, enabling any two amino acids to attend to one another. Such a framework facilitates unveiling complex co-evolutionary patterns, such as high-order correlations among AAs, that customized self-attentions might overlook.

4.2 1D Zero-/Few-Shot MSA Decoding

Leveraging the 2D evolutionary positional encoding, we further release the stacked MSA decoding task into the scalable 1D sequence generation framework, without compromising the integrity of co-evolutionary information. Specifically, we convert the MSA $M \in \mathbb{A}^{N \times L}$ into the flattened 1D sequence $M^f \in \mathbb{A}^{1 \times NL}$ along the row axis to ensure that we can generate the MSA sequentially during inference. Similarly, the 2D positional id matrix $P \in \mathbb{N}^{2 \times N \times L}$ is reshaped into a flattened format, $P^f \in \mathbb{N}^{1 \times 2 \times NL}$. This allows the model to conduct a simple auto-regressive generation process, as illustrated in Figure 2.

Discussions. Admittedly, introducing 2D positional encoding introduces higher time complexity in comparison to conventional customized attention mechanisms (from $O(N^2L) + O(NL^2)$ to $O(N^2L^2)$). However, it is worth noting that the original stacked nature of MSA poses challenges for integrating it with acceleration techniques used in large language models, such as Flash Attention [42, 43]. The 1D decoding framework, conversely, can be easily scaled to accommodate in-context learning frameworks, such as retrieval augmented generation, to further enhance the model’s generation capability and expand its application range. From a practical standpoint, the high parallelism of the 1D decoding framework demonstrates superior inference speed, benefiting from techniques like Flash Attention and KV-cache, while incurring negligible memory overhead compared to customized attention mechanisms. For further details, please refer to Appendix Section A.4.

5 The Training Pipeline of MSAGPT

The training pipeline involves three successive stages: **Stage 1: MSA Generative Pre-Training** to obtain the base MSA generation model; **Stage 2: Rejective Fine-tuning (RFT)** to instruct the base model with high-quality MSAs via AF2 annotations, which can reduce generation hallucinations ; **Stage 3: Reinforcement Learning from AlphaFold2 Feedback (RLAF)** to further enhance RFT model’s capabilities based on the feedback of AF2. (See Appendix Section A for training details.)

5.1 Stage 1: MSA Generative Pre-Training

Pre-Training Dataset. We utilize the Uniclust30 MSA dataset from OpenProteinSet [44], which is processed through an all-against-all search on Uniclust30 [45] using HHblits [46]. This results in approximately 16 million MSAs (See Appendix A.1 for Details).

Pre-training Objective. We adapt the language modeling objective [47] to the MSA generation task. The cross-entropy loss for modeling the intrinsic distribution of MSA $M^f \in \mathbb{A}^{1 \times NL}$ is defined as:

$$L_{ce} = \mathbb{E}_{M^f} \left[\sum_{i=0}^{N \times L} -\log p(M_i^f | M_{<i}^f, \theta) \right] \quad (1)$$

where $\mathbf{M}^f \in \mathbb{A}^{1 \times NL}$ is 1D flattened version of the input MSA and θ is the learned parameter.

5.2 Stage 2: Rejective Fine-tuning (RFT)

Noted that the pre-trained dataset inevitably contains noisy co-evolutionary patterns, such as large portions of deletions and insertions, which may mislead the base model to yield hallucinated cases, i.e., the linguistically reasonable but intrinsically unfaithful MSA. Thus we select highly-quality MSAs to further fine-tune the base model via a rejective sampling procedure based on the AF2-annotation.

RFT Dataset. We collect 120,780 protein sequences with structures from Protein Data Bank (PDB) [14]. For the sequence Q , we search its MSA $M \in \mathbb{A}^{N \times L}$ from UniClust30 [45] with HHblits [46]. Then we sample several MSA subsets $\mathbf{m} = \{m_1, m_2, \dots, m_i\}$ with replacement, where $m_i \in \mathbb{A}^{n \times L}$ and $n \ll N$. To assure the information density of the sampled data, we filter out the MSA with depth N fewer than $\lceil n \times i/2 \rceil$. Subsequently, we employ AF2 to score the sampled subset using the structure prediction accuracy $\mathbb{I}_{acc}(Q, m_i)$. Then the RFT dataset \mathcal{D}_{RFT} is defined as:

$$\mathcal{D}_{RFT} = \{(Q, m_i) | (\mathbb{I}_{acc}(Q, m_i)) > \theta_1 \cap (\mathbb{I}_{acc}(Q, m_i) - \mathbb{I}_{acc}(Q, -)) > \theta_2\} \quad (2)$$

where $\mathbb{I}_{acc}(Q, -)$ indicates the prediction accuracy without using MSAs. We set the sampling number $i = 10$, the depth of sampled MSA $n = 16$, $\theta_1 = 0.9$, and $\theta_2 = 0.2$, which results in approximately 60k samples. The base model is fine-tuned on \mathcal{D}_{RFT} with the same pre-training objective.

5.3 Stage 3: Reinforcement Learning from AlphaFold2 Feedback (RLAF)

We further employ AF2 as the reward model to perform the Reinforcement Learning with AF2 Feedback (RLAF) using Direct Preference Optimization [38] (DPO) to further guide the RFT model to decode meaningful structure-related MSA patterns that align with the preference of AF2.

RLAF Preference Dataset. For each query Q from the PDB, we use the RFT model to generate its MSA $M \in \mathbb{A}^{N \times L}$ in zero-shot manner. Then, we also sample several MSA subsets $\mathbf{m} = \{m_1, m_2, \dots, m_i\}$ and obtain the preference dataset $\mathcal{D}_{DPO} = \left\{Q^{(k)}, m_w^{(k)}, m_l^{(k)}\right\}_{k=1}^K$ as follows,

$$\mathcal{D}_{DPO} = \{(Q, m_w, m_l) | (\mathbb{I}_{acc}(Q, m_w) - \mathbb{I}_{acc}(Q, m_l)) > \theta_3\} \quad (3)$$

where we set the $\theta_3 = 0.3$, rendering the number of preference data $D_{DPO} = 11k$.

RLAF Training Objective. The adapted DPO loss is defined as:

$$L_{DPO} = \mathbb{E}_{(Q, m_w, m_l) \in \mathcal{D}_{DPO}} \left[-\log \sigma \left(\beta \log \frac{\pi_\theta(m_w|Q)}{\pi_{ref}(m_w|Q)} - \beta \log \frac{\pi_\theta(m_l|Q)}{\pi_{ref}(m_l|Q)} \right) \right] \quad (4)$$

where π_θ and π_{ref} are initialized by the RFT model and π_{ref} is frozen while π_θ is optimized. During the RLAF training phase, we found that merely using the DPO loss led to training instability. Thus we adopt the pre-training loss L_{CE} for the chosen answer m_w as a regularization term with the coefficient factor λ in the total loss to mitigate this issue. The total loss $L = L_{DPO} + \lambda L_{CE}$, $\lambda = 0.1$. Another critical coefficient β , which measures the penalty intensity for incorrect answers is set to $\beta = 0.1$.

6 Experiments

6.1 Setup

Benchmarked Dataset. We employ the datasets from CAMEO [23], CASP14&15, and PDB [14], which are esteemed benchmarks in protein structure analysis spanning a diverse array of biological protein families. For each protein sequence, we search its MSA on UniClust30 database [45] using HHblits [46]. Given our focus on addressing the challenge presented by cases with limited MSA information, we establish two specific benchmarks to represent the MSA-scarce conditions:

Table 1: The performance of structure prediction on three natural MSA-scarce benchmarks. avg. Depth represents the average depth of searched MSA across all query sequences. Compared with the base model, the RFT and DPO models achieve higher TM-Score while with lower pLDDT values. (See Appendix Table 5 for more results.)

Model	CAMEO (avg. Depth = 8.5)				CASP (avg. Depth = 4.6)				PDB (avg. Depth = 2.6)			
	Zero-Shot		Few-Shot		Zero-Shot		Few-Shot		Zero-Shot		Few-Shot	
	pLDDT	TM	pLDDT	TM	pLDDT	TM	pLDDT	TM	pLDDT	TM	pLDDT	TM
AF2 MSA	63.8	55.4	77.4	71.4	44.0	32.6	54.2	44.1	55.2	45.6	61.0	52.3
EvoDiff	68.0	59.1	70.1	60.1	45.0	29.5	50.6	38.1	54.8	45.2	59.1	47.2
MSA-Aug.	67.7	59.2	77.4	72.1	56.8	36.6	63.4	46.3	61.9	49.8	66.0	55.3
EvoGen	66.1	60.3	78.6	75.3	48.2	38.4	55.1	48.5	57.6	49.5	62.8	55.4
MSAGPT	70.8	61.4	80.8	75.2	59.0	39.8	65.4	51.0	68.6	53.4	71.3	59.6
+ RFT	68.0	60.5	79.8	76.4	56.8	40.2	64.0	53.6	66.8	53.4	70.3	60.1
+ DPO	68.9	62.7	80.2	76.7	54.2	43.7	62.7	57.0	64.5	53.6	68.0	59.7
	(+3.1)	(+2.4)	(+2.2)	(+1.4)	(+2.2)	(+5.3)	(+2.0)	(+8.5)	(+6.7)	(+3.8)	(+5.3)	(+4.7)

Table 2: Zero-shot evaluation on artificial MSA-scarce benchmark (GDT stands for GDT-TS).

Model	pTM	pLDDT TM	GDT	LDDT
AF2 MSA	28.9	69.6	42.0	62.4
EvoDiff	34.7	70.8	44.0	64.1
MSA-Aug.	33.2	74.3	45.8	65.9
EvoGen	31.7	71.6	44.6	64.6
MSAGPT	37.9	80.4	50.7	70.7
+RFT & DPO	37.9	79.7	51.6	71.4
				76.9

Table 3: Evaluation of selection methods.

Model	CAMEO	CASP	PDB
MSAGPT-DPO	76.7	57.0	59.7
+ pLDDT Selection	77.5	57.6	60.5

Table 4: Evaluation on transfer learning.

Model	CtP	SsP	LocP	MIB
w/o Virtual MSA	11.6	66.5	58.3	57.5
w/ Virtual MSA	13.1	69.0	56.4	60.3

Natural MSA-scarce Benchmark. We identify 200 protein sequences with the number of searched MSA fewer than 20 (8 from cameo, 13 from CASP14&15, 179 from PDB).

Artificial MSA-scarce Benchmark. We collect approximately 8k protein sequences based on the PDB released before 2024-01-22 without searching MSA.


All MSA from the test set are removed from the pre-train dataset.

Baselines. To assess the performance of MSAGPT, we adopt AF2 as the benchmark MSA-based PSP algorithm. For MSA generation baselines, we compare **MSAGPT**, its RFT-version and its DPO-version with two advanced MSA generation algorithms: **EvoDiff** [34], which utilizes the diffusion framework for controllable protein sequence generation. Specifically, we use the EvoDiff-MSA for the MSA generation, **MSA-Augmentor** [20], which utilizes a sequences-to-sequences pre-training architecture incorporating an encoder and a decoder based on the axial attention [17] and **EvoGen** [21], which employs a meta generative model framework with customized attention, leveraging guidance from AF2 to refine its MSA generation. As PoET [19] is designed for mutational scanning tasks, we don't take it as the baseline. Additionally, we include the reference model **AF2 MSA**, which utilizes all the searched natural MSA for prediction.


MSA Generation Pipeline. Given that MSAGPT can perform flexible zero- or few-shot MSA generation to accommodate different levels of available evolutionary information, we define two generation settings to evaluate models' performances under varying conditions:

Zero-Shot Generation. MSA generation is conducted using only the query sequence as input, emphasizing the model's ability to infer necessary evolutionary patterns without additional contexts.

Few-Shot Generation. All the searched natural MSA are viewed as the prompt to inform the few-shot MSA generation process. Then the generated MSA, combined with the initial prompts, serves as augmented data for structure predictions.



(a) Effects of MSA Depths



(b) Effects of Selection Methods

Figure 4: The effect of different MSA depths and selection methods. The X-axis indicates the different MSA depths. The Y-axis represents the TM-Score. The dashed line denotes the non-selection baseline.




Figure 5: Ablation study with positional embedding variants.

Evaluation Metric. We use several widely-used metrics to assess structural similarity between predicted structures and ground truth: TM-Score, GDT-TS, and LDDT. Additionally, we include pTM and pLDDT, the corresponding predicted metrics estimated by AF2. All metrics are normalized from 0 to 100 for comparison, with higher scores indicating higher confidence and usually a more accurate prediction (see Appendix Section B for details).

6.2 MSAGPT’s Virtual MSAs Reflect the Co-evolutionary Information

Table 1 and 2 showcase the comparative results in two benchmarks across different baselines. Notably, AF2 MSA, which relies solely on the limited searched MSA without incorporating virtual MSA, exhibits the worst performance. Predictions enhanced with MSA generated by MSA-Augmentor or EvoGen surpass the performance of AF2 MSA. This underscores the critical role of high-quality MSA in enhancing the accuracy of cutting-edge PSP algorithms. Overall, MSAGPT surpasses other advanced baselines by a large margin on both benchmarks, specifically achieving +1.4% improvement on CAMEO, +8.5% on CASP, and +4.7% on PDB, as measured by TM-Score, on the natural MSA-scarce benchmark. This significant improvement demonstrates not only the superior accuracy and effectiveness of MSAGPT but also its robustness in handling cases with noisy or low-quality MSA.

Moreover, compared with the base model, the RFT and DPO models achieve higher golden metric scores, that is, GDT, LDDT, and TM-Score, but with a lower predictive score, that is, the value of pTM and pLDDT. This discrepancy might arise from the presence of highly confident (according to pTM and pLDDT) but lower-scored decoys (according to TM-Score), as observed in [21], indicating that aligning with the preference dataset, which is filtered based on TM-Score, makes the model more inclined to generate truly informative MSA rather than hallucinated ones.

Statistically, MSAGPT effectively improves the prediction accuracy for 91.0% and 88.9% of protein sequences with limited MSA when compared to AF2 MSA on Zero-Shot and Few-shot scenarios, respectively. This significant finding highlights the potential of our MSAGPT framework to uncover and leverage co-evolutionary patterns within biosequences. Notably, we also discuss the scenario with abundant natural MSA in the Appendix Section B.2.

6.3 Rethinking the MSA Selection Strategy

We further study the effect of different depths of virtual MSA, as shown in Figure 4(a). We observe a trend where the relative improvement in structure prediction accuracy decreases as the depth of virtual MSA increases. The accuracy based on MSA with 64 MSA sequences even underperforms those based on only 16 or 32 sequences. We hypothesize that increasing the number of virtual MSA beyond a certain threshold may introduce a dilution effect, where the density of valuable co-evolutionary signals is compromised by the inclusion of the hallucinated generation noise. To alleviate this, we explore MSA selection strategies for filtering out low-quality, noise-inducing sequences while retaining those that contribute positively to the accuracy of structure predictions, as illustrated in Figure 4(b) (See Appendix Section C for details).

1D Sequence Similarity or Diversity Measure. We first arrange MSA by their similarity to the query sequence in descending order. The results reveal that prioritizing MSA based on their high similarity to the query, termed as *static similarity (STA-SIM)*, does not improve prediction

accuracy compared to the non-selection approach (N/A). On the contrary, the *static diversity (STA-DIV)* strategy, which favors MSA with lower similarity rankings, slightly outperforms the baseline, highlighting the importance of sequence diversity in enhancing MSA quality. Moreover, we employ the dynamic approach, initially selecting the most (or least) similar MSA to the query sequence and progressively incorporating additional MSA based on their average similarity to the cumulatively selected set, termed as *dynamic similarity (DYN-SIM)* and *dynamic diversity (DYN-DIV)*.

The results further confirm the advantage of fostering diversity within MSA rather than selecting only the sequences with high similarities to the query sequence. We also inspect the effectiveness of the widely-adopted MSA *trimming (TRIM)* strategy [21], which yields a similar TM-Score to the non-selection baseline, undermining its efficacy in selecting MSA with high quality.

3D Structure Affinity Measure. We assume that the generated sequence with high quality should exhibit structural congruity with the query sequence, thereby emitting strong co-evolutionary signals. To validate this, we rank sequences within MSA by their predicted tertiary structures according to the pTM, a predicted TM score [2], pLDDT, and TM-Score, from highest to lowest. These approaches, especially when guided by the pLDDT score, consistently select high-quality MSA, evidenced by the enhanced TM-Score. We compare the non-selection methods (N/A) and pLDDT selection methods on the three benchmarked datasets on few-shot generation scenarios in Table 3. This confirms our hypothesis that structural similarity plays a crucial role in effective MSA selections.

6.4 Transfer Learning of MSAGPT

Since protein structures largely dictate their functions, the virtual MSA, enhancing structure prediction, should similarly benefit other protein tasks. To validate this, we focus on two protein structural tasks: Contact Prediction (CtP) and Secondary Structural Prediction (SsP) and two protein functional tasks: Localization Prediction (LocP) and Metal Ion Binding (MIB) [11]. We sample 1,000 sequences from each benchmark and conduct 5-fold cross-validation (See Appendix Section B.3 for details).

Results. Table 4 demonstrates that incorporating MSA from MSAGPT consistently surpasses merely using the single sequence on most tasks. Yet, it achieves inferior performance on the LocP task, which agrees with the observation [48] that protein language models may not present scaling behavior on several protein functional or property prediction tasks. Nevertheless, the results show the great potential of MSAGPT to contribute to a wide range of tasks with generated MSA. We are motivated to explore additional transfer tasks to assess MSAGPT’s utility across various domains further.

6.5 Ablation Study

To understand the effect of various positional encoding strategies on capturing co-evolutionary patterns, we design four model variants: **1D_gpt**: Adopts the standard GPT positional encoding; **1D_2nd**: Utilizes only the second-dimensional of the 2D evolutionary positional encoding mechanism; **1D_1st**: Utilizes the first-dimensional positional encoding; **2D_full**: Implements the 2D evolutionary positional encoding mechanism (See Appendix Section B for details).

Results. Figure 5 showcases the TM-score distribution across different model variants. The 1D_gpt exhibits the lowest performance, attributed to its simplistic approach of treating the MSA as a concatenation of homologous sequences, thereby failing to discern any co-evolutionary patterns. Both the 1D_1st and 1D_2nd demonstrate significant improvement over 1D_gpt, by explicitly encoding column- or row-wise relationships within the MSA, respectively. Notably, the performance of 1D_1st is better than that of 1D_2nd, suggesting that column-wise covariance patterns play a more crucial role in structural predictions than row-wise patterns. This aligns with the understanding that the permutation of sequence order does not alter the covariance information among residue sites [17]. Remarkably, the 2D_full variant, which incorporates the proposed 2D evolutionary positional encoding, outperforms all other models, which underscores its effectiveness in capturing the intricate evolutionary information present in MSA.

7 Limitations

In this section, we discuss some limitations that should be resolved in future work.

Scaling behavior of MSAGPT. While we have showcased the effectiveness of MSAGPT in generating informative virtual MSA, it is important to note that our pre-training was conducted with a model containing 2.8 billion parameters. The performance and behavior of MSAGPT, when scaled concerning dataset size, model size, and total compute resources, remain unknown.

Transfer Learning on a wide range of tasks. While we have demonstrated the transferability of MSAGPT on several tasks, including protein structure prediction and protein function prediction, its performance on a broader range of tasks remains an open question. The ability of a model to transfer its learned knowledge and adapt to new tasks is a critical aspect of transfer learning. While MSAGPT has shown promising results on the tasks it was evaluated on, it is important to assess its performance on a more diverse set of tasks spanning various domains and problem types.

8 Border Impact

The aim of this paper is to improve the accuracy of protein structure prediction in cases with limited homologous sequences. The generated MSA also shows great potential to transfer to other protein-related tasks. By leveraging the information encoded in the generated MSAs, it is possible to enhance the performance of various protein-related tasks beyond structure prediction. However, the generative MSA may be misused to contaminate the high-quality nature MSA databases. Thus, it is necessary to train a classifier to distinguish the real from MSAGPT-generated MSA.

9 Conclusion

This paper introduces MSAGPT, a novel approach that prompts protein structure prediction via MSA generative pre-training, to enable accurate protein structure predictions in situations where co-evolutionary information is scarce. To meticulously characterize the co-evolutionary patterns within MSA, MSAGPT designs two innovative techniques: the 2D Evolutionary Positional Encoding scheme and the 1D Zero-/Few-Shot MSA Decoding mechanisms. The post-alignment learning from AlphaFold2 feedback further enhances the quality of MSA generation. Empirical experiments conducted on a variety of benchmarks have demonstrated MSAGPT’s robustness and effectiveness. In the future, we plan to apply MSAGPT to broader areas, particularly for tasks that heavily rely on co-evolutionary information, and investigate the aforementioned limitations.

Acknowledgments. This work has been supported by the NSFC for Distinguished Young Scholar 62425601, New Cornerstone Science Foundation through the XPLORER PRIZE and Tsinghua University Initiative Scientific Research Program.

References

- [1] Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, pages 1–3, 2024.
- [2] 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.
- [3] Trieu H Trinh, Yuhuai Wu, Quoc V Le, He He, and Thang Luong. Solving olympiad geometry without human demonstrations. *Nature*, 625(7995):476–482, 2024.
- [4] Minkyung Baek, Frank DiMaio, Ivan Anishchenko, Justas Dauparas, Sergey Ovchinnikov, Gyu Rie Lee, Jue Wang, Qian Cong, Lisa N Kinch, R Dustin Schaeffer, et al. Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557):871–876, 2021.
- [5] William R Pearson. An introduction to sequence similarity (“homology”) searching. *Current protocols in bioinformatics*, 42(1):3–1, 2013.
- [6] Nelson Perdigão, Julian Heinrich, Christian Stolte, Kenneth S Sabir, Michael J Buckley, Bruce Tabor, Beth Signal, Brian S Gloss, Christopher J Hammang, Burkhard Rost, et al. Unexpected

features of the dark proteome. *Proceedings of the National Academy of Sciences*, 112(52):15898–15903, 2015.

- [7] 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.
- [8] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.
- [9] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [10] Erik Nijkamp, Jeffrey A Ruffolo, Eli N Weinstein, Nikhil Naik, and Ali Madani. Progen2: exploring the boundaries of protein language models. *Cell Systems*, 14(11):968–978, 2023.
- [11] Bo Chen, Xingyi Cheng, Pan Li, Yangli-ao Geng, Jing Gong, Shen Li, Zhilei Bei, Xu Tan, Boyan Wang, Xin Zeng, et al. xtrimpoglm: unified 100b-scale pre-trained transformer for deciphering the language of protein. *arXiv preprint arXiv:2401.06199*, 2024.
- [12] Noelia Ferruz, Steffen Schmidt, and Birte Höcker. Protgpt2 is a deep unsupervised language model for protein design. *Nature communications*, 13(1):4348, 2022.
- [13] Baris E Suzek, Hongzhan Huang, Peter McGarvey, Raja Mazumder, and Cathy H Wu. Uniref: comprehensive and non-redundant uniprot reference clusters. *Bioinformatics*, 23(10):1282–1288, 2007.
- [14] Helen M Berman, John Westbrook, Zukang Feng, Gary Gilliland, Talapady N Bhat, Helge Weissig, Ilya N Shindyalov, and Philip E Bourne. The protein data bank. *Nucleic acids research*, 28(1):235–242, 2000.
- [15] Martin Steinegger, Milot Mirdita, and Johannes Söding. Protein-level assembly increases protein sequence recovery from metagenomic samples manyfold. *Nature methods*, 16(7):603–606, 2019.
- [16] Martin Steinegger and Johannes Söding. Clustering huge protein sequence sets in linear time. *Nature communications*, 9(1):2542, 2018.
- [17] Roshan M Rao, Jason Liu, Robert Verkuil, Joshua Meier, John Canny, Pieter Abbeel, Tom Sercu, and Alexander Rives. Msa transformer. In *International Conference on Machine Learning*, pages 8844–8856. PMLR, 2021.
- [18] He Zhang, Fusong Ju, Jianwei Zhu, Liang He, Bin Shao, Nanning Zheng, and Tie-Yan Liu. Co-evolution transformer for protein contact prediction. *Advances in Neural Information Processing Systems*, 34:14252–14263, 2021.
- [19] Timothy Truong Jr and Tristan Bepler. Poet: A generative model of protein families as sequences-of-sequences. *Advances in Neural Information Processing Systems*, 36, 2024.
- [20] Le Zhang, Jiayang Chen, Tao Shen, Yu Li, and Siqi Sun. Enhancing the protein tertiary structure prediction by multiple sequence alignment generation. *arXiv preprint arXiv:2306.01824*, 2023.
- [21] Jun Zhang, Sirui Liu, Mengyun Chen, Haotian Chu, Min Wang, Zidong Wang, Jiali Yu, Ningxi Ni, Fan Yu, Dechin Chen, et al. Unsupervisedly prompting alphafold2 for accurate few-shot protein structure prediction. *Journal of Chemical Theory and Computation*, 19(22):8460–8471, 2023.
- [22] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019.

[23] Jürgen Haas, Alessandro Barbato, Dario Behringer, Gabriel Studer, Steven Roth, Martino Bertoni, Khaled Mostaguir, Rafal Gumienny, and Torsten Schwede. Continuous automated model evaluation (cameo) complementing the critical assessment of structure prediction in casp12. *Proteins: Structure, Function, and Bioinformatics*, 86:387–398, 2018.

[24] Chengxin Zhang, Wei Zheng, SM Mortuza, Yang Li, and Yang Zhang. Deepmsa: constructing deep multiple sequence alignment to improve contact prediction and fold-recognition for distant-homology proteins. *Bioinformatics*, 36(7):2105–2112, 2020.

[25] Wei Zheng, Qiqige Wuyun, Yang Li, Chengxin Zhang, P Lydia Freddolino, and Yang Zhang. Improving deep learning protein monomer and complex structure prediction using deepmsa2 with huge metagenomics data. *Nature Methods*, pages 1–11, 2024.

[26] L Steven Johnson, Sean R Eddy, and Elon Portugaly. Hidden markov model speed heuristic and iterative hmm search procedure. *BMC bioinformatics*, 11:1–8, 2010.

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

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

[29] Ruidong Wu, Fan Ding, Rui Wang, Rui Shen, Xiwen Zhang, Shitong Luo, Chenpeng Su, Zuofan Wu, Qi Xie, Bonnie Berger, et al. High-resolution de novo structure prediction from primary sequence. *BioRxiv*, pages 2022–07, 2022.

[30] Ratul Chowdhury, Nazim Bouatta, Surojit Biswas, Christina Floristean, Anant Kharkar, Koushik Roy, Charlotte Rochereau, Gustaf Ahdritz, Joanna Zhang, George M Church, et al. Single-sequence protein structure prediction using a language model and deep learning. *Nature Biotechnology*, 40(11):1617–1623, 2022.

[31] 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.

[32] Ali Madani, Ben Krause, Eric R Greene, Subu Subramanian, Benjamin P Mohr, James M Holton, Jose Luis Olmos Jr, Caiming Xiong, Zachary Z Sun, Richard Socher, et al. Large language models generate functional protein sequences across diverse families. *Nature Biotechnology*, pages 1–8, 2023.

[33] Ahmed Elnaggar, Michael Heinzinger, Christian Dallago, Ghalia Rehawi, Yu Wang, Llion Jones, Tom Gibbs, Tamas Feher, Christoph Angerer, Martin Steinegger, et al. Prottrans: Toward understanding the language of life through self-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(10):7112–7127, 2021.

[34] Sarah Alamdari, Nitya Thakkar, Rianne van den Berg, Alex X Lu, Nicolo Fusi, Ava P Amini, and Kevin K Yang. Protein generation with evolutionary diffusion: sequence is all you need. *BioRxiv*, pages 2023–09, 2023.

[35] Jonathan Frazer, Pascal Notin, Mafalda Dias, Aidan Gomez, Joseph K Min, Kelly Brock, Yarin Gal, and Debora S Marks. Disease variant prediction with deep generative models of evolutionary data. *Nature*, 599(7883):91–95, 2021.

[36] Bo Chen, Ziwei Xie, Jiezhong Qiu, Zhaofeng Ye, Jinbo Xu, and Jie Tang. Improved the heterodimer protein complex prediction with protein language models. *Briefings in Bioinformatics*, 24(4):bbad221, 2023.

[37] Pascal Sturmels, Jesse Vig, Ali Madani, and Nazneen Fatema Rajani. Profile prediction: An alignment-based pre-training task for protein sequence models. *arXiv preprint arXiv:2012.00195*, 2020.

[38] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290*, 2023.

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

[40] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

[41] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.

[42] Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35:16344–16359, 2022.

[43] Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. *arXiv preprint arXiv:2307.08691*, 2023.

[44] Gustaf Ahdritz, Nazim Bouatta, Sachin Kadyan, Lukas Jarosch, Dan Berenberg, Ian Fisk, Andrew Watkins, Stephen Ra, Richard Bonneau, and Mohammed AlQuraishi. Openproteinset: Training data for structural biology at scale. *Advances in Neural Information Processing Systems*, 36, 2024.

[45] Milot Mirdita, Lars Von Den Driesch, Clovis Galiez, Maria J Martin, Johannes Söding, and Martin Steinegger. Uniclus databases of clustered and deeply annotated protein sequences and alignments. *Nucleic acids research*, 45(D1):D170–D176, 2017.

[46] Martin Steinegger, Markus Meier, Milot Mirdita, Harald Vöhringer, Stephan J Haunsberger, and Johannes Söding. Hh-suite3 for fast remote homology detection and deep protein annotation. *BMC bioinformatics*, 20(1):1–15, 2019.

[47] 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.

[48] Francesca-Zhoufan Li, Ava P Amini, Yisong Yue, Kevin K Yang, and Alex X Lu. Feature reuse and scaling: Understanding transfer learning with protein language models. *bioRxiv*, pages 2024–02, 2024.

[49] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

[50] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018.

Appendix

Table of Contents

A	Training Settings and Hyper-parameter Studies.	14
A.1	Pre-Training	14
A.2	RFT	15
A.3	RLAF	16
A.4	Inference Efficiency	17
B	Experimental Settings	18
B.1	Evaluation Details	18
B.2	Experiments on MSA-abundant Scenario	18
B.3	Setup of Transferability of MSAGPT to Other Tasks	18
B.4	Setup of Ablation Study	18
C	Selection Strategy Details and pLDDT Evaluation	19
D	Protein Structure Prediction Analysis compared with natural MSA	19
E	Protein Structure Prediction Improvement after DPO	20

A Training Settings and Hyper-parameter Studies.

The overall training pipeline is illustrated in Figure 6.




Figure 6: The overall training pipeline and the illustration of preference dataset construction process for SFT and DPO learning stages.

A.1 Pre-Training

To obtain high-quality MSA, we first screen out clusters with sequence lengths from 25 to 2000, and only retain sequences with the minimum identity of 30% and the largest proportion of gap tokens no more than 10%. The clusters with more than 10 sequences are left. As the MSA obtained by MSA search tools, inevitably contain noisy co-evolutionary patterns, such as large portions of deletions,

Table 5: The performance (pTM, GDT, LDDT) of structure prediction on three natural MSA-scarce benchmarks.

Model	CAMEO (avg. Depth = 8.5)						CASP (avg. Depth = 4.6)						PDB (avg. Depth = 2.6)					
	Zero-Shot			Few-Shot			Zero-Shot			Few-Shot			Zero-Shot			Few-Shot		
	pTM	GDT	LDDT	pTM	GDT	LDDT	pTM	GDT	LDDT	pTM	GDT	LDDT	pTM	GDT	LDDT	pTM	GDT	LDDT
AF2 MSA	41.3	60.8	66.7	58.2	75.5	80.3	31.2	29.5	40.5	40.8	41.6	51.9	40.3	43.7	56.6	46.2	50.4	62.8
EvoDiff	54.2	62.4	68.1	48.1	63.0	67.7	40.2	27.3	38.4	44.1	35.6	46.7	47.4	43.3	55.9	45.5	45.0	57.1
MSA-Aug.	47.2	62.3	68.3	58.4	75.1	80.6	39.1	33.2	42.7	44.3	43.2	51.9	44.8	47.6	59.8	48.8	52.8	65.2
EvoGen	45.0	63.5	68.8	61.8	78.8	82.8	36.0	34.8	44.9	42.8	45.1	54.6	44.3	47.0	59.4	49.1	53.0	65.4
MSAGPT	51.4	64.3	69.4	63.6	78.8	83.2	40.0	36.3	45.5	47.7	47.8	56.0	48.6	50.6	53.3	53.7	57.0	69.2
+ RFT	48.1	63.5	68.1	62.4	79.0	83.1	41.2	36.1	45.5	49.8	49.8	56.0	48.6	50.5	62.4	53.6	57.2	69.0
+ DPO	49.4	64.4	68.7	63.0	79.4	83.5	40.0	39.5	48.1	49.8	53.1	60.0	47.6	50.9	63.2	52.7	57.8	69.5

insertions, and gaps. Many previous works aim to filter high-quality MSAs by clustering or ranking them based on predefined rules. One primary rule is to find MSA sequences most similar to the main sequence with fewer gaps, as these are more likely to represent informative co-evolutionary patterns. Following this idea, we sample MSA sequences using similarity-based weighted sampling, where sequences more similar to the query protein are more likely to be selected and ranked higher. Moreover, we also randomly shuffle the sequences in the selected MSA by a certain proportion to avoid injecting the order bias.

Regarding the backbone of MSAGPT, we employ the standard transformer decoder framework [47, 49] and train the model with 2.8 billion parameters owning 36 layers, 2560 embedding size, and 40 attention heads. We employ batches of 48 MSAs with each MSA containing 12,288 residues. We follow BF16 mixed-precision pre-training strategy. We use AdamW [50] as our optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.95$, $\text{eps} = 10^{-8}$ and a learning rate of 1.2×10^{-4} . We use a cosine learning rate schedule, with a warmup of the first 2.5% steps, and decay the final learning rate down to 10% of the peak learning rate. We use a weight decay of 0.1 and gradient clipping of 1.0 without dropout. For the tokenization of the protein data, we use the residue-level tokenizer which is adopted in several PLMs [28, 11, 10]. To save the GPU memory and accelerate the pre-training process, we substitute the standard self-attention module with the Flash Attention-v1 [42] in each layer. All models are trained on 24 A800 GPUs for 254k updates, consuming about 150 billion tokens. This process consumes approximately 150 billion tokens, requiring around 2.7×10^{18} floating point operations (FLOPs).

Pre-trained Dataset Figure 7 illustrates the length and depth distribution of the pre-training dataset. Moreover, we implement a thorough filtering process to eliminate any potential data leakage. Specifically, we remove all MSAs of sequences in the test sets (CAMEO, CASP, and PDB) from the pre-training dataset. Furthermore, we ensure that any sequence in the pre-training set with a similarity greater than 0.9 to a sequence in the test set is excluded. To validate this filtering process, we used the HHblits [46] tool to retrieve sequences from the test set and calculate their maximal similarity distribution with sequences in the pre-training dataset. The results, illustrated in Figure 9(b) show that the maximum similarity is 0.89, confirming that there is no data leakage in the pre-training dataset.

A.2 RFT

We fine-tune the base model using the pre-training cross-entropy loss on \mathcal{D}_{RFT} with training only one epoch. Specifically, we adopt the same experimental settings as that used in the pre-training stage, except for the learning rate of 1.0×10^{-5} by default. Following the pre-training phase, the model undergoes a rejective fine-tuning process, which is more energy-efficient. This stage is executed on 8 x A800 GPUs for a single epoch for about two days.

RFT Dataset Filtering Threshold When curating the RFT dataset, we first sample multiple MSA subsets for each protein structure, and select high-quality MSA subsets based on the following criteria: (1) the absolute structure prediction accuracy using the MSA subset, as measured by TM-score, should be larger than θ_1 , and (2) the relative improvement of the prediction accuracy after using the MSA subset, as compared to single sequence prediction, should be larger than θ_2 . We set $\theta_1 = 0.9$, and experiment with different θ_2 values, as shown in table 7. The RFT model trained with dataset




Figure 7: The length and depth distribution of the pre-training dataset.




Figure 8: The correlation between total token length (the protein sequence length multiplied by the number of generated MSAs) and the inference time (minutes). In most cases (total token length $< 20K$), the generation time of MSAGPT is lower than the AF2 search pipeline requiring more than 30 minutes. The result shows MSAGPT can generate substantial sequence lengths within practical time, thus affirming its scalability and efficiency.

filtered by $\theta_2 = 0.2$ yields the best result, indicating that the relative information gain provided by MSA is a valuable indicator for curating high quality datasets for RFT. Moreover, $\theta_2 = 0.5$ results in a 20% decrease in dataset size, leading to inferior RFT model performance, highlighting the necessity of an intricate balance between data quality and data volume.

A.3 RLAf

We fine-tune the RFT model using the DPO algorithm on \mathcal{D}_{DPO} with training only one epoch. Specifically, we adopt the batch size of 1 with each MSA subset containing a maximum of 16,384 residues. We also use AdamW [50] with the learning rate of 1.0×10^{-6} by default. We linearly warmup the learning rate from 0 to 1.0×10^{-6} over the first 0.1% steps. This stage is also executed on 8 x A800 GPUs for a single epoch for about one day

Table 6: Performance comparison between different data source and filtering threshold values.

Data Source	nature(0.2)	nature(0.3)	generated(0.3)	nature(0.3)+generated(0.4)
MSAGPT+ RLAF	62.6	64.5	63.5	62.7

Table 7: Performance comparison between different relative improvement threshold θ_2 values.

Threshold θ_2	0	0.2	0.5
MSAGPT+ RFT	61.2	62.5	61.3




Figure 9: (a) The distribution of MSA depth of benchmarked datasets and (b) the similarity distribution of sequences in the test set, as retrieved from the pre-training set using HHblits.

RLAF Dataset. We conducted experiments with different data sources and filtering thresholds θ_3 —defined as the minimum relative improvement of the good case over the bad case in DPO data pairs—for the RLAF dataset, as detailed in Table 6. Utilizing only natural MSA subsets sampled from PDB, we found that higher θ_3 values lead to improved model performance, suggesting a correlation between the disparity within data pairs and DPO effectiveness. Interestingly, the quality of MSA subsets generated by the RFT model surpasses that of natural MSA subsets at a θ_3 of 0.2. However, the performance declines when natural MSAs are mixed with generated MSAs, compared to using a single data source during training. This indicates that maintaining distribution homogeneity is crucial for effective DPO alignment.

A.4 Inference Efficiency

Generally, it’s vital to consider not just the immediate resource consumption during pre-training and post-alignment, but also the long-term utilization of these models. Once pre-trained, MSAGPT demonstrates significant efficiency, capable of generating protein sequences with up to 100,000 amino acids in under 8 hours. This efficiency underscores the model’s value, especially when amortized over its application lifespan and subsequent fine-tunings for specific tasks.

Regarding the scalability of the MSAGPT. We present the inference time with different total lengths (measured by protein sequence length multiply the number of generated sequences.), as shown in Figure 8.

The result showcases MSAGPT’s ability to generate substantial sequence lengths within practical time frames, thus affirming its scalability and efficiency.

Table 8: The paired Student’s t-test between MSAGPT and other baselines on three benchmarks based on the TM-Score, where the p-value less than 0.05 indicates the result is said to be statistically significant.

Model	CAMEO (avg. Depth = 8.5)		CASP (avg. Depth = 4.6)		PDB (avg. Depth = 2.6)	
	Zero-Shot	Few-Shot	Zero-Shot	Few-Shot	Zero-Shot	Few-Shot
AF2 MSA	0.014	0.023	0.008	0.007	5e-7	8e-9
EvoDiff	0.016	1e-5	5e-4	7e-5	0.012	1e-8
MSA-Aug.	0.023	0.014	0.044	0.015	6e-6	1e-7
EvoGen	0.038	0.027	0.067	0.016	1e-8	1e-9
MSAGPT	-	-	-	-	-	-

Table 9: The results of 5-fold cross-validation performance between with or without virtual MSA generated by MSAGPT on four protein-related tasks.

Model	1	2	3	4	5	AVG
	Top (L/5)	ACC	ACC	ACC	ACC	-
w/o Virtual MSA (CtP)	11.4	14.3	10.7	9.9	11.8	11.6
w/ Virtual MSA (CtP)	14.1	13.7	13.4	11.8	12.3	13.1
w/o Virtual MSA (SsP)	67.7	65.8	64.0	68.9	66.2	66.5
w/ Virtual MSA (SsP)	70.5	67.8	67.5	70.5	69.0	69.0
w/o Virtual MSA (LocP)	56.0	64.5	48.0	59.0	57.0	58.3
w/ Virtual MSA (LocP)	47.0	58.5	53.5	64.0	59.0	56.4
w/o Virtual MSA (MIB)	58.0	53.5	49.5	59.0	67.5	57.5
w/ Virtual MSA (MIB)	61.5	57.0	63.0	53.0	67.0	60.3

B Experimental Settings

B.1 Evaluation Details

We employ three golden metrics: TM-Score, a widely-used metric for assessing the structural similarity between predicted structures and ground truth, LDDT, the local distance difference test score measures how well local interactions in a reference structure are conserved the protein model being assessed, GDT-TS, the global distance test to represent “total score”, is a measure of similarity between two protein structures with known amino acid correspondences (e.g. identical amino acid sequences) but different tertiary structures, and two predicted metrics: pLDDT, the predicted local distance difference test measuring the local confidence of per-residue and pTM, an integrated measure of how well AlphaFold2 has predicted the overall structure. All metrics are normalized from 0 to 100 for comparison, with higher scores indicating higher confidence and usually a more accurate prediction. where 1 indicates a perfect match between two structures. Each run across 3 independent runs. For each run, we adopt the different temperatures ($T \in \{0.8, 1.0\}$) along with different nucleus sampling factors ($P \in \{0.8, 1.0\}$), experimenting with varying the number of generated MSAs in 8, 16, 32, and 64. The final performance is determined by first identifying the predicted structure with the highest accuracy across these different depths, and then averaging the results across the test set.

B.2 Experiments on MSA-abundant Scenario

We compare the results of query sequences with abundant natural MSAs to those with abundant natural MSAs augmented by MSAGPT’s generated MSAs on CAMEO set. For this comparison, we sample 128, 256, and 512 sequences from both the natural MSAs and the generated MSAs, as shown in Table 10. These results indicate that the inclusion of generated MSAs has no significant effect on the performance in MSA-abundant conditions, which is consistent with previous findings that when more than 64 MSAs as input, AF2 predicts a “converged” structure.

B.3 Setup of Transferability of MSAGPT to Other Tasks

We utilized the MSA Transformer [17] as the backbone model with the task-specific heads. We finetune MSA transformer with the head with $lr = 3e - 5$ and batchsize = 16 on all experiments. All the task benchmarks are obtained following the pipeline in [11]. For each task, we sample 1000 protein sequences with the corresponding labels. Then we use MSAGPT-DPO to generate 32 virtual MSAs with the generation strategy $T=0.8$ and $P=0.8$. Both setups are trained briefly (for one epoch) for 5-fold cross-validation as shown in Table 9, and we report the average performance.

B.4 Setup of Ablation Study

Experiments are conducted based on models with 150 million parameter size encompassing 30 layers, 20 attention heads, 640 embedding dimensions. These models are trained across approximately 30 billion tokens, amounting to around 40k training steps, maintaining consistency in hyper-parameter

Table 10: Performance comparison in MSA-abundant scenarios across all 194 cases in the CAMEO datasets.

# Natural MSA	# Virtual MSA	TM	GDT-TS
0	0	39.1	28.5
128	0	85.1	82.0
128	128	85.3	82.2
256	0	86.2	83.8
256	256	86.0	83.5
512	0	86.5	84.2
512	512	86.4	84.0

settings with MSAGPT, except for variations in the positional encoding mechanism. The efficacy of each variant is assessed through zero-shot MSA generation on the CASP test set.

C Selection Strategy Details and pLDDT Evaluation

To evaluate the effectiveness of different selection strategies, we extracted 4, 8, 12, 16, 24, and 32 sequences from 48 zero-shot generated MSA for each method and computed the median TM-scores (Figure 4(b)) and pLDDT scores (Figure 10) across 33 test cases. The strategies are detailed below.

Static Similarity / Static Diversity Strategy: We select the top-k generated MSA with the highest / lowest sequence identity to the query sequence. Sequence identity is determined by the percentage of identical residues between the two sequences.

Dynamic Similarity / Static Diversity Strategy: Starting with the MSA most / least similar to the query sequence, we sequentially incorporate MSA into the selected set based on the highest or lowest average sequence identity with all sequences already included, until reaching a total of k MSA.

Trimming Strategy: Suggested by EvoGen, this method filters out MSA with less than 50% coverage or sequence identity to the query sequence above 90% or below 20%. Subsequently, it iteratively selects the MSA with the closest sequence identity to the query and an average sequence identity below 90% with all the chosen MSA.

pTM / pLDDT / TM Score Strategy: For each generated MSA, we remove the gaps and predict its structure using AF2. The structures are then ranked based on the pTM score (as reported by AF2), the pLDDT score (as reported by AF2), or the TM score compared to the query sequence’s ground truth structure (calculated by US Align), and the MSA corresponding to the top-k structures for each metric are selected accordingly.

D Protein Structure Prediction Analysis compared with natural MSA

We present a detailed visual comparison of protein structures predicted by AlphaFold2 (AF2) utilizing MSA augmented by MSAGPT, against those predicted with natural MSA. This comparison, as depicted in Figure 12, highlights the remarkable capability of MSAGPT in enhancing the accuracy of structure predictions to levels that closely rival, and in some cases surpass, those based on naturally derived MSA.

We delve into a visualized analysis of protein structures predicted using AlphaFold2 (AF2) with MSA augmented by our proposed model (MSAGPT), alongside those augmented by EvoGen and




Figure 10: **The pLDDT curves across different selection methods.** Dashed red line represents using all generated sequences of a given depth. Solid lines represent selecting a subset of a given depth from 48 generated sequences with a specific strategy. The curves are smoothed using the Exponential Moving Average with alpha=0.3.




Figure 11: Visualization of improved structure prediction compared with baseline models. **Yellow:** Ground truth; **Pink:** Predictions based on MSA generated by MSAGPT; **Blue:** Predictions from MSA generated by EvoGen; **Green:** Predictions utilizing MSA generated by MSA-Augmenter.

MSA-Augmenter. This comparison, visualized in Figure 11, encompasses a spectrum of proteins, ranging from short sequences with relatively simple structures, like 7mnv_B, to long sequences with complex configurations, such as 7tdv_B. It includes proteins characterized by multiple beta sheets, exemplified by 7ywg_B, as well as those rich in alpha helices, such as 7tdv_B. Across these diverse cases, MSAGPT significantly surpasses both EvoGen and MSA-Augmenter, improving the TM score to a maximum of 0.9.

By detailed examination, we observe that while the MSA augmented by the baseline models assist AF2 in accurately predicting local structures and folds, they fall short in aligning the global composition and orientation with the ground truth structure, which is effectively addressed by MSA generated by MSAGPT. The local structures, which are generally more discernible from the spatial arrangements of adjacent amino acids, contrast with the global structures whose prediction relies heavily on comprehensively understanding the co-evolutionary information within MSA. These co-evolutionary patterns, indicating proximity in three-dimensional space through simultaneous mutations at multiple positions, are crucial for accurate global structure prediction. These findings underscore MSAGPT's impressive capability to comprehensively capture and utilize co-evolutionary information, thereby significantly enhancing the accuracy of protein structure predictions. More visualization cases about the predictions based on MSA generated by MSAGPT and the predictions based on the natural MSA are illustrated in Appendix D.

E Protein Structure Prediction Improvement after DPO

Figure 13 represents the comparison before and after the DPO training, depicting notable enhancements in structure prediction accuracy. Figure 14 and 15 provide an in-depth analysis of the generated MSA for each case. Specifically, residues 43, 53, 71-79, 105-111, 122, 132 and 157-166 in the




Figure 12: Visualization of improved structure prediction compared with nature MSA. Yellow: Ground truth; Pink: Predictions based on MSA generated by MSAGPT; Blue: Predictions from MSA generated by natural MSA.




Figure 13: **Visualization of improved structure prediction after DPO.** **Yellow:** Ground truth; **Blue:** Predictions based on MSA generated by MSAGPT; **Pink:** Predictions based on MSA generated by MSAGPT-DPO.;

MSA of 7wme_A, along with residues 22-27, 53, and 73 in the MSA of 7sxb_A, display distinct characteristics pre- and post-DPO training.




Figure 14: **Residue Distribution of Generated MSA for 7wme_A.** The red box indicates natural MSA used as prompts during generation. The blue box indicates generated MSA. Residues are colored using the clustal scheme by Jalview.




Figure 15: Residue Distribution of Generated MSA for 7sxb_A. The red box indicates natural MSA used as prompts during generation. The blue box indicates generated MSA. Residues are colored using the clustal scheme by Jalview.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: **[Yes]**

Justification: We clearly made statements to claim the contributions in the abstract and introduction sections.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: **[Yes]**

Justification: We discussed several limitations of our work in Section [7](#).

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: **[NA]**

Justification: This work is primarily focused on empirical experiments and extensive validation rather than providing theoretical results. Therefore, the paper does not include a set of assumptions or proofs for theoretical results since it is not the main emphasis of the research.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide the implementation details in the Appendix and the released model weight at <https://github.com/THUDM/MSAGPT> to ensure the reproduction of all the experimental results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide the implementation details in the Appendix to ensure the reproduction of all the experimental results. Moreover, we also released model weight at <https://github.com/THUDM/MSAGPT> to the community.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide the implementation details in the main content or the Appendix to ensure the reproduction of all the experimental results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We have presented the results of the paired Student’s t-test between our model and other baselines on three benchmarks (Table 1) in the Appendix Table 8.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: We provide the implementation details in the Appendix and the released model weight at <https://github.com/THUDM/MSAGPT> to ensure the reproduction of all the experimental results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: We have reviewed and obeyed the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[Yes\]](#)

Justification: We have already discussed the broader impact in Section 8.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [\[Yes\]](#)

Justification: We have already discussed the broader impact with potential risk in Section 8.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [\[Yes\]](#)

Justification: We have mentioned all the sources of used data, code, and models, and given the credited to the corresponding authors or organizations.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [\[Yes\]](#)

Justification: We release the inference script and model weight at <https://github.com/THUDM/MSAGPT>.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [\[NA\]](#)

Justification: We don't include human subjects or crowdsourcing in this work.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [\[NA\]](#)

Justification: We don't include human subjects or crowdsourcing in this work.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.