# Trainable subnetworks reveal insights into structure knowledge organization in protein language models

Ria Vinod[1], Ava P. Amini[2], Lorin Crawford[2,*], and Kevin K. Yang[2,*]

[1]Center for Computational Molecular Biology, Brown University, Providence, RI, USA
[2]Microsoft Research, Cambridge, MA, USA
[*]Corresponding Email(s): lcrawford@microsoft.com; yang.kevin@microsoft.com

**Abstract**

Protein language models (PLMs) pretrained via a masked language modeling objective have proven effective across a range of structure-related tasks, including high-resolution structure prediction. However, it remains unclear to what extent these models factorize protein structural categories among their learned parameters. In this work, we introduce trainable subnetworks, which mask out the PLM weights responsible for language modeling performance on a structural category of proteins. We systematically trained 39 PLM subnetworks targeting both sequence- and residue-level features at varying degrees of resolution using annotations defined by the CATH taxonomy and secondary structure elements. Using these PLM subnetworks, we assessed how structural factorization in PLMs influences downstream structure prediction. Our results show that PLMs are highly sensitive to sequence-level features and can predominantly disentangle extremely coarse or fine-grained information. Furthermore, we observe that structure prediction is highly responsive to factorized PLM representations and that small changes in language modeling performance can significantly impair PLM-based structure prediction capabilities. Our work presents a framework for studying feature entanglement within pretrained PLMs and can be leveraged to improve the alignment of learned PLM representations with known biological concepts.

## Introduction

Understanding protein structure is essential for deciphering biological function, as a protein's structure governs its molecular stability, interactions, and activity. Recent protein language models (PLMs) trained purely on sequence data have been shown to learn representations that implicitly encode structural information [1–4]. These models have proven effective in a wide range of protein engineering tasks including structure prediction [5–7], function annotation [8, 9], mutation effect estimation [10], and even the design of novel proteins [11–13].

Many PLMs are pretrained with a self-supervised masked language modeling (MLM) objective, where models are tasked with predicting the amino acid identities of randomly masked tokens in a sequence [1, 2, 14, 15]. Since protein structure is fundamentally determined by amino acid sequence, PLMs can implicitly encode structural information in their weights. ESM-2, a family of protein language models trained on evolutionary-scale data, showed that performance on the language modeling task can predict the quality of a PLM-predicted structure [1, 7]. PLM representations are now widely used as inputs to models that predict structural coordinates [16], and scaling analyses have shown that improvements in MLM performance improve single-sequence structure prediction accuracy [7, 17].

The simplicity and efficacy of pretraining language models on protein sequences have led to growing interest in understanding their internal mechanisms, establishing interpretability as an important research direction. With respect to the structure features that PLMs learn, prior

work has shown that protein contact information is stored in attention maps [18], that PLMs learn co-evolving motifs which emulate an understanding of fundamental biophysics [19], and that sparse latent features from PLMs capture known functional biophysical properties and motifs [20, 21].

To date, however, it remains unclear (i) if and how structure information is factorized and stored in the learned PLM weights, and (ii) whether this factorization affects performance on downstream structure prediction tasks. One approach to uncover whether such factorization of concepts in pretrained language models exist is via subnetwork discovery [22]. Subnetworks are sparse computational subgraphs of a pretrained model's weights that are responsible for performance on a specific task or class of inputs. In the natural language domain, subnetwork discovery has been widely used to uncover linguistic properties—such as semantics, syntax, and relational entities—learned during pretraining by defining concepts and localizing weights, neurons, or layers that encode them [23–27].

In this work, we focus on leveraging subnetworks to interrogate the factorization of structural categories of proteins in the pretrained PLM, ESM-2 (Fig. 1). Our goal was to find PLM subnetworks—sparse subgraphs of the original model weights—that when isolated, suppress the model's ability to make correct predictions on one class of inputs while preserving MLM performance on all other classes of inputs. As part of our main analysis, we systematically trained 39 such PLM subnetworks to suppress either residue- or sequence-level structural information across different scales of the CATH hierarchy. Our results reveal that structural categories are indeed encoded in a factorized manner within PLM weights and that, although non-suppressed inputs achieve ESM-2-level perplexity, structure prediction is still statistically significantly perturbed. Together, our subnetworks approach and the suite of analyses we perform provide insight into how PLMs organize structural features among learned parameters.

# Methods

## Suppression and maintenance inputs are defined by structural annotations

Protein secondary structure refers to the local spatial arrangement of amino acid residue backbone atoms. Each residue in a protein sequence adopts a secondary structure, which can be classified as an alpha helix, beta sheet, or loop. The resulting structural arrangements enable further categorization of proteins at the sequence level, grouping them by shared architectural and evolutionary features (Fig. 1A). In this work, we define structural categories at the residue level (alpha helix, beta sheet, loop) using DSSP [28] and at the sequence-level (Class, Architecture, Topology, Homologous Superfamily) using the CATH taxonomy (Fig. 1D).

## Notation

We denote a pretrained PLM as $f(\mathbf{x}, \boldsymbol{\theta})$, where $\mathbf{x}$ represents the input sequence and $\boldsymbol{\theta}$ are the model weights. The objective of our training procedure is to predict a binary mask, $\mathbf{m} \in \{0, 1\}^K$, where $K$ denotes the number of parameters in the PLM that we seek to learn the mask over. A subnetwork can be obtained by taking the Hadamard product between the binary mask and pretrained PLM weights, $f(\mathbf{x}, \mathbf{m} \odot \boldsymbol{\theta})$. Subnetworks are trained independently to suppress a structural category of sequence-level or residue-level inputs (i.e., all suppressed inputs belong to the same category defined by structural annotation). We define a suppression input sequence as $\mathbf{x}^*$ and all other input sequences (i.e., maintenance inputs) as $\mathbf{x}'$ (Fig. 1A; suppressed inputs are red, maintained inputs are black). For residue-level suppression, let $\mathcal{A}$ be an annotation from DSSP that contains a set of $J$ positions corresponding to residues in {alpha helix, beta sheet, loop}. We define the $J$ suppression inputs for this DSSP annotation as $\mathbf{x}_{(J^*)}$. The maintenance inputs are all remaining complementary positions of residues with annotation $\mathcal{A}$, which are then defined as $\mathbf{x}_{(J')}$.

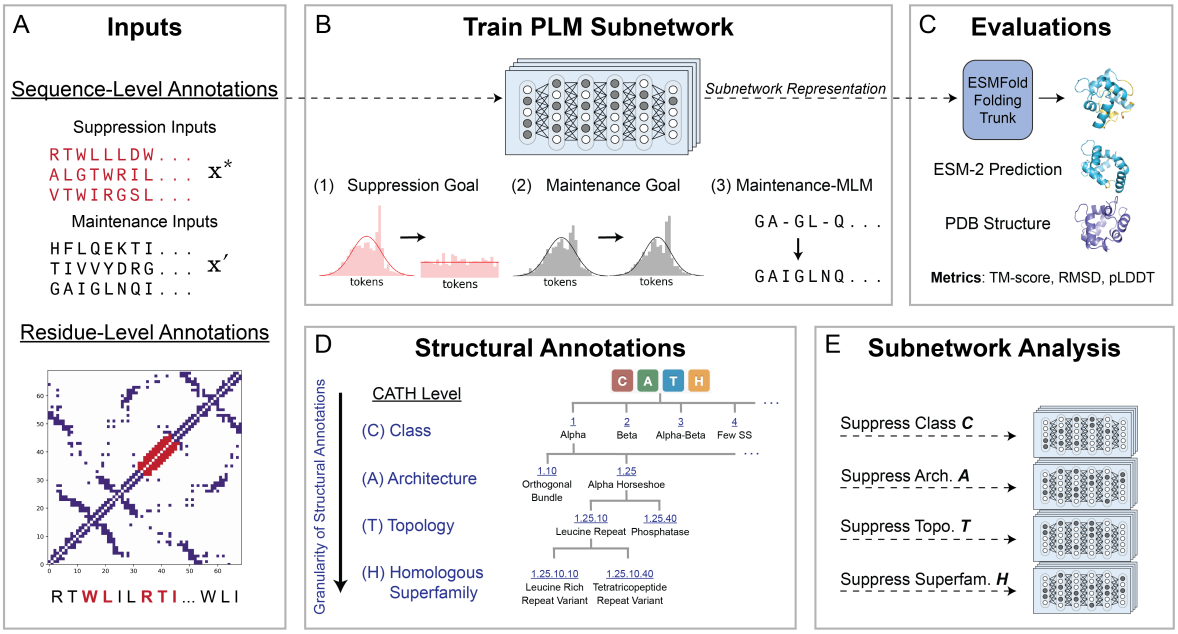


**Figure 1. Summary of subnetwork training procedure and main set of evaluations and analysis. (A)** Suppression inputs can be defined at either the protein sequence-level (based on CATH annotations) or the residue-level (based on 3-way secondary structure annotations). Suppressed inputs are red; maintained inputs are black. **(B)** PLM subnetworks are trained via a weighted multi-objective differentiable mask learning scheme, which defines suppression and maintenance goals for a binary mask learning procedure. **(C)** The subnetwork sequence representation is passed as input to the ESMFold folding trunk to predict a 3D structure. The resulting 3D structure can be compared to the ESM-2 prediction for the same sequence, both of which are independently aligned to ground-truth structures from the Protein Data Bank (PDB). **(D)** The CATH database offers a hierarchical classification of protein structural annotations across four levels, grouping proteins by their fold patterns and evolutionary relationships. Each successive level provides increasingly fine-grained structural annotations. **(E)** Independent subnetworks are trained to suppress different structural categories.

## Subnetwork training objective

To obtain a subnetwork $f(\mathbf{x}, \mathbf{m} \odot \boldsymbol{\theta})$, we learn a binary mask $\mathbf{m}$ using a weighted loss that consists of the following components (Fig. 1B):

1. **Suppression goal.** If the mask is working appropriately, the subnetwork should struggle to reconstruct its suppressed inputs accurately. In other words, a well-calibrated predictive distribution over the suppressed input tokens should be approximately uniformly distributed with respect to the vocabulary $V$. As a result, we define the suppression loss as minimizing the Kullback–Leibler (KL) divergence between (i) the predictive distribution of the subnetwork over the tokens corresponding to the suppression inputs and (ii) a uniform reference distribution over tokens in the vocabulary (denoted as $\mathcal{U}_V$). For sequence-level suppression, this corresponds to

$$\mathcal{L}_{\text{supp}} = \text{KL}\left(f(\mathbf{x}^*, \mathbf{m} \odot \boldsymbol{\theta}), \mathcal{U}_V\right). \tag{1}$$

   Similarly, for residue-level suppression,

$$\mathcal{L}_{\text{supp}} = \text{KL}\left(f(\mathbf{x}_{(J^*)}, \mathbf{m} \odot \boldsymbol{\theta}), \mathcal{U}_V\right) \tag{2}$$

   where, again, the index $J^*$ is used to represent all the positions corresponding to residues with DSSP annotation $\mathcal{A}$.

2. **Maintenance-KL goal.** Even in the presence of a mask, the subnetwork should preserve the predictive behavior of the full PLM on maintenance inputs. Therefore, as a maintenance goal, we also aim to minimize the KL divergence between (i) the predictive distribution of the subnetwork over the tokens corresponding to the maintenance inputs and (ii) the predictive distribution for the pretrained PLM over the same elements. For sequence-level suppression, this is expressed as the following

$$\mathcal{L}_{\text{maint}} = \text{KL}\left(f(\mathbf{x}', \mathbf{m} \odot \boldsymbol{\theta}), f(\mathbf{x}', \boldsymbol{\theta})\right). \tag{3}$$

Similarly, for the residue-level suppression, we write

$$\mathcal{L}_{\text{maint}} = \text{KL}\left(f(\mathbf{x}_{(J')}, \mathbf{m} \odot \boldsymbol{\theta}), f(\mathbf{x}_{(J')}, \boldsymbol{\theta})\right) \tag{4}$$

where, again, the index $J'$ denotes the set of positions not included in a given residue annotation $\mathcal{A}$ from DSSP.

3. **Maintenance masked language modeling (MLM) goal.** In practice, the maintenance-KL goal is insufficient because it only enforces similarity between the subnetwork and the original PLM output distributions, without ensuring that the subnetwork retains the ability to assign the correct probabilities to predicted tokens. Prior work demonstrated via a set of ablations that all 3 loss components are necessary to achieve the desired subnetwork behavior on suppression and maintenance inputs [22], and found that omitting either maintenance-KL or maintenance-MLM increased perplexity on the maintenance inputs. We therefore introduce an additional maintenance-MLM loss, which ensures that the subnetwork can still allocate the appropriate probability mass to the right corresponding tokens, preserving its overall language modeling behavior. To include an MLM objective on the maintenance inputs, we randomly select 15% of sequence positions over which the MLM loss is computed. Of these $M$ positions, 80% are replaced with a mask token, 10% are randomly replaced, and 10% are unchanged. Namely,

$$\mathcal{L}_{\text{MLM}} = \sum_{i \in M} \log p_{\mathbf{m} \odot \boldsymbol{\theta}}\left(x'_i \mid \mathbf{x}'_{(M')}\right) \tag{5}$$

where $M'$ are unmasked positions (i.e., true amino acid identities). To perform residue-level suppression, we again randomly choose a set of $M$ masked positions with the same mask-and-mutate scheme; however, this time, we compute the MLM loss with respect to the masked indices that overlap with the positions corresponding to a particular residue-level annotation $\mathcal{A}$, which we denote as $M \cap J$. This can be expressed as

$$\mathcal{L}_{\text{MLM}} = \sum_{i \in M \cap J'} \log p_{\mathbf{m} \odot \boldsymbol{\theta}}\left(x_i \mid \mathbf{x}_{(M' \cap J')}\right) \tag{6}$$

where $M \cap J'$ denotes intersection of indices that are in $M$ and not in $J$, and $M' \cap J'$ represents the indices that are not in either set.

Overall, the final weighted training objective is then represented as the sum

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{supp}} + \lambda_2 \mathcal{L}_{\text{maint}} + \lambda_3 \mathcal{L}_{\text{MLM}} \tag{7}$$

where $\boldsymbol{\lambda} = \{\lambda_1, \lambda_2, \lambda_3\}$ are hyperparameters. A description of how we select training hyperparameters can be found in S1 Text.

## Differentiable weight masking for subnetworks

Following previous work [23, 24, 29], we adopt a differentiable weight masking scheme to learn $\mathbf{m}$. Here, each binary mask parameter $m_i \in \{0, 1\}$ is sampled from a Gumbel distribution. We learn a logit $l_i \in \mathbb{R}$ for every $i$-th parameter and obtain a continuous mask score over the unit interval via the following Gumbel softmax transformation

$$s_i = \sigma\left\{-\frac{1}{\tau}\left[l_i + \log\left(\frac{U_1}{1 - U_1}\right)\right]\right\} \tag{8}$$

where $\sigma(\cdot)$ is the sigmoid function, $\tau$ is a temperature scaling hyperparameter, and $U_1 \sim \mathcal{U}(0, 1)$ is a random variable drawn from a standard uniform distribution. This Gumbel noise introduces stochasticity into the logit sampling process which results in exploring different binary mask configurations during training. We backpropagate through the collection of continuous mask scores $\mathbf{s}$ and threshold values to obtain binarized mask values $\mathbf{m}$ using the following

$$m_i = [\mathbb{1}(s_i > T) - s_i]_{\text{detach}} + s_i, \tag{9}$$

where $\mathbb{1}(\cdot)$ is an indicator function, $T$ is the mask score threshold, and $[\cdot]_{\text{detach}}$ prevents back-propagation through the discrete values in $\mathbf{m}$. This thresholding operation enables differentiable training by allowing gradients to flow through the continuous mask scores, while still computing the loss with respect to binary predictions [26, 30]. The sparsity is defined as the proportion of zeros in the learned binary mask—and therefore in the subnetwork—which is calculated as

$$\text{sparsity}(\mathbf{m}) = \frac{1}{K}\sum_{i=1}^{K} \mathbb{1}(m_i = 0), \tag{10}$$

where $\mathbf{m} = [m_1, \ldots, m_K]^\top \in \{0, 1\}^K$ and, again, $K$ is the total number of mask parameters. In our results, we represent this number as a percentage and multiply the value above by 100%. A description of how we select weight masking hyperparameters can be found in the S1 Text.

## Model architecture and datasets

In our main set of analyses, we learned subnetworks in the ESM-2 650 million parameter model [7]. While other binary masking approaches often focus on the final layers to capture fine-grained concepts or properties, we chose to learn the mask over the full model, i.e., all layers, which prevents any reliance on weak early-layer signals and spurious late-layer correlations. For additional validation, we also apply our subnetworks approach to three state-of-the-art PLMs of varying size, architectures, and pretraining tasks: ProtBERT-UR100 [2], CARP-640M [14], and Dayhoff-170M-UR90 [31]. Details on the weight masking scheme, choice of which hidden layers to mask, hyperparameters, and mask initialization values for each of these models can be found in S1 Text).

For training and evaluation, we used the CATH S20 version 4.3.0 release [32]. This dataset is comprised of CATH domains, which are sequences clustered at a maximum of 20% pairwise sequence identity with at least 60% alignment overlap. This ensures low redundancy while maintaining structural and functional diversity across domains. Every domain has an annotation at successive levels of the CATH hierarchy: Class (C), Architecture (A), Topology (T), and Homologous Superfamily (H). We used DSSP [28] to obtain reduced three-way residue-level secondary structure categories, which is the procedure from Ahdritz et al. [33]. During training, we filtered down to the set of CATH domains with length between 64 and 1024 residues (the maximum context window of ESM-2) and with available PDB structures. Altogether, this left 8886 CATH domains. For each subnetwork, we randomly split this data into 70% for training, 20% for validation, and 10% for a held out set. Subnetworks were learned according to the

suppression and maintenance inputs present in the train split. At evaluation time, we performed MLM evaluations over all splits, which allowed for aggregating performance on suppression and maintenance inputs that are both seen during training via the train split, but also unseen inputs in the validation and test splits. We performed the structure prediction evaluations with the ESMFold folding trunk on the validation and test sets to limit computational cost. We provide details on CATH data and annotation frequencies in Fig. S1.

# Results

We used trained subnetworks to assess whether structural categories are factorized in PLM weights and, in turn, whether this factorization affects downstream structure prediction. To this end, we first evaluated 39 trained subnetworks to assess MLM performance on suppression and maintenance inputs. Second, we passed subnetwork representations as inputs to the ESMFold folding trunk to evaluate the effect of the suppression on structure prediction capabilities (see Fig. 1C). In both sets of evaluations, we compared the effects of sequence-level suppression and residue-level suppression.

## Sparse subnetworks enable successful factorizations in ESM-2 weights

A structural category of proteins is considered to be factorized in PLM weights if we can identify a subnetwork that selectively reduces MLM performance on suppression inputs while maintaining performance on maintenance inputs, relative to the full PLM. We measured MLM performance using perplexity (i.e., the exponential of the negative log-likelihood), computed separately on the suppressed and maintained inputs for both the subnetworks and the pretrained ESM-2 model (Fig. 2A).

By our training procedure, a subnetwork can achieve this differential performance by identifying a sparse subgraph of PLM weights where parameters that encode information about the suppressed input category are zeroed out. As the subnetwork training procedure does not explicitly promote any sparse regularization, the percentage of learned sparsity is intrinsic to any discovered implicit factorization. We quantified the mean learned sparsity (defined in Eq. (10) and reported as a percentage) across structural levels in Fig. 2B. As suppressed categories become more fine-grained within the CATH hierarchy, the learned sparsity decreases, indicating that smaller structural categories can be factorized by zeroing out fewer parameters. This trend suggests a correlation between the learned percent sparsity of the subnetwork and the frequency of structural annotations (i.e., annotation granularity) in the data.

The trained subnetworks consistently yield higher perplexity on the suppression inputs, implying that structure-relevant representations can be selectively factorized in the PLM weight space (Fig. 2C). Targeted suppression of residues belonging to a secondary structure (i.e., alpha helices or beta strands) results in an approximate two-fold increase in perplexity of these residues. Suppression of broader sequence-level annotations, such as CATH Class where the suppression set sizes are largest, yields predictive distributions resembling random noise, with perplexity approaching 20—the value expected from uniform guessing over 20 amino acids. A similar increase in perplexity is observed when suppressing Homologous Superfamilies, despite much smaller suppression sets ($< 100$ sequences), suggesting that factorization is strongest at the coarsest and finest levels of annotation granularity. In contrast, structural categories are more weakly factorized at intermediate CATH levels like Architecture and Topology, as evidenced by lower perplexities on suppression inputs as compared to that of other CATH levels. One possible explanation for suppression being most effective at the Class and Homologous Superfamily levels is that these categories represent the most distinct structural and evolutionary boundaries in protein sequence space. Coarse Class categories capture global, fundamental secondary structure composition (i.e., mainly Alpha, mainly Beta, or "mixed" Alpha-Beta classes),

while fine-grained superfamilies reflect localized evolutionary relationships, both of which have been shown to form well-separated regions in PLM embedding space [1, 15, 34]. In contrast, intermediate categories such as Architecture and Topology are structurally heterogeneous, combining folds that share only partial motifs but differ in geometry, making them inherently less separable and thus more difficult for subnetworks to isolate into distinct groups. When training is repeated across multiple ESM-2 seeds under the same suppression objective, the resulting subnetworks exhibit highly similar sparsity patterns and achieve nearly identical performance on both suppression and maintenance inputs, indicating that optimization converges to stable and functionally equivalent solutions (Fig. S2 and Table S2). The configurations of all trained subnetworks are reported in Table S3, and the corresponding MLM performance is reported in Table S4.

As a control, we trained two types of subnetworks. The first control was to suppress $N$ randomly selected sequences, where values of $N$ were chosen to mimic the size of a CATH sequence category. The second control was to suppress randomly selected residues. We evaluated the residue control subnetwork separately on alpha helices and beta sheets. Both the random sequence and random residue subnetworks fail to produce differential performance on the suppressed inputs, with the only exception being the sequence control with $N = 2000$, which resulted in a 2.6-point increase in perplexity on suppression inputs relative to ESM-2 (8.8). In contrast, subnetworks trained to suppress CATH-Class level categories of sequences are able to attain a perplexity of greater than 30, while achieving maintenance goals. We reasoned that any increase in perplexity in the sequence-controls arises because the mask learning process removes parameters that encode features broadly shared across sequences. Since randomly selected sequences do not share meaningful structural characteristics, the subnetwork cannot identify parameters specific to suppressed structural information and instead must zero out weights important for general MLM performance, which can lead to a uniform increase in perplexity across all inputs.

We also applied our subnetworks approach to three additional state-of-the-art PLMs of varying size, architectures, and pretraining tasks, each having been trained on UniRef data: ProtBERT-UR100 (420M parameters) [2], a transformer masked language model with a BERT architecture; CARP-640M (640M parameters) [14], a convolutional neural network masked language model where transformer layers are replaced by ByteNet dilated CNN blocks; and Dayhoff-170M-UR90 (170M parameters) [31], an efficient hybrid state-space-model transformer trained with an autoregressive objective. We present these analyses as additional results in Fig. S3. The subnetworks in these additional PLMs show similar trends to ESM-2, in that they are discovered by pruning less than 3% of model parameters, are most strongly factorized at the Class and Homologous Superfamily levels, and factorize categories of sequences to a larger extent when compared to residues.

Together, the percent sparsity and MLM evaluations show that subnetworks can increase perplexity in a targeted manner by identifying sparse factorizations of structural information in PLMs, in a manner that is consistent with the continuous nature of protein structural diversity.

## Subnetworks perturb structure prediction accuracy on both suppressed and maintained inputs

Having evaluated subnetworks on the language modeling task, we then assessed how suppression via a subnetwork affects structure prediction accuracy. ESMFold [7] introduced a folding trunk—a simplified version of AlphaFold2's Evoformer [35]—that converts language model representations into 3D structures. Using this frozen trunk allows the isolation of subnetwork-induced changes in sequence representations and enables investigation into how they affect structure prediction accuracy. Since the folding trunk serves as a fixed decoder, any degradation in template modeling (TM) score or predicted local distance difference test (pLDDT), or increased root mean square deviation (RMSD), directly reflects a loss of relevant structural

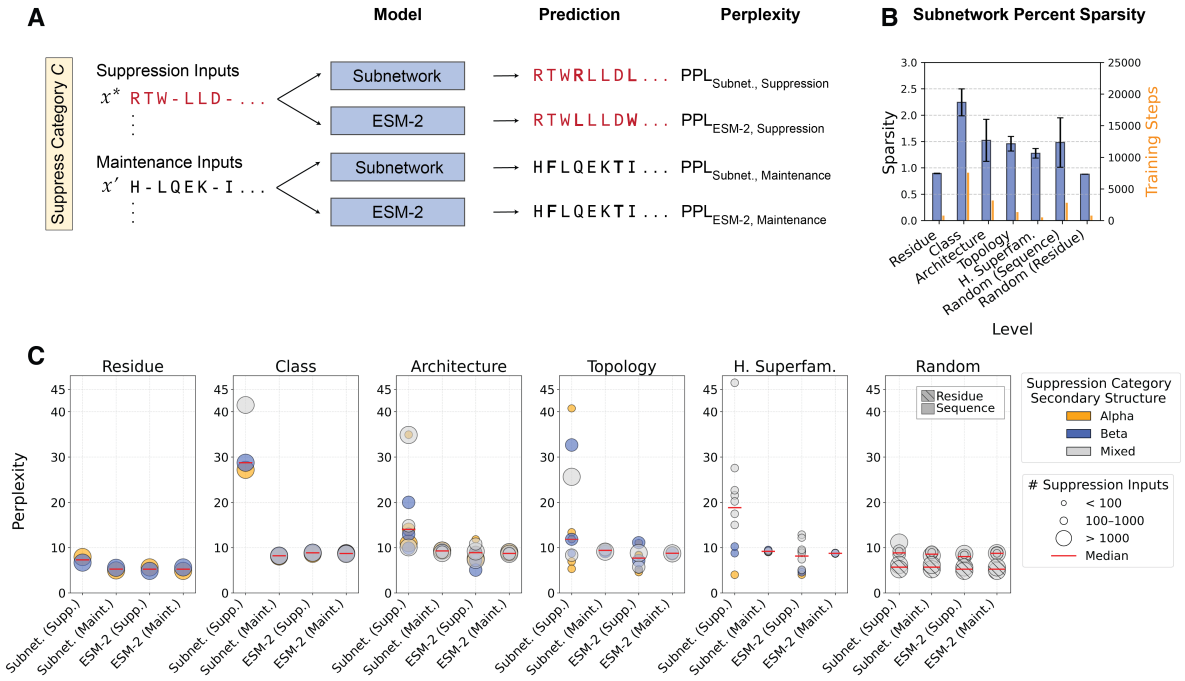


**Figure 2. Ablating less than 3% of ESM-2 parameters increases perplexity on suppressed inputs and reveals successful factorizations of structural categories in PLM weights. (A)** For a given subnetwork, suppression and maintenance inputs are defined by their structural annotation. At evaluation time, the same sequence is fed as separate inputs to the subnetwork and ESM-2. ESM-2 performance on the same suppression and maintenance inputs, labeled as "ESM-2 Supp." and "ESM-2 Maint." respectively, provides a baseline for the subnetwork. (B) Learned sparsity of each subnetwork across categories of structural annotations (reported mean with ± standard deviation depicted by the error bars). The right y-axis shows the average number of optimization steps in yellow. **(C)** Perplexity (y-axis) on test sequences across suppression and maintenance input categories (x-axis). Each point represents an independently trained subnetwork, where marker size indicates number of suppression inputs in the structural category and color indicates its predominant secondary structure type. Suppression and maintenance inputs are defined per subnetwork; ESM-2 points reflect baseline perplexity on those same inputs.

information in the sequence representations. For each input sequence, we extracted sequence representations from the subnetwork and from ESM-2, and we passed them as separate inputs to the ESMFold folding trunk for structure prediction. Both predicted structures were independently aligned to the ground truth PDB structure to obtain a TM score, RMSD, and pLDDT for each model's prediction. This procedure was repeated on all suppression inputs $\mathbf{x}^*$ and all maintenance inputs $\mathbf{x}'$ in the validation datasets (Fig. 3A).

First, we evaluated whether the subnetwork leads to a decrease in structure prediction accuracy relative to the ESM-2 performance on these inputs (Fig. 3B). For each input, we computed the change in RMSD where $\Delta\text{RMSD} = \text{RMSD}_{\text{Subnet., PDB}} - \text{RMSD}_{\text{ESM-2, PDB}}$. These differences are consistently larger on suppression inputs and minor on maintenance inputs (Fig. 3B). This suggests that the subnetwork is able to ablate structural information pertaining to the suppressed inputs in PLM weights. We then performed separate paired t-tests for the inputs in the suppression and maintenance sets, comparing $\text{RMSD}_{\text{Subnet., PDB}}$ and $\text{RMSD}_{\text{ESM-2, PDB}}$. Both tests result in significant $p$-values, suggesting that despite targeting only suppression inputs, the subnetwork also significantly affects structure prediction on maintenance inputs. That is,

$\Delta\text{RMSD}_{\text{supp.}}$ and $\Delta\text{RMSD}_{\text{maint.}}$ are both statistically significant. We repeated this procedure using the TM-score and pLDDT and observed a similar trend (Fig. S4A-B). We reason that the combination of pronounced effects on suppression inputs and minimal changes in structure prediction accuracy on maintenance inputs—together with the significant $p$-values—indicates that the frozen ESMFold trunk is sensitive to subtle representational shifts, rather than reflecting a lack of modularization in sequence representations. The subnetwork mask learning procedure was not optimized to preserve structure prediction accuracy; instead, it was designed to partition the PLM's weight space based on sequence-level characterization and MLM performance. Consequently, small variations in downstream structural metrics are expected, and this analysis mainly serves to reveal how sequence-level factorizations influence structure prediction capabilities.

Next, we assessed whether the magnitude of these differences is greater for suppression inputs than for maintenance inputs across CATH and secondary structure categories (Fig. 3C). That is, we seek to confirm that $|\Delta\text{RMSD}_{\text{supp.}}| - |\Delta\text{RMSD}_{\text{maint.}}| > 0$ (Fig. 3C). These differences are consistently positive, suggesting that a subnetwork always results in a greater increase in RMSD on the suppression inputs when compared to maintenance inputs. We also confirm that the differences in magnitude are consistently statistically significant by applying a two-sample Kolmogorov–Smirnov (KS) test between $|\Delta\text{RMSD}_{\text{supp.}}|$ and $|\Delta\text{RMSD}_{\text{maint.}}|$, rejecting the null hypothesis that the suppression and maintenance inputs exhibit the same distributions of difference in magnitudes of RMSD increase. We repeated the same analysis for TM-score and pLDDT and observed the same trend (Fig. S4C–D). We also observe that the difference in magnitudes of change in each metric is, on average, larger for Alpha-Beta proteins, followed by Mainly Beta and then Mainly Alpha proteins. We reason that this could be due to stricter structural constraints imposed by the mixed secondary-structure composition of Alpha-Beta folds, which couple alpha helices and beta sheets into interdependent architectures. Consistent with this interpretation, we used ProteinMPNN perplexities to evaluate sequence-level structural constraints and find a supporting trend: Alpha-Beta proteins exhibit the lowest perplexities, followed by Mainly Beta and then Mainly Alpha proteins, indicating that mixed Alpha-Beta folds have the most restricted sequence variability and strongest structure–sequence coupling (Table S5). All structure prediction metrics and $p$-values are reported in Table S6.

## Structure prediction is more sensitive to suppressed sequence features

Alpha helices and beta sheets are local secondary structure elements that, when predominant in a sequence, define the overall CATH Class of a protein domain. Sequences consisting of majority alpha helices form "Mainly Alpha" domains, while those rich in beta sheets form "Mainly Beta" domains at the CATH Class level. This connection allows us to directly compare the change in pLDDT between structures predicted by ESM-2 and those predicted by subnetworks trained to suppress (i) sequences labeled by CATH Class (e.g., Mainly Alpha or Mainly Beta), and (ii) residues belonging to specific secondary structures (e.g., alpha helices or beta sheets). We illustrated predicted structures for alpha suppression subnetworks in Fig. 3D and beta suppression subnetworks in Fig. 3E, where in each panel the top row is a suppressed input and the bottom row is a maintained input. While suppressing at the residue level based on secondary structure annotations only affects the mean pLDDT of the full predicted structure, suppressing sequences leads to no clear predicted fold or structure pattern. Since pLDDT is a per residue metric, we computed residue-specific pLDDT of predicted structures for both residue- and sequence-level suppression with alpha type (Fig. 3F) and beta type (Fig. 3G). Suppressing sequences still leads to a lower residue-specific pLDDT of the target residue type for both alpha and beta suppression subnetworks. This suggests that factorization of structural information in PLM weights is more sensitive to sequence-level suppression and, by extension, that PLMs more effectively model distributions at the sequence level.

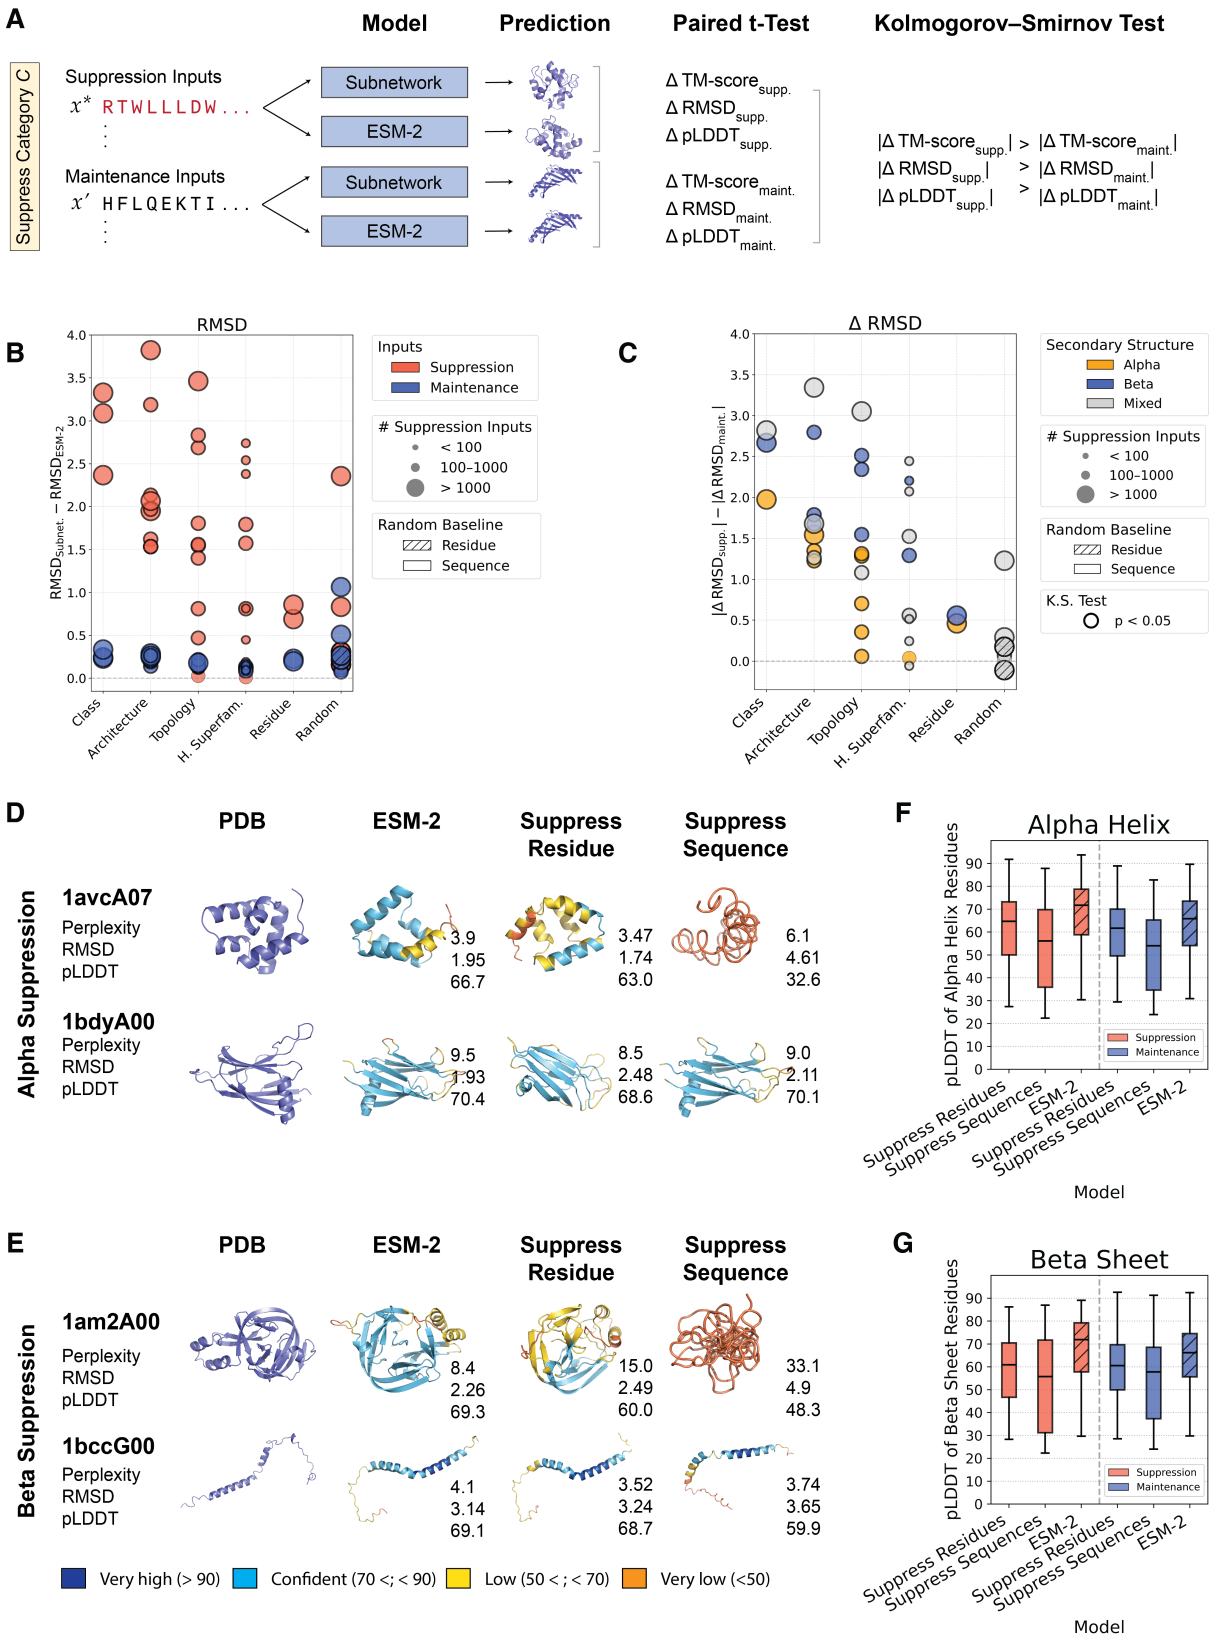
A
Model
Prediction
Paired t-Test
Kolmogorov–Smirnov Test
Suppress Category C
Suppression Inputs
Maintenance Inputs
Subnetwork
ESM-2
Δ TM-score supp.
Δ RMSD supp.
Δ pLDDT supp.
Δ TM-score maint.
Δ RMSD maint.
Δ pLDDT maint.
|Δ TM-score supp.| > |Δ TM-score maint.|
|Δ RMSD supp.| > |Δ RMSD maint.|
|Δ pLDDT supp.| > |Δ pLDDT maint.|
B
RMSD
Inputs
Suppression
Maintenance
# Suppression Inputs
< 100
100–1000
> 1000
Random Baseline
Residue
Sequence
Class
Architecture
Topology
H. Superfam.
Residue
Random
C
Δ RMSD
Secondary Structure
Alpha
Beta
Mixed
K.S. Test
p < 0.05
D
PDB
ESM-2
Suppress Residue
Suppress Sequence
Alpha Suppression
1avcA07
Perplexity
RMSD
pLDDT
3.9
1.95
66.7
3.47
1.74
63.0
6.1
4.61
32.6
1bdyA00
Perplexity
RMSD
pLDDT
9.5
1.93
70.4
8.5
2.48
68.6
9.0
2.11
70.1
F
Alpha Helix
pLDDT of Alpha Helix Residues
Suppression
Maintenance
Suppress Residues
Suppress Sequences
ESM-2
Model
E
PDB
ESM-2
Suppress Residue
Suppress Sequence
Beta Suppression
1am2A00
Perplexity
RMSD
pLDDT
8.4
2.26
69.3
15.0
2.49
60.0
33.1
4.9
48.3
1bccG00
Perplexity
RMSD
pLDDT
4.1
3.14
69.1
3.52
3.24
68.7
3.74
3.65
59.9
Very high (> 90)
Confident (70 <; < 90)
Low (50 < ; < 70)
Very low (<50)
G
Beta Sheet
pLDDT of Beta Sheet Residues
Suppression
Maintenance
Suppress Residues
Suppress Sequences
ESM-2
Model

**Figure 3. Subnetworks consistently impair structure prediction capabilities on suppressed inputs. (A)** Sequence representations for a given input are obtained from the subnetwork and ESM-2 and fed as inputs to the ESMFold folding trunk to produce two predicted structures. Both structures are independently aligned to the ground truth PDB structure, and metrics are calculated for evaluations. **(B)** RMSD differences between the subnetwork and ESM-2 baseline (y-axis) across structural levels (x-axis). Each point represents the RMSD increase for suppression inputs (red) or maintenance inputs (blue) relative to ESM-2. Marker size indicates the number of suppression inputs for each subnetwork. Bold outlines of markers indicate significant paired t-test $p$-values ($p < 0.05$). **(C)** Difference in absolute RMSD changes from the ESM-2 baseline (y-axis), stratified by structural level (x-axis). Each point corresponds to an individual subnetwork and shows the difference between the magnitudes of suppression and maintenance $\Delta$RMSD values. Marker size reflects the number of suppressed inputs in the subnetwork, and color indicates secondary structure type. Bold outlines of markers indicate significance by Kolmogorov–Smirnov (KS) test ($p < 0.05$). **(D–E)** Visualizations and evaluation metrics for subnetwork predictions on select validation targets under (D) alpha suppression and (E) beta suppression. For each example, the ground truth PDB structure, ESM-2 prediction, and subnetwork predictions under residue-level and sequence-level suppression are shown. Each prediction is annotated with perplexity, mean RMSD, and mean pLDDT scores. **(F–G)** Distributions of predicted per-residue pLDDT scores across (F) alpha helix and (G) beta sheet residues across suppression and maintenance conditions, for subnetworks and baseline ESM-2 predictions. Each box shows the median and interquartile range.

# Discussion

We introduce subnetwork discovery as a post-hoc mechanistic approach to explore the modularity of learned representations in pretrained protein language models (PLMs). Because subnetworks are extracted without any fine-tuning, they expose relationships already present in the pretrained weights. The method is therefore a lightweight alternative to concept-level and other supervision-heavy interpretable pretraining schemes [36]. Our experiments show that PLMs can effectively disentangle both coarse and fine-grained CATH structural categories.

We chose ESM-2 for our main analyses because in addition to it being widely adopted and well-studied, it has an associated structure module that allows us to directly probe structure information in language modeling representations. We chose the 650M variant because it offers strong performance on both language modeling and structure prediction related tasks with reasonable computational cost. Learning a mask over model parameters effectively doubles GPU memory usage, and the 650M variant allowed us to train each subnetwork efficiently on a single H100 machine. Smaller ESM-2 variants yield less reliable structure prediction performance, while larger variants would be prohibitively expensive for the full set of experiments we presented. Since the performance of ESM-2 has been shown to scale with model size [7, 17], we expect the subnetwork behavior observed here to exhibit similar scaling trends. In fact, our results offer a mechanistic explanation for this behavior: if structural information is organized within localized subsets of weights, then increasing model capacity expands the number of structural categories that can be represented, enhancing both representational and predictive performance. This interpretation is consistent with the view that PLMs encode statistics of co-evolving residues [19], as larger models trained on more diverse categories of sequences can store more subnetworks. We further find that subnetwork behavior is consistent across PLMs of different sizes, architectures, and training objectives, including ProtBERT-UR100 [2], CARP-640M [14], and Dayhoff-170M-UR90 [31]. This suggests that structural factorization is a general property of protein representation learning rather than an artifact of a specific model or training setup. The learned masks themselves provide further insight into where structural

information is stored within a given PLM (Figs. S5-S8). In ESM-2, the subnetwork primarily zeros out parameters in mid-to-late layers (12–33), which aligns well with prior work showing that PLMs learn secondary structure features through superposition, with later layers becoming increasingly specialized for this purpose [20, 21]. These layers likely correspond to higher-level abstractions of secondary and tertiary structure, analogous to how natural language models progressively learn linguistic hierarchy by encoding lexical and syntactic features in lower layers and semantic abstractions in higher layers [37, 38].

Our training procedure operates directly in sequence space, and our results highlight the importance of broader sequence context in protein representation learning: we find that sequence-level categories are much more strongly factorized than residue-types. We found that while factorized PLM representations are similarly organized in the folding trunk weights of ESMFold, structure prediction accuracy is significantly perturbed when using representations from a subnetwork. The modest perturbations in structure prediction for maintenance inputs underscore that perplexity alone, typically the hallmark metric to evaluate PLMs, is not a sufficient measure of representational quality: even when a subnetwork reproduces ESM-2-level perplexity, small shifts in embedding distributions can propagate through the frozen ESMFold trunk and alter predicted structures.

Despite these successes, our work has some limitations. First, although comprehensive, our analysis is confined to the CATH dataset and to fundamental secondary structure categories. While assessing generalization to multi-domain proteins and complexes is an interesting future direction, the goal of this work is to characterize how structural categories are factorized in PLM weights at the level of their basic, clearly labeled secondary structure components. This approach relies on annotations which can be costly and time-consuming to curate, but are increasingly available with rapidly expanding with deep-learning based pipelines [39]. Second, our evaluation of structure prediction relies on the ESMFold folding trunk as the only available decoder capable of translating PLM representations into 3D coordinates. Since ESMFold is specifically designed for ESM-2 embeddings, this restricts our ability to perform comparable analyses with other protein language models such as ProtBERT, CARP, and Dayhoff. Extending the folding trunk framework to models with different architectures or pretraining objectives would provide valuable insight into how these subnetwork representations directly influence structure prediction capabilities.

In sum, we envision that our approach will help accelerate PLM development and interpretability, as our framework can be readily extended to any labeled protein sequence dataset. It provides a lens to understand how factorizations in representation space affect downstream structural prediction tasks. As masked language model-based PLMs continue to scale, understanding how these models synthesize and organize biological information will become increasingly important.

# Code and data availability

All code is available under an open-source MIT license at https://github.com/microsoft/plm_subnetworks. Training data, model configurations and checkpoints, and results are available via the Zenodo link in the repository.

# Acknowledgments

This research was conducted using a combination of computational resources and services provided by Microsoft Research and the Center for Computation and Visualization at Brown University. This research was also supported in part by a David & Lucile Packard Fellowship for Science and Engineering awarded to LC. Any opinions, findings, and conclusions or recommen-

dations expressed in this material are those of the author(s) and do not necessarily reflect the views of any of the funders.

## Author contributions

RV, APA, LC, and KKY conceived the study. RV developed the methods, software, and performed the analyses. LC provided resources. APA, LC, and KKY supervised the project. All authors interpreted the results and wrote the manuscript.

## Declaration of interests

The authors have declared that no competing interests exist.

## References

[1] Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Jason Liu, Demi Guo, Myle Ott, C. Lawrence Zitnick, Jerry Ma, and Rob Fergus. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15):e2016239118, April 2021. doi: 10.1073/pnas.2016239118. URL https://www.pnas.org/doi/10.1073/pnas.2016239118. Publisher: Proceedings of the National Academy of Sciences.

[2] Ahmed Elnaggar, Michael Heinzinger, Christian Dallago, Ghalia Rehawi, Yu Wang, Llion Jones, Tom Gibbs, Tamas Feher, Christoph Angerer, Martin Steinegger, Debsindhu Bhowmik, and Burkhard Rost. ProtTrans: Towards Cracking the Language of Life's Code Through Self-Supervised Learning, May 2021. URL https://www.biorxiv.org/content/10.1101/2020.07.12.199554v3. Pages: 2020.07.12.199554 Section: New Results.

[3] Roshan Rao, Jason Liu, Robert Verkuil, Joshua Meier, John F. Canny, Pieter Abbeel, Tom Sercu, and Alexander Rives. MSA Transformer, August 2021. URL https://www.biorxiv.org/content/10.1101/2021.02.12.430858v3. Pages: 2021.02.12.430858 Section: New Results.

[4] Michael Heinzinger, Ahmed Elnaggar, Yu Wang, Christian Dallago, Dmitrii Nechaev, Florian Matthes, and Burkhard Rost. Modeling aspects of the language of life through transfer-learning protein sequences. *BMC bioinformatics*, 20(1):723, December 2019. ISSN 1471-2105. doi: 10.1186/s12859-019-3220-8.

[5] Ruidong Wu, Fan Ding, Rui Wang, Rui Shen, Xiwen Zhang, Shitong Luo, Chenpeng Su, Zuofan Wu, Qi Xie, Bonnie Berger, Jianzhu Ma, and Jian Peng. High-resolution de novo structure prediction from primary sequence, July 2022. URL https://www.biorxiv.org/content/10.1101/2022.07.21.500999v1. Pages: 2022.07.21.500999 Section: New Results.

[6] Ratul Chowdhury, Nazim Bouatta, Surojit Biswas, Charlotte Rochereau, George M. Church, Peter K. Sorger, and Mohammed AlQuraishi. Single-sequence protein structure prediction using language models from deep learning, August 2021. URL https://www.biorxiv.org/content/10.1101/2021.08.02.454840v1. Pages: 2021.08.02.454840 Section: New Results.

[7] Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, Allan dos Santos Costa, Maryam

Fazel-Zarandi, Tom Sercu, Salvatore Candido, and Alexander Rives. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637): 1123–1130, March 2023. doi: 10.1126/science.ade2574. URL https://www.science.org/doi/10.1126/science.ade2574. Publisher: American Association for the Advancement of Science.

[8] Roshan Rao, Nicholas Bhattacharya, Neil Thomas, Yan Duan, Xi Chen, John Canny, Pieter Abbeel, and Yun S. Song. Evaluating Protein Transfer Learning with TAPE, June 2019. URL http://arxiv.org/abs/1906.08230. arXiv:1906.08230 [cs].

[9] Samir Char, Nathaniel Corley, Sarah Alamdari, Kevin K Yang, and Ava P Amini. Protnote: a multimodal method for protein–function annotation. *Bioinformatics*, 41(5):btaf170, 04 2025. ISSN 1367-4811. doi: 10.1093/bioinformatics/btaf170. URL https://doi.org/10.1093/bioinformatics/btaf170.

[10] Joshua Meier, Roshan Rao, Robert Verkuil, Jason Liu, Tom Sercu, and Alexander Rives. Language models enable zero-shot prediction of the effects of mutations on protein function, November 2021. URL https://www.biorxiv.org/content/10.1101/2021.07.09.450648v2. Pages: 2021.07.09.450648 Section: New Results.

[11] Erik Nijkamp, Jeffrey Ruffolo, Eli N. Weinstein, Nikhil Naik, and Ali Madani. ProGen2: Exploring the Boundaries of Protein Language Models, June 2022. URL http://arxiv.org/abs/2206.13517. arXiv:2206.13517 [cs].

[12] Sarah Alamdari, Nitya Thakkar, Rianne van den Berg, Neil Tenenholtz, Robert Strome, Alan M. Moses, Alex X. Lu, Nicolò Fusi, Ava P. Amini, and Kevin K. Yang. Protein generation with evolutionary diffusion: sequence is all you need, November 2024. URL https://www.biorxiv.org/content/10.1101/2023.09.11.556673v2. Pages: 2023.09.11.556673 Section: New Results.

[13] Thomas Hayes, Roshan Rao, Halil Akin, Nicholas J. Sofroniew, Deniz Oktay, Zeming Lin, Robert Verkuil, Vincent Q. Tran, Jonathan Deaton, Marius Wiggert, Rohil Badkundri, Irhum Shafkat, Jun Gong, Alexander Derry, Raul S. Molina, Neil Thomas, Yousuf A. Khan, Chetan Mishra, Carolyn Kim, Liam J. Bartie, Matthew Nemeth, Patrick D. Hsu, Tom Sercu, Salvatore Candido, and Alexander Rives. Simulating 500 million years of evolution with a language model, December 2024. URL https://www.biorxiv.org/content/10.1101/2024.07.01.600583v2. Pages: 2024.07.01.600583 Section: New Results.

[14] Kevin K. Yang, Nicolo Fusi, and Alex X. Lu. Convolutions are competitive with transformers for protein sequence pretraining. *Cell Systems*, 15(3):286–294.e2, March 2024. ISSN 2405-4712. doi: 10.1016/j.cels.2024.01.008. URL https://www.sciencedirect.com/science/article/pii/S2405471224000292.

[15] Ethan C. Alley, Grigory Khimulya, Surojit Biswas, Mohammed AlQuraishi, and George M. Church. Unified rational protein engineering with sequence-based deep representation learning. *Nature Methods*, 16(12):1315–1322, December 2019. ISSN 1548-7105. doi: 10.1038/s41592-019-0598-1.

[16] Bo Chen, Xingyi Cheng, Pan Li, Yangli-ao Geng, Jing Gong, Shen Li, Zhilei Bei, Xu Tan, Boyan Wang, Xin Zeng, Chiming Liu, Aohan Zeng, Yuxiao Dong, Jie Tang, and Le Song. xTrimoPGLM: Unified 100B-Scale Pre-trained Transformer for Deciphering the Language of Protein, December 2024. URL http://arxiv.org/abs/2401.06199. arXiv:2401.06199 [q-bio].

[17] 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, February 2024. URL https://www.biorxiv.org/content/10.1101/2024.02.05.578959v2. Pages: 2024.02.05.578959 Section: New Results.

[18] Jesse Vig, Ali Madani, Lav R. Varshney, Caiming Xiong, Richard Socher, and Nazneen Fatema Rajani. BERTology Meets Biology: Interpreting Attention in Protein Language Models, March 2021. URL http://arxiv.org/abs/2006.15222. arXiv:2006.15222 [cs].

[19] Zhidian Zhang, Hannah K. Wayment-Steele, Garyk Brixi, Haobo Wang, Dorothee Kern, and Sergey Ovchinnikov. Protein language models learn evolutionary statistics of interacting sequence motifs. *Proceedings of the National Academy of Sciences*, 121(45): e2406285121, November 2024. doi: 10.1073/pnas.2406285121. URL https://www.pnas.org/doi/10.1073/pnas.2406285121. Publisher: Proceedings of the National Academy of Sciences.

[20] Elana Simon and James Zou. InterPLM: Discovering Interpretable Features in Protein Language Models via Sparse Autoencoders, January 2025. URL https://www.biorxiv.org/content/10.1101/2024.11.14.623630v2. Pages: 2024.11.14.623630 Section: New Results.

[21] Etowah Adams, Liam Bai, Minji Lee, Yiyang Yu, and Mohammed AlQuraishi. From Mechanistic Interpretability to Mechanistic Biology: Training, Evaluating, and Interpreting Sparse Autoencoders on Protein Language Models, February 2025. URL https://www.biorxiv.org/content/10.1101/2025.02.06.636901v1. Pages: 2025.02.06.636901 Section: New Results.

[22] Deniz Bayazit, Negar Foroutan, Zeming Chen, Gail Weiss, and Antoine Bosselut. Discovering Knowledge-Critical Subnetworks in Pretrained Language Models. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6549–6583, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.376. URL https://aclanthology.org/2024.emnlp-main.376/.

[23] Boxi Cao, Hongyu Lin, Xianpei Han, Le Sun, Lingyong Yan, Meng Liao, Tong Xue, and Jin Xu. Knowledgeable or Educated Guess? Revisiting Language Models as Knowledge Bases. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli, editors, *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1860–1874, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.146. URL https://aclanthology.org/2021.acl-long.146/.

[24] Victor Sanh, Thomas Wolf, and Alexander M. Rush. Movement Pruning: Adaptive Sparsity by Fine-Tuning, October 2020. URL http://arxiv.org/abs/2005.07683. arXiv:2005.07683 [cs].

[25] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a Single Network to Multiple Tasks by Learning to Mask Weights, March 2018. URL http://arxiv.org/abs/1801.06519. arXiv:1801.06519 [cs].

[26] Róbert Csordás, Sjoerd van Steenkiste, and Jürgen Schmidhuber. Are Neural Nets Modular? Inspecting Functional Modularity Through Differentiable Weight Masks, March 2021. URL http://arxiv.org/abs/2010.02066. arXiv:2010.02066 [cs].

[27] Xiongyi Zhang, Jan-Willem van de Meent, and Byron Wallace. Disentangling Representations of Text by Masking Transformers. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 778–791, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.60. URL https://aclanthology.org/2021.emnlp-main.60/.

[28] W. Kabsch and C. Sander. Dictionary of protein secondary structure: pattern recognition of hydrogen-bonded and geometrical features. *Biopolymers*, 22(12):2577–2637, December 1983. ISSN 0006-3525. doi: 10.1002/bip.360221211.

[29] Steven Cao, Victor Sanh, and Alexander M. Rush. Low-Complexity Probing via Finding Subnetworks, April 2021. URL http://arxiv.org/abs/2104.03514. arXiv:2104.03514 [cs].

[30] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation, August 2013. URL http://arxiv.org/abs/1308.3432. arXiv:1308.3432 [cs].

[31] Kevin K. Yang, Sarah Alamdari, Alex J. Lee, Kaeli Kaymak-Loveless, Samir Char, Garyk Brixi, Carles Domingo-Enrich, Chentong Wang, Suyue Lyu, Nicolo Fusi, Neil Tenenholtz, and Ava P. Amini. The Dayhoff Atlas: scaling sequence diversity for improved protein generation, July 2025. URL https://www.biorxiv.org/content/10.1101/2025.07.21.665991v1. ISSN: 2692-8205 Pages: 2025.07.21.665991 Section: New Results.

[32] CA Orengo, AD Michie, S Jones, DT Jones, MB Swindells, and JM Thornton. CATH – a hierarchic classification of protein domain structures. *Structure*, 5(8):1093–1109, August 1997. ISSN 0969-2126. doi: 10.1016/S0969-2126(97)00260-8. URL https://www.sciencedirect.com/science/article/pii/S0969212697002608.

[33] Gustaf Ahdritz, Nazim Bouatta, Christina Floristean, Sachin Kadyan, Qinghui Xia, William Gerecke, Timothy J. O'Donnell, Daniel Berenberg, Ian Fisk, Niccolò Zanichelli, Bo Zhang, Arkadiusz Nowaczynski, Bei Wang, Marta M. Stepniewska-Dziubinska, Shang Zhang, Adegoke Ojewole, Murat Efe Guney, Stella Biderman, Andrew M. Watkins, Stephen Ra, Pablo Ribalta Lorenzo, Lucas Nivon, Brian Weitzner, Yih-En Andrew Ban, Peter K. Sorger, Emad Mostaque, Zhao Zhang, Richard Bonneau, and Mohammed AlQuraishi. OpenFold: Retraining AlphaFold2 yields new insights into its learning mechanisms and capacity for generalization, August 2023. URL https://www.biorxiv.org/content/10.1101/2022.11.20.517210v3. Pages: 2022.11.20.517210 Section: New Results.

[34] Tristan Bepler and Bonnie Berger. Learning protein sequence embeddings using information from structure. *CoRR*, abs/1902.08661, 2019. URL http://arxiv.org/abs/1902.08661.

[35] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, August 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03819-2. URL https://www.nature.com/articles/s41586-021-03819-2. Publisher: Nature Publishing Group.

[36] Aya Abdelsalam Ismail, Tuomas Oikarinen, Amy Wang, Julius Adebayo, Samuel Stanton, Taylor Joren, Joseph Kleinhenz, Allen Goodman, Héctor Corrada Bravo, Kyunghyun Cho, and Nathan C. Frey. Concept Bottleneck Language Models For protein design, December 2024. URL http://arxiv.org/abs/2411.06090. arXiv:2411.06090 [cs].

[37] Ian Tenney, Dipanjan Das, and Ellie Pavlick. BERT rediscovers the classical NLP pipeline. *CoRR*, abs/1905.05950, 2019. URL http://arxiv.org/abs/1905.05950.

[38] Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. Linguistic knowledge and transferability of contextual representations. *CoRR*, abs/1903.08855, 2019. URL http://arxiv.org/abs/1903.08855.

[39] Andy M. Lau, Nicola Bordin, Shaun M. Kandathil, Ian Sillitoe, Vaishali P. Waman, Jude Wells, Christine A. Orengo, and David T. Jones. Exploring structural diversity across the protein universe with the encyclopedia of domains. *Science*, 386(6721):eadq4946, 2024. doi: 10.1126/science.adq4946. URL https://www.science.org/doi/abs/10.1126/science.adq4946.

# S1 Supplementary Text

## S1.1 Subnetwork training details

For each PLM, we conducted a grid search to select optimal hyperparameter values to achieve suppression in all subnetworks. We set the below hyperparameters to the same values across all subnetworks learned within the same model. We include the results of the grid search on subnetwork performance in the accompanying code and data (see https://github.com/microsoft/plm_subnetworks for more information). The hyperparameters that vary by PLM are listed below and summarized in Table S1.

- **Mask initialization ($s_{init}$).** Each mask logit $l_i$ is initialized such that $\sigma(l_i) = p_{\text{keep},0}$, where $p_{\text{keep},0}$ is the desired initial retention probability (e.g., 0.96) of the $i$-th parameter. This is achieved by setting $l_i = \log\left[p_{\text{keep},0}/(1 - p_{\text{keep},0})\right]$, so that each parameter begins training with a high probability of being retained. This initialization enables the model to start from a near-complete network and learn to suppress parameters selectively over time. We choose $p_{\text{keep},0}$ to ensure that the subnetwork MLM loss initially matches the PLM average loss over all inputs, while introducing less than approximately 1% sparsity. This yields a mask score distribution centered near zero, placing most sigmoid activations in a high-sensitivity regime and allowing flexible updates toward maintenance and suppression goals. We explored the effects of this random initialization of mask scores in Fig. S2.
- **Suppression-KL loss weight ($\lambda_{\text{supp.}}$).** Applied using Eq. (1) and Eq. (2).
- **Maintenance-KL loss weight ($\lambda_{\text{maint.}}$).** Applied using Eq. (3) and Eq. (4).
- **Maintenance-MLM loss weight ($\lambda_{\text{MLM}}$).** Applied using Eq. (5) and Eq. (6).
- **Mask temperature ($\tau$).** We apply a cosine decay schedule to the temperature $\tau$, decreasing it smoothly from a maximum of $\tau_{\text{init}}$ to $\tau_{\text{final}}$ over 100 epochs. This schedule enables high-entropy exploration early in training and progressively sharper binarization of the mask as training progresses.
- **Mask threshold ($T$).** In practice, the binarization threshold $T$ in Eq. (9) varies by PLM.

## S1.2 Learning rate

We use a maximum learning rate of `peak_lr` with a reverse OneCycle learning rate schedule. The learning rate remains constant at `peak_lr` during an initial plateau phase, defined by `1 / lr_plateau_epochs`, and then decays to `peak_lr` divided by 30 using cosine annealing. The learning rate schedule spans the full number of estimated training steps, as determined by PyTorch Lightning. Only mask parameters are optimized, using the AdamW optimizer with $\beta = (0.9, 0.999)$, $\epsilon = 10^{-6}$, and zero weight decay. The maximum value of `peak_lr` observed in training is $10^{-1}$.

## S1.3 Compute

We train each subnetwork on one NVIDIA H100 GPU. We curate training batches of size 16 with a 0.25 frequency of suppression inputs in a batch such that subnetworks training converges over fewer epochs. We report the number of suppression inputs per category and the number training steps in Table S3.

# Supplementary Figures

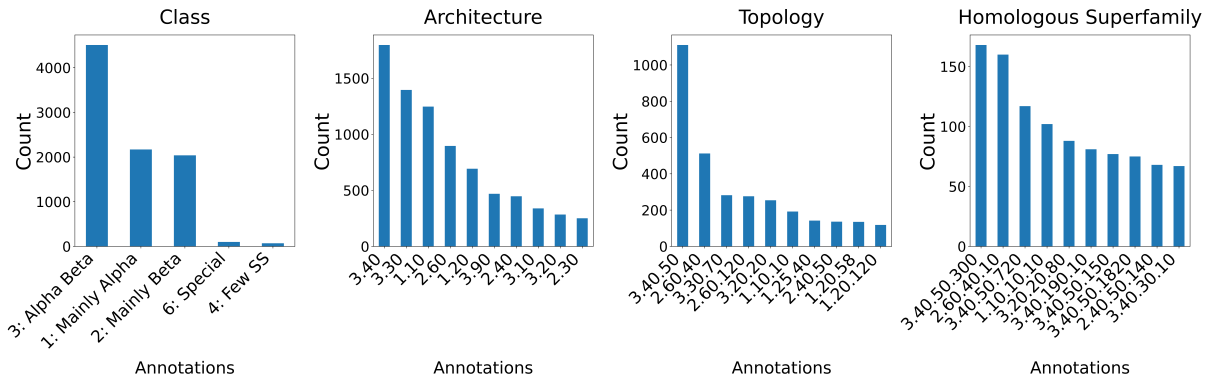


**(A) Annotation frequencies by CATH levels.** Each CATH domain is annotated with a label at the Class, Architecture, Topology, and Homologous Superfamily levels. Bar plots show the counts (y-axis) of the top 10 most frequent annotations (x-axis) at each level of the CATH hierarchy.

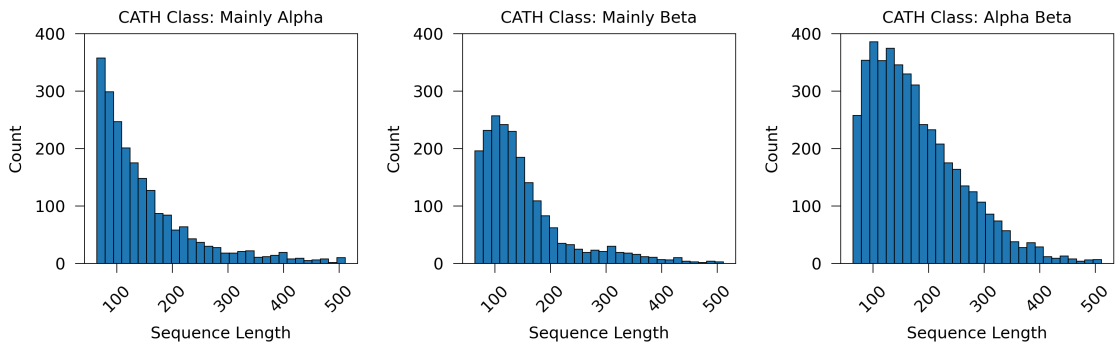


**(B) Sequence length distributions stratified by CATH level.** Histograms show the counts (y-axis) of domain sequence lengths (x-axis) for each CATH Class: Mainly Alpha, Mainly Beta, Alpha Beta.

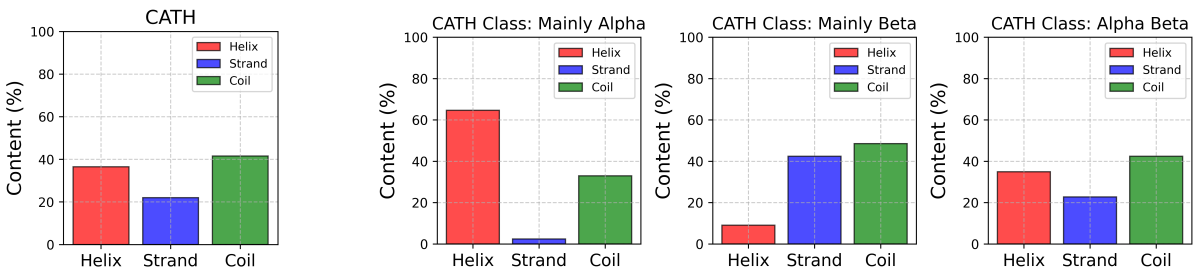


**(C) Secondary structure composition of CATH domains.** (Left) Average fraction of residues annotated as helix, strand, or coil across all CATH domains, based on DSSP annotations [28]. (Right) Breakdown of secondary structure content stratified by CATH class. DSSP 8-state annotations are mapped to 3-state labels: H, G, I → H; E, B → E; T, S, - → L. Helix is H, beta strand is E, and loop is L.

**Figure S1. CATH annotation characteristics.**

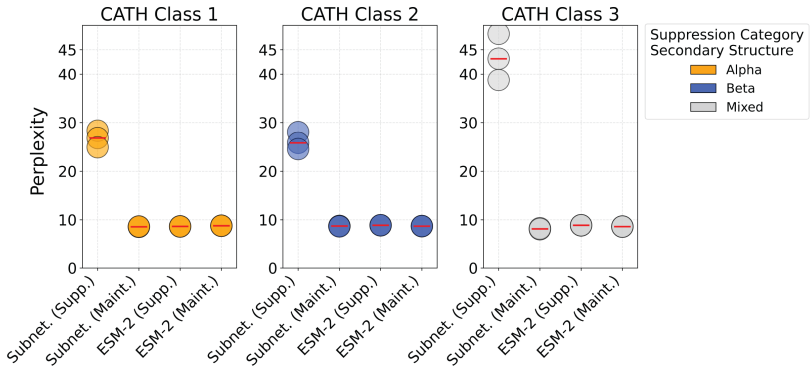


**Figure S2. ESM-2 CATH Class-level subnetwork language modeling performance across multiple seeds.** Three subnetworks were independently trained for each CATH Class suppression target (Mainly Alpha, Mainly Beta, Alpha-Beta) to assess the reproducibility of mask learning given random initialization of mask scores. Each point represents the validation perplexity of a subnetwork stratified by category of inputs.

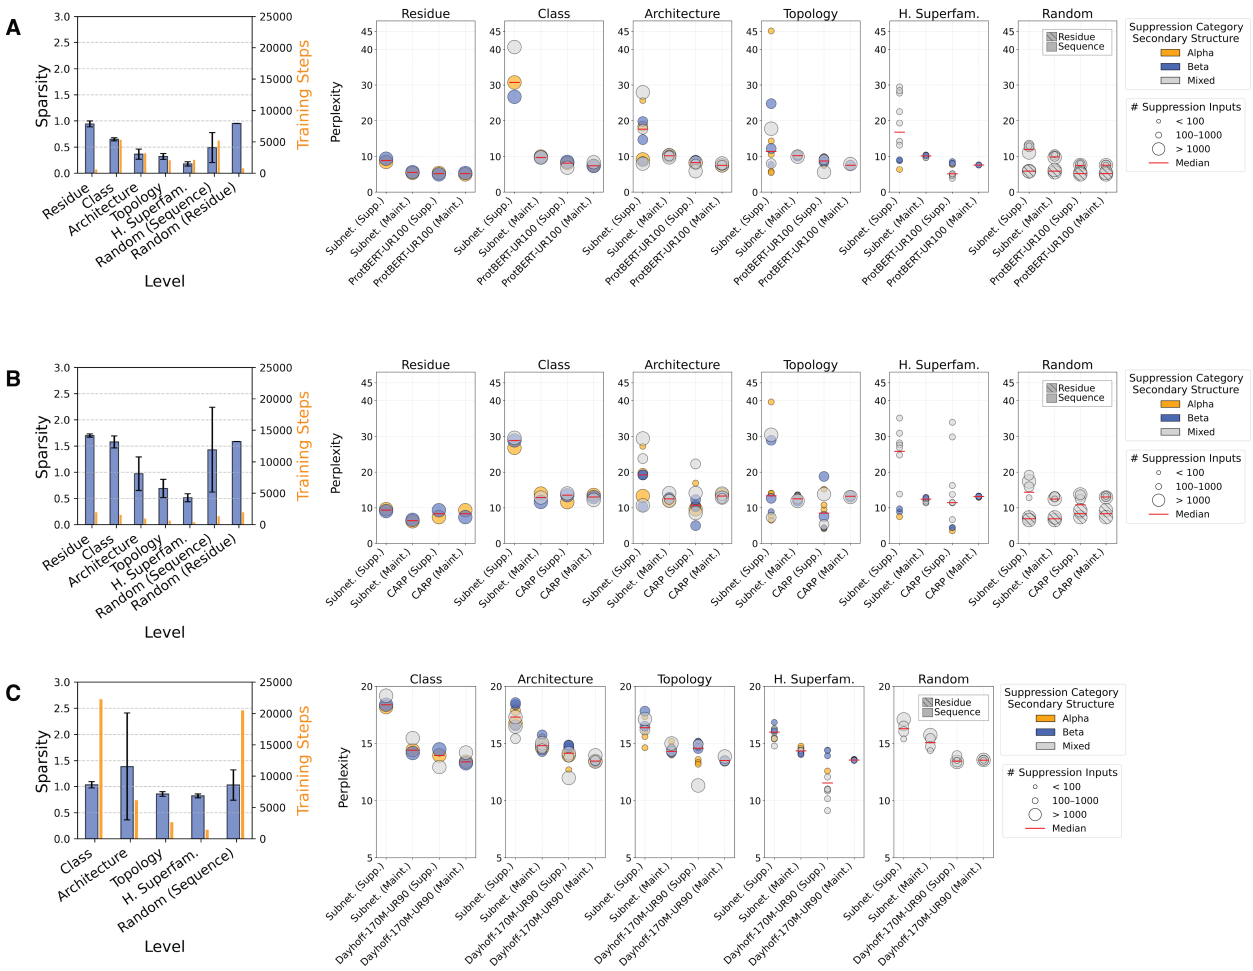


**Figure S3. Subnetwork language modeling performance on three additional PLMs of varying size and architectures.** To investigate whether subnetworks exist in pretrained PLMs other than ESM-2, we applied our approach to three additional models of varying size and architectures trained on UniRef data: ProtBERT-UR100, a transformer masked language model with a BERT architecture [2]; CARP-640M, a convolutional neural network masked language model (transformer layers are replaced by ByteNet dilated CNN blocks) [14]; and Dayhoff-170M-UR90, an efficient hybrid state-space-model transformer trained with and autoregressive objective [31]. **(A) ProtBERT-UR100 (420M).** Left: Learned percent sparsity by category of learned subnetworks in ProtBERT-UR100. Right: Masked language modeling performance on suppression and maintenance categories of sequences for each subnetwork. **(B) CARP-640M.** Left: Learned percent sparsity by category of learned subnetworks in CARP-640M. Right: Masked language modeling performance on suppression and maintenance categories of sequences for each subnetwork. **(C) Dayhoff-170M-UR90.** Left: Learned percent sparsity by category of learned subnetworks in Dayhoff-170M-UR90. Right: Autoregressive language modeling performance on suppression and maintenance categories of sequences for each subnetwork. Residue subnetworks were omitted from Dayhoff-170M-UR90 results because causal perplexity cannot be computed selectively over individual residues.

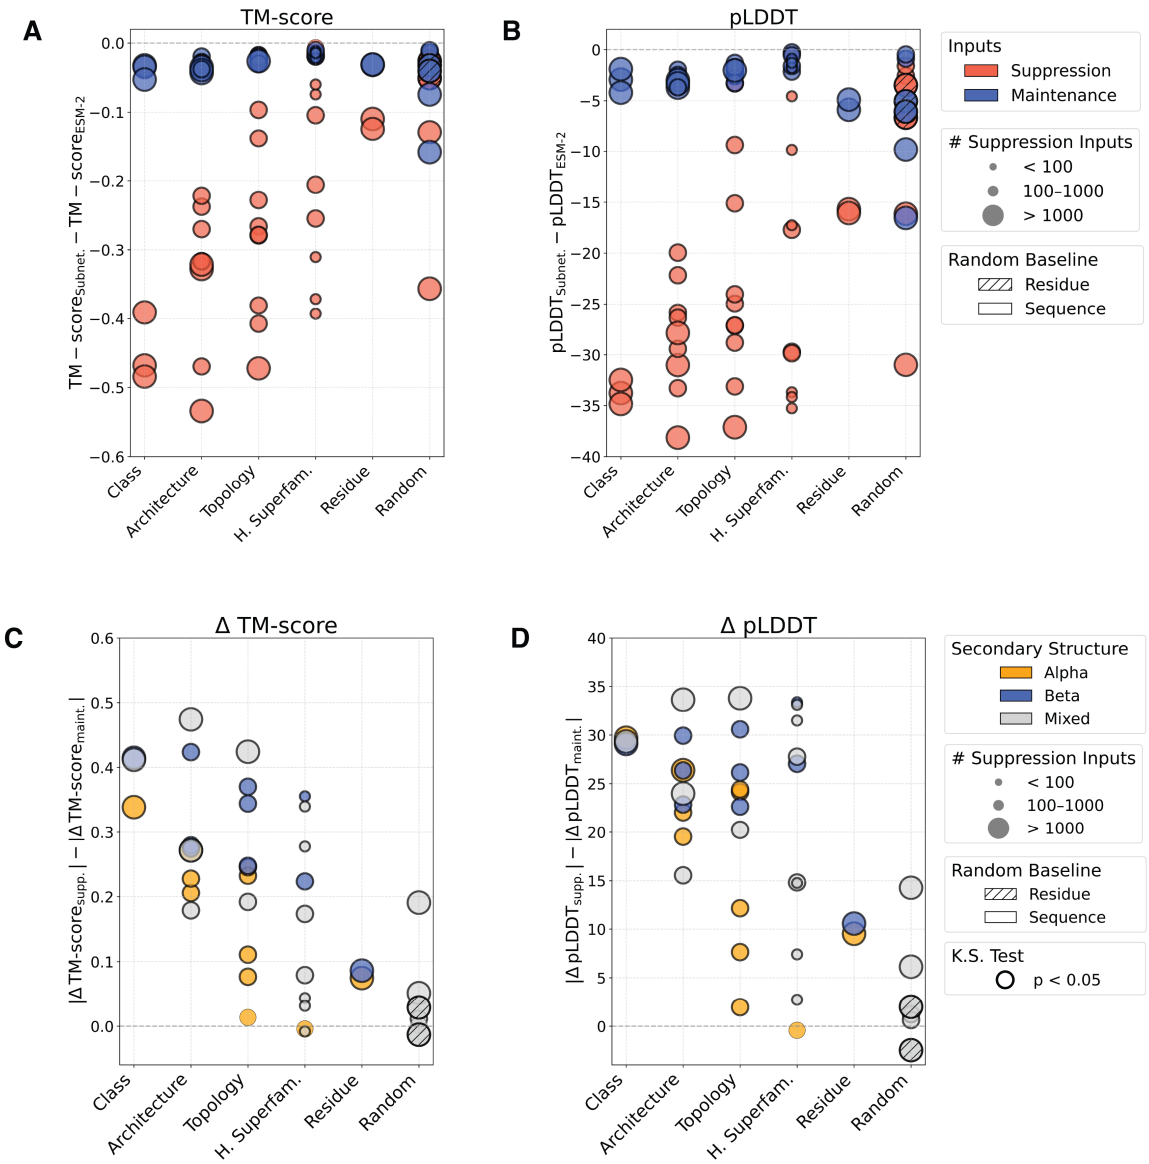


**Figure S4. ESM-2 650M subnetwork-predicted TM-score and pLDDT with the ESM-2 folding trunk. (A–B)** Structural prediction differences (y-axis) between subnetworks and the ESM-2 baseline shown for **(A)** TM-score and **(B)** pLDDT across structural levels (x-axis). Each point represents change in metrics for suppression inputs (red) or maintenance inputs (blue) for a subnetwork relative to ESM-2. Marker size indicates the number of suppression inputs for each subnetwork. Bold outlines of markers indicate statistically significant paired t-test $p$-values ($p < 0.05$). **(C–D)** Difference in absolute structure prediction metric changes from the ESM-2 baseline (y-axis), stratified by structural level, for **(C)** TM-score and **(D)** pLDDT. Each point corresponds to an individual subnetwork and shows the difference between suppression and maintenance $\Delta$-values. Marker size reflects the number of suppressed inputs in the subnetwork, and color indicates secondary structure type. Bold outlines of markers indicate statistical significance by Kolmogorov–Smirnov (KS) test ($p < 0.05$).

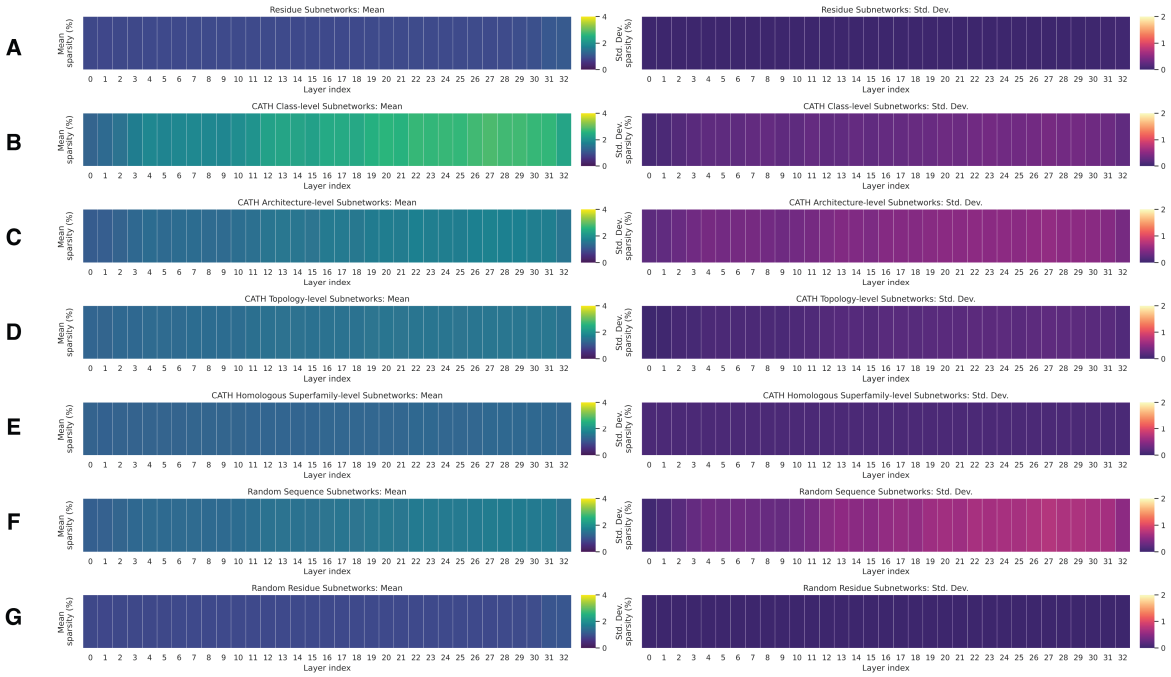


**Figure S5. Mask interpretation of ESM-2 650M.** Mean and standard deviation percent of parameters pruned by layer for subnetworks grouped at the levels of **(A)** residue, **(B)** CATH class, **(C)** CATH architecture, **(D)** CATH topology, **(E)** CATH homologous superfamily, **(F)** random sequence suppression, and **(G)** random residue suppression.

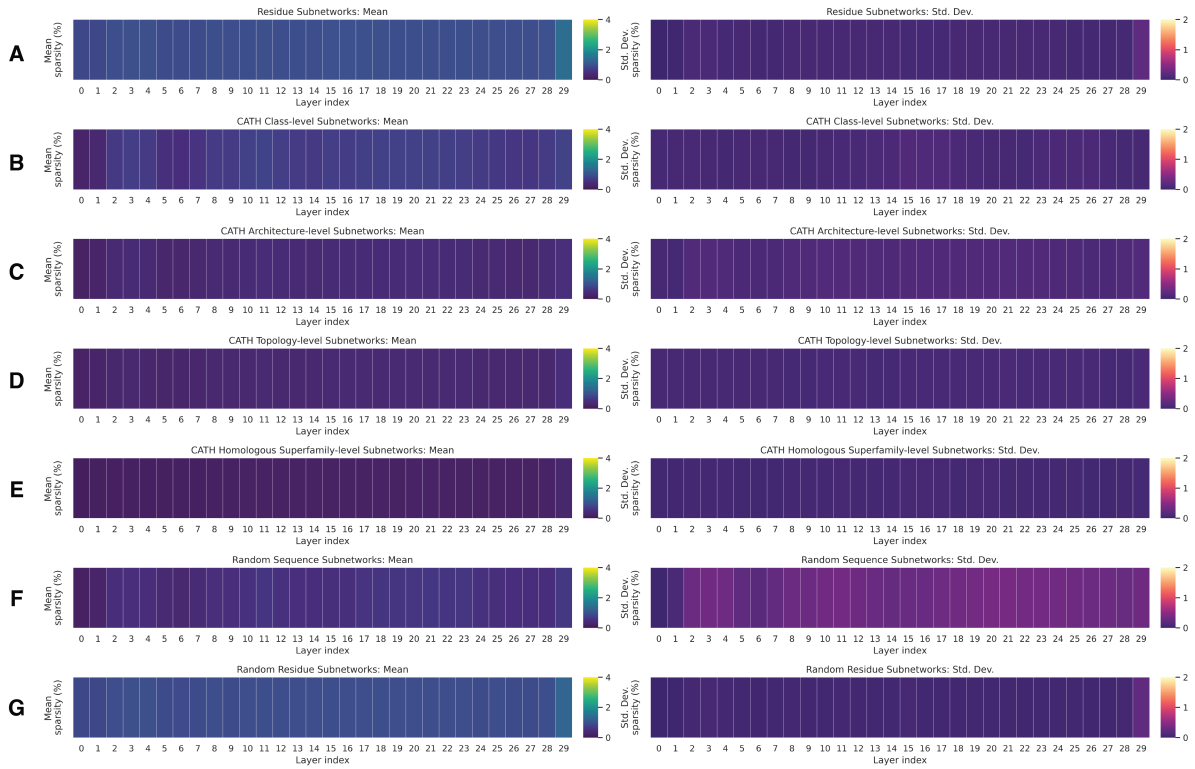


**Figure S6. Mask interpretation of ProtBERT-UR100.** Mean and standard deviation percent of parameters pruned by layer for subnetworks grouped at the levels of **(A)** residue, **(B)** CATH class, **(C)** CATH architecture, **(D)** CATH topology, **(E)** CATH homologous superfamily, **(F)** random sequence suppression, and **(G)** random residue suppression.

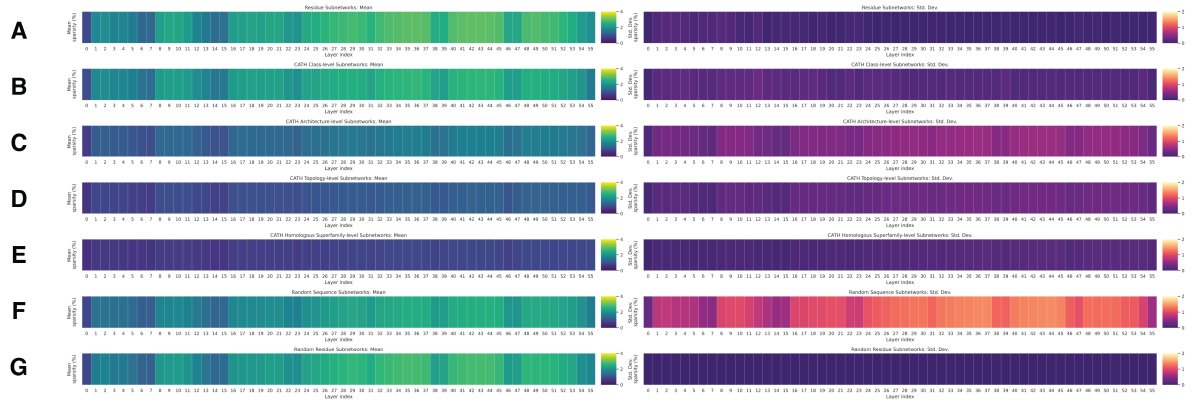


**Figure S7. Mask interpretation of CARP-640M.** Mean and standard deviation percent of parameters pruned by layer for subnetworks grouped at the levels of **(A)** residue, **(B)** CATH class, **(C)** CATH architecture, **(D)** CATH topology, **(E)** CATH homologous superfamily, **(F)** random sequence suppression, and **(G)** random residue suppression.

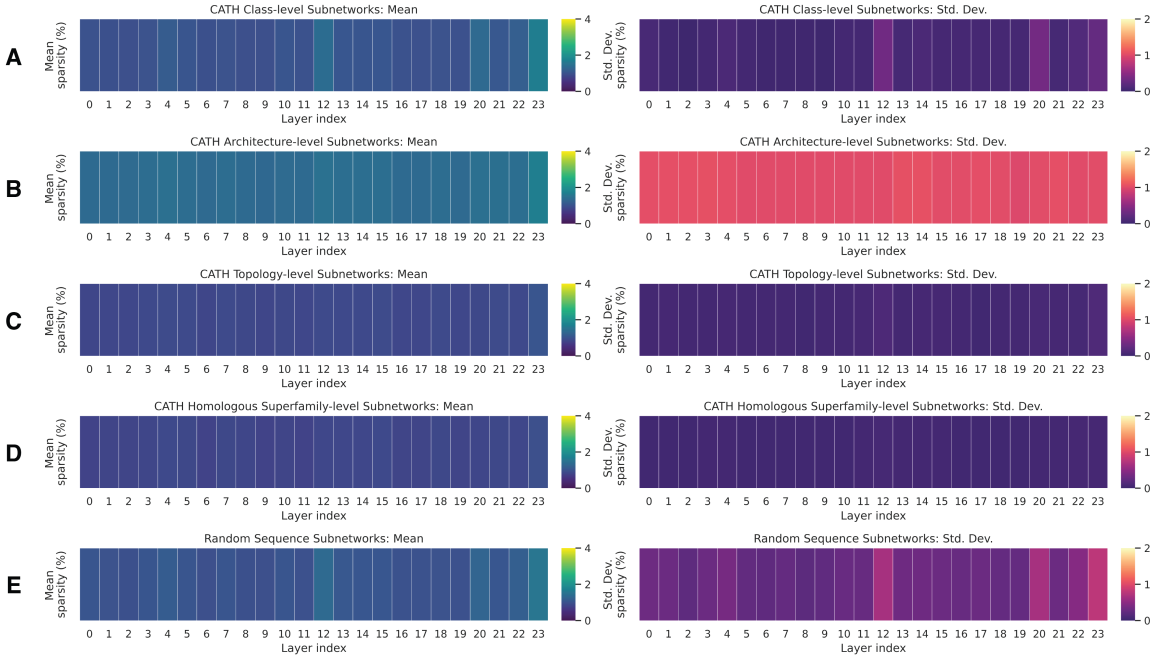


**Figure S8. Mask interpretation of Dayhoff-170M-UR90.** Mean and standard deviation percent of parameters pruned by layer for subnetworks grouped at the levels of **(A)** CATH class, **(B)** CATH architecture, **(C)** CATH topology, **(D)** CATH homologous superfamily, and **(E)** random sequence suppression.

# Supplementary Tables

**Table S1. Training and hyperparamter configurations and masked modules for subnetwork learning across four pretrained protein language models.** Each PLM subnetwork differs in which modules are masked, according to the model architecture, following evidence that knowledge is localized in representational modules [22]. Masking is thus applied to self-attention or convolutional projections while leaving embeddings, bias, and normalization layers intact. Subnetworks trained within a single PLM reuse the same hyperparameter configuration shown below. ESM-2 650M, ProtBERT-UR100, and CARP-640M are masked-language-model PLMs trained with the same MLM objective, whereas Dayhoff-170M-UR90 is an autoregressive model trained with a next-token prediction objective.

| Model | Masked Modules | Layers | $s_{\text{init}}$ | $\lambda_{\text{maint}}$ | $\lambda_{\text{supp}}$ | $\lambda_{\text{MLM}}$ | $\tau_{\text{init}}$ | $\tau_{\text{final}}$ | $T$ |
|---|---|---|---|---|---|---|---|---|---|
| ESM-2 650M | Self-attention layers | 34 | 0.96 | 7 | 10 | 1 | 3 | 0.01 | 0.40 |
| ProtBERT-UR100 | Self-attention layers | 30 | 0.998 | 9 | 10 | 1 | 1 | 0.1 | 0.43 |
| CARP-640M | Convolutional kernels | 56 | 0.995 | 9 | 10 | 1 | 1 | 0.1 | 0.40 |
| Dayhoff-170M-UR90 | Self-attention + Mamba projections | 24 | 0.99 | 4 | 4 | 4 | 1 | 0.1 | 0.40 |

**Table S2. Mean and standard deviation of ESM-2 CATH Class-level subnetwork masked language modeling performance across multiple seeds.** Three subnetworks were independently trained for each CATH Class suppression target (Mainly Alpha, Mainly Beta, Alpha-Beta) to assess the reproducibility of mask learning given random initialization of mask scores. Reported are the mean ± standard deviation of masked language modeling perplexity for subnetworks and baseline ESM-2 perplexity stratified by categories of inputs.

| **Category** | **Target** | **Sparsity** | **Training Step** | **Subnet. Supp.** | **Subnet. Maint.** | **ESM Supp.** | **ESM Maint.** | **t-Test ($p$) Supp.** | **t-Test ($p$) Maint.** |
|---|---|---|---|---|---|---|---|---|---|
| Class | 1 | 2.47 | 7,600 | $24.9 \pm 72.5$ | $8.6 \pm 17.4$ | $8.6 \pm 15.2$ | $8.8 \pm 14.1$ | $< 1e{-}16$ | 3.7e-05 |
| | 1 | 2.45 | 7,640 | $26.8 \pm 82.5$ | $8.6 \pm 16.1$ | $8.6 \pm 15.2$ | $8.8 \pm 14.1$ | $< 1e{-}16$ | 2.4e-12 |
| | 1 | 2.48 | 7,560 | $28.3 \pm 82.9$ | $8.6 \pm 16.2$ | $8.6 \pm 15.2$ | $8.8 \pm 14.1$ | $< 1e{-}16$ | 2.2e-13 |
| | 2 | 2.4 | 7,040 | $28.0 \pm 70.2$ | $8.7 \pm 16.0$ | $8.9 \pm 15.7$ | $8.7 \pm 14.0$ | $< 1e{-}16$ | 6.5e-03 |
| | 2 | 2.36 | 7,160 | $25.8 \pm 62.8$ | $8.7 \pm 16.3$ | $8.9 \pm 15.7$ | $8.7 \pm 14.0$ | $< 1e{-}16$ | 1.5e-02 |
| | 2 | 2.4 | 7,120 | $24.6 \pm 59.8$ | $8.7 \pm 15.4$ | $8.9 \pm 15.7$ | $8.7 \pm 14.0$ | $< 1e{-}16$ | 3.0e-01 |
| | 3 | 2.56 | 7,880 | $48.3 \pm 106.4$ | $8.3 \pm 22.4$ | $8.9 \pm 13.7$ | $8.6 \pm 15.2$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3 | 2.58 | 7,900 | $43.2 \pm 89.9$ | $8.1 \pm 17.0$ | $8.9 \pm 13.7$ | $8.6 \pm 15.2$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3 | 2.59 | 8,000 | $38.8 \pm 84.3$ | $8.1 \pm 18.0$ | $8.9 \pm 13.7$ | $8.6 \pm 15.2$ | $< 1e{-}16$ | $< 1e{-}16$ |

**Table S3. Overview of subnetwork definitions.** Each subnetwork is trained to selectively suppress either a specific type of residue belonging to a secondary structure or a set of sequences that belong to the same CATH category classification. In our study, we consider only the top 10 most frequent labels in each CATH category, and report the number of sequences and predominant type of secondary structure content in each category.

| **Category** | **Target** | **# Seqs.** | **Type** |
|---|---|---|---|
| Residue | Helix | 36.5% of residues | Alpha |
| | Sheet | 21.9% of residues | Beta |
| Class | 1 | 2169 | Alpha |

| Category | Target | # Seqs. | Type |
|---|---|---|---|
| | 2 | 2038 | Beta |
| | 3 | 3129 | Mixed |
| Architecture | 1.10 | 1247 | Alpha |
| | 1.25 | 182 | Alpha |
| | 1.20 | 693 | Alpha |
| | 2.30 | 251 | Beta |
| | 2.40 | 447 | Beta |
| | 2.60 | 896 | Beta |
| | 3.30 | 1396 | Mixed |
| | 3.40 | 1797 | Mixed |
| | 40 | 469 | Mixed |
| Topology | 1.10.10 | 192 | Alpha |
| | 1.10.287 | 111 | Alpha |
| | 1.20.58 | 135 | Alpha |
| | 1.20.120 | 118 | Alpha |
| | 1.25.40 | 142 | Alpha |
| | 2.40.50 | 136 | Beta |
| | 2.60.40 | 512 | Beta |
| | 2.60.120 | 276 | Beta |
| | 3.30.70 | 282 | Mixed |
| | 3.40.50 | 1110 | Mixed |
| H. Superfam. | 3.40.50.300 | 168 | Mixed |
| | 2.60.40.10 | 160 | Beta |
| | 3.40.50.720 | 117 | Mixed |
| | 1.10.10.10 | 102 | Alpha |
| | 3.20.20.80 | 88 | Mixed |
| | 3.40.190.10 | 81 | Mixed |
| | 3.40.50.150 | 77 | Mixed |
| | 3.40.50.1820 | 75 | Mixed |
| | 2.40.50.140 | 68 | Beta |
| | 3.40.30.10 | 67 | Mixed |
| Random | Sequences | 100 | Mixed |
| | | 200 | Mixed |
| | | 1000 | Mixed |
| | | 2000 | Mixed |
| Random | Residue | Random | Mixed |

**Table S4. ESM-2 650M subnetwork and baseline language modeling performance.** Below we report the mean and standard deviation of the subnetwork and ESM-2 baseline perplexities (illustrated in Fig. 2C) stratified by categories of inputs. The $t$-test is performed on paired inputs of the subnetwork performance and ESM-2 baseline performance and $p$-values are reported below. The random residue control subnetwork is trained to suppress random residues, but we independently evaluate and report the MLM performance of this subnetwork on alpha helices and beta sheets. All per-sequence metrics are available in CSVs in the code repository.

| Category | Target | Sparsity | Training Step | Subnet. Supp. | Subnet. Maint. | ESM-2 Supp. | ESM-2 Maint. | t-Test ($p$) Supp. | t-Test ($p$) Maint. |
|---|---|---|---|---|---|---|---|---|---|
| Residue | Helix | 0.90 | 784 | $7.9 \pm 5.0$ | $5.0 \pm 5.2$ | $5.6 \pm 4.0$ | $4.9 \pm 5.1$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | Sheet | 0.89 | 784 | $6.7 \pm 6.1$ | $5.6 \pm 4.1$ | $4.9 \pm 5.1$ | $5.6 \pm 4.0$ | $< 1e{-}16$ | $< 1e{-}16$ |
| Class | 1 | 2.13 | 7,680 | 27.2±80.3 | $8.0 \pm 15.4$ | $8.6 \pm 15.2$ | $8.8 \pm 14.1$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 2 | 2.06 | 7,200 | 28.7±74.1 | $8.2 \pm 15.5$ | $8.9 \pm 15.7$ | $8.7 \pm 14.0$ | $< 1e{-}16$ | $< 1e{-}16$ |

| Category | Target | Sparsity | Training Step | Subnet. Supp. | Subnet. Maint. | ESM-2 Supp. | ESM-2 Maint. | t-Test ($p$) Supp. | t-Test ($p$) Maint. |
|---|---|---|---|---|---|---|---|---|---|
| | 3 | 2.53 | 7,820 | 41.5±91.4 | $8.2 \pm 18.5$ | $8.9 \pm 13.7$ | $8.6 \pm 15.2$ | $< 1e{-}16$ | $< 1e{-}16$ |
| Architecture | 1.10 | 2.02 | 4,440 | 11.0±21.7 | $9.3 \pm 16.4$ | $7.6 \pm 15.2$ | $8.9 \pm 14.3$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 1.20 | 1.25 | 3,660 | 14.1±28.0 | $9.3 \pm 17.6$ | $9.0 \pm 12.4$ | $8.7 \pm 14.5$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 1.25 | 1.31 | 640 | 34.9±70.6 | $9.1 \pm 16.8$ | 11.9±19.6 | $8.7 \pm 14.2$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 2.30 | 1.05 | 1,680 | $9.9 \pm 12.7$ | $9.5 \pm 17.0$ | $5.1 \pm 5.1$ | $8.8 \pm 14.5$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 2.40 | 1.10 | 2,280 | 13.0±21.4 | $9.3 \pm 17.1$ | $7.2 \pm 10.5$ | $8.8 \pm 14.6$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 2.60 | 1.71 | 3,000 | 20.0±37.6 | $9.2 \pm 16.5$ | $8.9 \pm 14.1$ | $8.7 \pm 14.4$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3.30 | 1.68 | 4,880 | 10.1±14.6 | $9.3 \pm 17.6$ | $7.3 \pm 10.6$ | $9.0 \pm 15.0$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3.40 | 2.18 | 6,200 | 34.9±67.5 | $8.8 \pm 18.2$ | $9.2 \pm 14.0$ | $8.6 \pm 14.5$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3.90 | 1.41 | 1,640 | 14.9±28.0 | $9.0 \pm 16.3$ | 10.8±19.4 | $8.6 \pm 14.1$ | $< 1e{-}16$ | $< 1e{-}16$ |
| Topology | 1.10.10 | 1.51 | 1,280 | $6.9 \pm 9.6$ | $9.7 \pm 17.5$ | $4.6 \pm 6.7$ | $8.8 \pm 14.5$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 1.10.287 | 1.37 | 720 | $5.3 \pm 5.7$ | $9.4 \pm 17.1$ | $5.2 \pm 6.1$ | $8.8 \pm 14.5$ | $4.5e-01$ | $< 1e{-}16$ |
| | 1.20.120 | 1.38 | 800 | 11.9±14.5 | $9.5 \pm 16.9$ | $8.3 \pm 8.5$ | $8.7 \pm 14.5$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 1.20.58 | 1.41 | 960 | 13.4±19.8 | $9.6 \pm 17.3$ | $8.4 \pm 10.8$ | $8.7 \pm 14.4$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 1.25.40 | 1.34 | 780 | 40.8±94.2 | $9.2 \pm 16.7$ | 11.0±19.2 | $8.7 \pm 14.3$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 2.40.50 | 1.36 | 720 | $8.7 \pm 9.7$ | $9.3 \pm 16.4$ | $5.2 \pm 5.8$ | $8.8 \pm 14.5$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 2.60.120 | 1.37 | 1,000 | 32.7±49.6 | $9.1 \pm 17.9$ | 11.1±15.0 | $8.7 \pm 14.4$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 2.60.40 | 1.52 | 1,800 | 11.9±15.3 | $9.4 \pm 17.5$ | $7.1 \pm 9.7$ | $8.8 \pm 14.6$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3.30.70 | 1.52 | 1,500 | $8.3 \pm 9.4$ | $9.6 \pm 18.1$ | $5.7 \pm 6.3$ | $8.8 \pm 14.6$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3.40.50 | 1.80 | 3,840 | 25.6±40.2 | $9.1 \pm 17.7$ | $8.8 \pm 11.1$ | $8.7 \pm 14.8$ | $< 1e{-}16$ | $< 1e{-}16$ |
| H. Superfam. | 1.10.10.10 | 1.26 | 300 | $4.0 \pm 3.4$ | $9.0 \pm 16.5$ | $4.1 \pm 3.9$ | $8.8 \pm 14.5$ | $4.1e-02$ | $< 1e{-}16$ |
| | 2.40.50.140 | 1.38 | 480 | 10.2±11.0 | $9.5 \pm 17.4$ | $4.6 \pm 3.6$ | $8.8 \pm 14.4$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 2.60.40.10 | 1.04 | 840 | $8.8 \pm 9.1$ | $9.2 \pm 16.4$ | $4.9 \pm 4.3$ | $8.8 \pm 14.5$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3.20.20.80 | 1.26 | 360 | 46.4±55.6 | $9.1 \pm 16.8$ | 12.9±14.5 | $8.7 \pm 14.4$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3.40.190.10 | 1.29 | 420 | 22.7±30.1 | $9.2 \pm 16.5$ | $8.9 \pm 8.6$ | $8.7 \pm 14.4$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3.40.30.10 | 1.32 | 480 | 15.1±13.5 | $9.3 \pm 17.5$ | $5.1 \pm 4.4$ | $8.8 \pm 14.4$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3.40.50.150 | 1.31 | 480 | 17.5±24.0 | $9.2 \pm 17.5$ | $9.5 \pm 12.1$ | $8.7 \pm 14.4$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3.40.50.1820 | 1.27 | 360 | 20.2±21.8 | $9.1 \pm 16.5$ | 12.2±12.6 | $8.7 \pm 14.4$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3.40.50.300 | 1.30 | 560 | 21.6±28.1 | $9.1 \pm 16.3$ | $9.2 \pm 10.2$ | $8.7 \pm 14.5$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 3.40.50.720 | 1.32 | 660 | 27.6±49.2 | $9.2 \pm 16.7$ | $7.4 \pm 7.8$ | $8.8 \pm 14.5$ | $< 1e{-}16$ | $< 1e{-}16$ |
| Random | Helix | 0.87 | 196 | $6.2 \pm 4.3$ | $5.2 \pm 5.2$ | $5.6 \pm 4.0$ | $4.9 \pm 5.1$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | Sheet | 0.87 | 196 | $5.2 \pm 5.2$ | $6.2 \pm 4.3$ | $4.9 \pm 5.1$ | $5.6 \pm 4.0$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 100 seqs. | 1.17 | 420 | $9.2 \pm 13.2$ | $8.9 \pm 15.7$ | $8.4 \pm 12.0$ | $8.7 \pm 14.4$ | $1.8e-09$ | $< 1e{-}16$ |
| | 200 seqs. | 1.21 | 780 | $9.2 \pm 13.2$ | $9.0 \pm 15.7$ | $8.3 \pm 11.8$ | $8.7 \pm 14.4$ | $< 1e{-}16$ | $< 1e{-}16$ |
| | 1000 seqs. | 1.37 | 2,520 | $8.5 \pm 13.4$ | $8.9 \pm 16.2$ | $7.8 \pm 11.5$ | $8.8 \pm 14.6$ | $< 1e{-}16$ | $2.1e-04$ |
| | 2000 seqs. | 2.17 | 7,500 | 11.2±21.3 | $8.2 \pm 15.6$ | $8.6 \pm 14.0$ | $8.8 \pm 14.5$ | $< 1e{-}16$ | $< 1e{-}16$ |

**Table S5. ProteinMPNN perplexities on CATH PDBs (Mainly Alpha, Mainly Beta, and Alpha-Beta classes).** Reported values correspond to mean, standard deviation, minimum, and maximum perplexities across all CATH domains within each structural CATH Class.

| Category | Count (n) | Mean | Std. Dev. | Min | Max |
|---|---|---|---|---|---|
| Mainly Alpha | 10506 | 4.49 | 3.02 | 1.51 | 44.00 |
| Mainly Beta | 8493 | 3.71 | 2.15 | 1.56 | 35.19 |
| Alpha-Beta | 19092 | 3.50 | 1.97 | 1.65 | 57.61 |

**Table S6. ESM-2 650M subnetwork and baseline structure prediction performance using the ESMFold (650M) folding trunk on the validation datasets**. For TM-score, RMSD, and pLDDT, we report the mean ± standard deviation of the subnetwork performance across all sequences within each category of suppression and maintenance inputs. ESM-2 (650M) performance on the same categories is reported as the PLM baseline. To quantify the differences in subnetwork and PLM performance, we perform a paired t-test on all (i) suppression inputs and (ii) maintenance inputs, computing the difference in $\Delta\text{metric} = \text{metric}_{\text{Subnet.}} - \text{metric}_{\text{ESM-2.}}$. We then perform a Kolmogorov–Smirnov (KS) test on these differences to assess whether the distribution of $|\Delta_{\text{metric,supp}}|$ is significantly greater than that of $|\Delta_{\text{metric,maint}}|$. Our evaluation scheme is illustrated in Fig. 3A. We report $p$-values for both statistical tests for each subnetwork; significant $p$-values are in bold. For residue-level suppression, we evaluate structure prediction on CATH class categories of mainly alpha and mainly beta sequences as a proxy for evaluating alpha helix and beta strand performance. We report performance on the only residue-specific metric of pLDDT in Fig. 3F-G. The random residue suppression subnetwork, i.e. residue-control, is one subnetwork but we evaluated it separately on alpha helices and beta sheets. All per-sequence structure prediction metrics are provided via CSVs in our code repository.

| Category | Target | Metric | Subnet. Supp. | Subnet. Maint. | ESM-2 Supp. | ESM-2 Maint. | t-Test ($p$) Supp. | t-Test ($p$) Maint. | K.S.-test ($p$) |
|---|---|---|---|---|---|---|---|---|---|
| Residue | Helix | RMSD | 3.62 ± 1.19 | 3.05 ± 1.09 | 2.93 ± 1.07 | 2.83 ± 1.09 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.55 ± 0.20 | 0.68 ± 0.19 | 0.66 ± 0.20 | 0.71 ± 0.19 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 50.81±13.89 | 59.49±12.79 | 66.51±12.82 | 65.38±12.69 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | Sheet | RMSD | 3.74 ± 1.43 | 3.05 ± 1.12 | 2.89 ± 1.19 | 2.85 ± 1.11 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.57 ± 0.21 | 0.66 ± 0.20 | 0.70 ± 0.19 | 0.69 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 47.79±13.55 | 60.98±12.72 | 63.86±13.49 | 65.87±12.65 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| Class | 1 | RMSD | 5.22 ± 1.42 | 3.05 ± 1.10 | 2.86 ± 1.12 | 2.82 ± 1.10 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.27 ± 0.12 | 0.68 ± 0.19 | 0.66 ± 0.21 | 0.71 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 33.31 ± 8.37 | 62.30±13.07 | 67.09±13.15 | 65.24±12.92 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 2 | RMSD | 6.02 ± 1.28 | 3.08 ± 1.10 | 2.93 ± 1.15 | 2.84 ± 1.09 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.23 ± 0.11 | 0.66 ± 0.20 | 0.69 ± 0.19 | 0.69 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 30.53 ± 7.66 | 64.25±12.91 | 63.02±13.59 | 66.16±12.66 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 3 | RMSD | 6.13 ± 1.39 | 3.26 ± 1.15 | 2.81 ± 1.06 | 2.93 ± 1.13 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.23 ± 0.12 | 0.61 ± 0.20 | 0.72 ± 0.19 | 0.67 ± 0.21 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 30.80 ± 7.97 | 60.54±13.99 | 65.64±12.55 | 64.75±13.30 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| Arch. | 1.10 | RMSD | 4.80 ± 1.20 | 3.12 ± 1.16 | 2.85 ± 1.13 | 2.83 ± 1.11 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.32 ± 0.13 | 0.66 ± 0.21 | 0.65 ± 0.20 | 0.71 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 35.91 ± 8.20 | 61.82±14.58 | 66.94±12.45 | 65.53±12.96 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 1.20 | RMSD | 4.59 ± 1.27 | 3.05 ± 1.16 | 3.06 ± 1.08 | 2.84 ± 1.10 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.36 ± 0.17 | 0.66 ± 0.21 | 0.63 ± 0.20 | 0.70 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 39.83±13.38 | 62.42±14.80 | 65.69±13.68 | 65.36±13.01 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 1.25 | RMSD | 4.29 ± 1.40 | 3.00 ± 1.10 | 2.67 ± 1.04 | 2.86 ± 1.10 | **4.6e-12** | **< 1e−16** | **< 1e−16** |

| Category | Target | Metric | Subnet. Supp. | Subnet. Maint. | ESM-2 Supp. | ESM-2 Maint. | t-Test ($p$) Supp. | t-Test ($p$) Maint. | K.S.-test ($p$) |
|---|---|---|---|---|---|---|---|---|---|
| | | TM-score | 0.54 ± 0.22 | 0.67 ± 0.20 | 0.78 ± 0.15 | 0.69 ± 0.20 | **7.5e-12** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 49.23±15.56 | 63.32±13.03 | 71.39±11.65 | 65.30±12.65 | **5.7e-16** | **< 1e−16** | **< 1e−16** |
| | 2.30 | RMSD | 4.66 ± 1.21 | 3.06 ± 1.17 | 2.69 ± 0.95 | 2.86 ± 1.12 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.37 ± 0.18 | 0.67 ± 0.20 | 0.68 ± 0.18 | 0.70 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 37.72±12.77 | 62.99±14.17 | 67.14±13.07 | 65.61±13.01 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 2.40 | RMSD | 5.12 ± 1.45 | 3.00 ± 1.15 | 2.99 ± 1.07 | 2.81 ± 1.09 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.34 ± 0.20 | 0.67 ± 0.20 | 0.66 ± 0.19 | 0.70 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 35.83±12.39 | 63.38±13.86 | 62.20±14.02 | 66.23±12.59 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 2.60 | RMSD | 6.03 ± 1.09 | 3.07 ± 1.15 | 2.84 ± 1.16 | 2.86 ± 1.11 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.24 ± 0.12 | 0.66 ± 0.21 | 0.71 ± 0.18 | 0.69 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 29.96 ± 6.60 | 63.17±13.73 | 63.26±13.72 | 65.66±12.94 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 3.30 | RMSD | 4.93 ± 1.24 | 3.05 ± 1.14 | 2.87 ± 1.14 | 2.80 ± 1.09 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.34 ± 0.16 | 0.67 ± 0.20 | 0.66 ± 0.20 | 0.71 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 36.39±11.38 | 63.25±13.87 | 64.26±14.23 | 66.24±12.61 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 3.40 | RMSD | 6.42 ± 1.24 | 3.18 ± 1.14 | 2.60 ± 0.98 | 2.92 ± 1.13 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.24 ± 0.11 | 0.63 ± 0.20 | 0.77 ± 0.16 | 0.67 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 30.33 ± 6.87 | 61.31±13.65 | 68.47±10.68 | 64.63±13.30 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 3.90 | RMSD | 4.85 ± 1.49 | 3.10 ± 1.22 | 3.32 ± 1.38 | 2.84 ± 1.11 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.42 ± 0.21 | 0.66 ± 0.21 | 0.65 ± 0.23 | 0.70 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 41.17±12.83 | 61.81±15.13 | 61.12±13.68 | 65.50±13.06 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| Topo. | 1.10.10 | RMSD | 3.90 ± 1.10 | 3.05 ± 1.15 | 2.35 ± 0.90 | 2.87 ± 1.13 | **9.0e-14** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.41 ± 0.16 | 0.67 ± 0.20 | 0.68 ± 0.18 | 0.70 ± 0.20 | **1.3e-14** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 44.81±14.34 | 62.72±14.01 | 71.96±10.84 | 65.30±13.04 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 1.10.287 | RMSD | 2.80 ± 0.93 | 2.98 ± 1.08 | 2.77 ± 0.86 | 2.85 ± 1.08 | 7.6e-01 | **< 1e−16** | **3.2e-02** |
| | | TM-score | 0.56 ± 0.17 | 0.68 ± 0.19 | 0.59 ± 0.18 | 0.70 ± 0.20 | **6.2e-03** | **< 1e−16** | 6.1e-02 |
| | | pLDDT | 67.32±13.60 | 63.70±12.75 | 70.62±11.32 | 65.74±12.57 | **5.9e-03** | **< 1e−16** | **4.6e-02** |
| | 1.20.120 | RMSD | 4.10 ± 1.35 | 3.00 ± 1.11 | 3.28 ± 1.13 | 2.83 ± 1.09 | **1.4e-04** | **< 1e−16** | **2.8e-05** |
| | | TM-score | 0.47 ± 0.19 | 0.67 ± 0.20 | 0.61 ± 0.20 | 0.70 ± 0.20 | **7.8e-07** | **< 1e−16** | **9.6e-11** |
| | | pLDDT | 48.32±14.04 | 63.08±13.69 | 63.44±12.66 | 65.64±12.96 | **1.6e-07** | **< 1e−16** | **1.2e-10** |
| | 1.20.58 | RMSD | 3.40 ± 0.98 | 3.00 ± 1.10 | 2.93 ± 1.13 | 2.85 ± 1.10 | **3.3e-04** | **< 1e−16** | **4.3e-04** |
| | | TM-score | 0.50 ± 0.16 | 0.68 ± 0.20 | 0.60 ± 0.21 | 0.70 ± 0.20 | **1.4e-05** | **< 1e−16** | **5.6e-08** |
| | | pLDDT | 55.58±12.59 | 63.37±13.01 | 64.98±14.31 | 65.58±12.65 | **7.5e-06** | **< 1e−16** | **5.6e-07** |
| | 1.25.40 | RMSD | 4.21 ± 1.60 | 2.98 ± 1.13 | 2.65 ± 1.14 | 2.85 ± 1.11 | **1.3e-08** | **< 1e−16** | **1.0e-12** |
| | | TM-score | 0.49 ± 0.25 | 0.68 ± 0.20 | 0.77 ± 0.15 | 0.70 ± 0.20 | **5.2e-10** | **< 1e−16** | **2.1e-16** |
| | | pLDDT | 47.61±17.00 | 63.50±13.65 | 74.72±10.08 | 65.48±13.00 | **1.7e-13** | **< 1e−16** | **< 1e−16** |
| | 2.40.50 | RMSD | 4.66 ± 1.19 | 2.98 ± 1.12 | 2.86 ± 1.11 | 2.86 ± 1.10 | **1.1e-11** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.35 ± 0.17 | 0.68 ± 0.20 | 0.63 ± 0.21 | 0.69 ± 0.20 | **6.2e-11** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 37.35±10.71 | 63.53±13.34 | 62.32±14.64 | 65.39±12.84 | **2.3e-12** | **< 1e−16** | **< 1e−16** |
| | 2.60.120 | RMSD | 5.75 ± 1.65 | 3.00 ± 1.15 | 2.91 ± 1.15 | 2.86 ± 1.12 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.36 ± 0.22 | 0.67 ± 0.21 | 0.74 ± 0.17 | 0.69 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 33.70±10.54 | 63.74±13.81 | 62.51±12.81 | 65.47±13.04 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 2.60.40 | RMSD | 5.29 ± 1.24 | 3.02 ± 1.14 | 2.61 ± 0.91 | 2.85 ± 1.10 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.32 ± 0.17 | 0.67 ± 0.20 | 0.73 ± 0.15 | 0.69 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 32.83 ± 8.78 | 64.14±13.58 | 65.96±12.34 | 65.44±12.82 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 3.30.70 | RMSD | 4.12 ± 1.11 | 3.05 ± 1.14 | 2.72 ± 1.13 | 2.84 ± 1.11 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.44 ± 0.18 | 0.67 ± 0.20 | 0.67 ± 0.20 | 0.70 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 41.78±13.24 | 62.50±13.90 | 65.82±14.80 | 65.76±12.80 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 3.40.50 | RMSD | 5.99 ± 1.16 | 3.09 ± 1.13 | 2.53 ± 0.88 | 2.91 ± 1.14 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.32 ± 0.14 | 0.66 ± 0.20 | 0.80 ± 0.14 | 0.68 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 32.65 ± 6.56 | 62.69±13.36 | 69.77 ± 9.16 | 64.73±13.38 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| H. Supfam. | 1.10.10.10 | RMSD | 2.30 ± 0.99 | 2.92 ± 1.10 | 2.28 ± 0.89 | 2.83 ± 1.09 | 8.8e-01 | **< 1e−16** | 1.8e-01 |
| | | TM-score | 0.71 ± 0.17 | 0.69 ± 0.20 | 0.72 ± 0.16 | 0.70 ± 0.20 | 4.7e-01 | **< 1e−16** | 7.0e-02 |
| | | pLDDT | 72.62±11.95 | 65.18±12.92 | 73.16±11.74 | 65.74±12.81 | 9.3e-02 | **< 1e−16** | 4.5e-01 |
| | 2.40.50.140 | RMSD | 5.03 ± 0.96 | 2.99 ± 1.09 | 2.49 ± 0.90 | 2.85 ± 1.08 | **1.1e-08** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.32 ± 0.14 | 0.68 ± 0.20 | 0.72 ± 0.17 | 0.70 ± 0.20 | **6.3e-08** | **< 1e−16** | **2.4e-14** |

| Category | Target | Metric | Subnet. Supp. | Subnet. Maint. | ESM-2 Supp. | ESM-2 Maint. | t-Test ($p$) Supp. | t-Test ($p$) Maint. | K.S.-test ($p$) |
|---|---|---|---|---|---|---|---|---|---|
| | | pLDDT | 34.61 ± 8.13 | 63.92±13.16 | 68.78±12.37 | 65.73±12.64 | **1.1e-07** | **< 1e−16** | **< 1e−16** |
| | 2.60.40.10 | RMSD | 3.75 ± 1.26 | 3.03 ± 1.14 | 2.17 ± 0.69 | 2.89 ± 1.14 | **7.2e-13** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.53 ± 0.19 | 0.67 ± 0.20 | 0.79 ± 0.10 | 0.69 ± 0.20 | **5.4e-15** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 41.57±12.07 | 63.52±13.34 | 71.28 ± 9.64 | 65.13±13.02 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 3.20.20.80 | RMSD | 3.45 ± 0.94 | 2.96 ± 1.11 | 2.63 ± 0.66 | 2.86 ± 1.10 | **6.4e-06** | **< 1e−16** | **1.8e-07** |
| | | TM-score | 0.79 ± 0.10 | 0.68 ± 0.20 | 0.87 ± 0.07 | 0.69 ± 0.20 | **1.2e-05** | **< 1e−16** | **6.3e-08** |
| | | pLDDT | 49.76 ± 9.16 | 63.98±13.30 | 67.03 ± 7.20 | 65.46±12.91 | **8.5e-10** | **< 1e−16** | **< 1e−16** |
| | 3.40.190.10 | RMSD | 5.22 ± 1.21 | 3.00 ± 1.14 | 2.84 ± 1.06 | 2.89 ± 1.15 | **5.1e-11** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.32 ± 0.11 | 0.67 ± 0.20 | 0.63 ± 0.12 | 0.69 ± 0.20 | **2.6e-09** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 32.96 ± 5.12 | 64.06±13.44 | 66.66 ± 6.61 | 65.12±13.23 | **9.4e-16** | **< 1e−16** | **< 1e−16** |
| | 3.40.30.10 | RMSD | 4.49 ± 1.63 | 2.95 ± 1.09 | 1.75 ± 0.48 | 2.86 ± 1.11 | **1.7e-06** | **< 1e−16** | **1.2e-14** |
| | | TM-score | 0.46 ± 0.21 | 0.68 ± 0.20 | 0.84 ± 0.08 | 0.69 ± 0.20 | **1.7e-06** | **< 1e−16** | **2.5e-15** |
| | | pLDDT | 41.38±14.66 | 64.21±12.76 | 76.66 ± 7.22 | 65.51±12.75 | **1.7e-09** | **< 1e−16** | **< 1e−16** |
| | 3.40.50.150 | RMSD | 2.85 ± 1.23 | 2.94 ± 1.11 | 2.40 ± 0.77 | 2.84 ± 1.10 | **2.2e-03** | **< 1e−16** | **1.5e-03** |
| | | TM-score | 0.79 ± 0.16 | 0.68 ± 0.20 | 0.85 ± 0.10 | 0.70 ± 0.20 | **5.8e-04** | **< 1e−16** | **2.8e-04** |
| | | pLDDT | 62.15±12.19 | 63.90±12.84 | 72.03 ± 9.41 | 65.70±12.82 | **8.7e-10** | **< 1e−16** | **< 1e−16** |
| | 3.40.50.1820 | RMSD | 2.57 ± 0.70 | 2.95 ± 1.13 | 2.39 ± 0.68 | 2.87 ± 1.12 | **4.9e-05** | **< 1e−16** | **1.7e-02** |
| | | TM-score | 0.86 ± 0.07 | 0.68 ± 0.20 | 0.88 ± 0.06 | 0.69 ± 0.20 | **2.9e-06** | **< 1e−16** | **2.5e-03** |
| | | pLDDT | 66.94 ± 8.82 | 64.39±13.30 | 71.50 ± 6.95 | 65.29±13.09 | **7.4e-08** | **< 1e−16** | **8.1e-10** |
| | 3.40.50.300 | RMSD | 3.44 ± 1.11 | 2.99 ± 1.12 | 2.63 ± 0.76 | 2.87 ± 1.12 | **2.0e-08** | **< 1e−16** | **1.7e-10** |
| | | TM-score | 0.68 ± 0.15 | 0.67 ± 0.20 | 0.79 ± 0.13 | 0.69 ± 0.20 | **2.2e-08** | **< 1e−16** | **3.9e-12** |
| | | pLDDT | 50.62±12.41 | 63.22±13.53 | 68.34 ± 8.00 | 65.36±13.26 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 3.40.50.720 | RMSD | 4.26 ± 1.41 | 2.98 ± 1.11 | 2.47 ± 0.93 | 2.87 ± 1.11 | **3.5e-08** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.56 ± 0.19 | 0.68 ± 0.20 | 0.77 ± 0.16 | 0.70 ± 0.20 | **5.4e-09** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 42.35±11.23 | 65.14±13.51 | 72.20±10.85 | 65.42±12.87 | **1.2e-14** | **1.3e-04** | **< 1e−16** |
| Random | 100 seqs. | RMSD | 3.06 ± 1.17 | 2.92 ± 1.11 | 2.92 ± 1.17 | 2.84 ± 1.11 | **1.0e-03** | **< 1e−16** | 5.2e-01 |
| | | TM-score | 0.66 ± 0.20 | 0.69 ± 0.20 | 0.67 ± 0.21 | 0.70 ± 0.20 | **3.7e-04** | **< 1e−16** | 5.8e-01 |
| | | pLDDT | 63.15±12.98 | 65.23±13.21 | 64.74±12.34 | 65.71±12.95 | **3.4e-07** | **< 1e−16** | **5.3e-06** |
| | 1000 seqs. | RMSD | 3.64 ± 1.24 | 3.37 ± 1.25 | 2.81 ± 1.05 | 2.86 ± 1.12 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.57 ± 0.21 | 0.62 ± 0.21 | 0.70 ± 0.19 | 0.69 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 49.65±13.98 | 55.49±14.77 | 65.79±12.85 | 65.32±13.06 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | 200 seqs. | RMSD | 3.00 ± 1.16 | 2.95 ± 1.13 | 2.79 ± 1.12 | 2.85 ± 1.13 | **4.6e-10** | **< 1e−16** | **1.5e-03** |
| | | TM-score | 0.67 ± 0.20 | 0.68 ± 0.20 | 0.70 ± 0.20 | 0.70 ± 0.20 | **2.4e-12** | **< 1e−16** | **1.6e-04** |
| | | pLDDT | 63.04±14.38 | 64.62±13.40 | 65.62±13.21 | 65.50±13.01 | **5.6e-13** | **< 1e−16** | **1.3e-06** |
| | 2000 seqs. | RMSD | 5.21 ± 1.30 | 3.93 ± 1.48 | 2.85 ± 1.08 | 2.86 ± 1.12 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | TM-score | 0.34 ± 0.15 | 0.53 ± 0.22 | 0.70 ± 0.20 | 0.69 ± 0.20 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | | pLDDT | 34.59 ± 6.67 | 48.86±15.22 | 65.59±12.74 | 65.41±13.08 | **< 1e−16** | **< 1e−16** | **< 1e−16** |
| | Helix | RMSD | 3.23 ± 1.17 | 3.02 ± 1.12 | 2.92 ± 1.14 | 2.81 ± 1.09 | **< 1e−16** | **< 1e−16** | **2.1e-08** |
| | | TM-score | 0.60 ± 0.20 | 0.68 ± 0.20 | 0.65 ± 0.21 | 0.71 ± 0.19 | **< 1e−16** | **< 1e−16** | **1.4e-12** |
| | | pLDDT | 59.77±14.46 | 60.31±13.52 | 66.43±13.05 | 65.37±12.75 | **< 1e−16** | **< 1e−16** | **1.6e-13** |
| | Sheet | RMSD | 3.04 ± 1.20 | 3.08 ± 1.12 | 2.89 ± 1.17 | 2.82 ± 1.08 | **1.0e-16** | **< 1e−16** | **4.9e-09** |
| | | TM-score | 0.67 ± 0.20 | 0.66 ± 0.20 | 0.70 ± 0.19 | 0.70 ± 0.20 | **< 1e−16** | **< 1e−16** | **2.4e-07** |
| | | pLDDT | 60.29±14.11 | 60.14±13.66 | 63.77±13.53 | 66.21±12.57 | **< 1e−16** | **< 1e−16** | **< 1e−16** |