Scaling down for efficiency: Medium-sized protein language models perform well at transfer learning on realistic datasets

Luiz C. Vieira[a], **Morgan L. Handojo**[a], **and Claus O. Wilke**[a,1]

This manuscript was compiled on January 28, 2025

Protein language models (pLMs) can offer deep insights into evolutionary and structural properties of proteins. While larger models, such as the 15 billion parameter model ESM-2, promise to capture more complex patterns in sequence space, they also present practical challenges due to their high dimensionality and high computational cost. We systematically evaluated the performance of various pLMs across multiple biological datasets to assess the impact of model size on transfer learning. Surprisingly, we found that larger models not necessarily outperform smaller ones, in particular when data is limited. Medium-sized models, such as ESM-2 650M and ESM C 600M, demonstrated consistently good performance, falling only slightly behind their larger counterparts—ESM-2 15B and ESM C 6B—despite being many times smaller. Additionally, we compared various methods of compressing embeddings prior to transfer learning, and we found that mean embeddings consistently outperformed other compression methods. In summary, ESM C 600M with mean embeddings offers an optimal balance between performance and efficiency, making it a practical and scalable choice for transfer learning in realistic biological applications.

ESM | pLM embeddings | Feature compression | Transfer learning

Significance Statement

This work challenges the common belief that larger language models always yield better results, here in the context of protein biochemistry. By systematically comparing transformer models of different sizes in transfer learning tasks, we demonstrate that medium size models, such as ESM C 600M, frequently perform as well as or better than larger variants, especially when data is limited. These findings provide an efficient strategy for machine learning-based protein analysis. Smaller and more efficient models help democratize cutting-edge AI approaches, making them more accessible to researchers with limited computational resources.

Machine learning (ML) has significantly advanced the field of protein biochemistry (1–4). By leveraging large datasets and sophisticated models, ML techniques have improved the accuracy of protein property predictions (5–8). Transformer models, such as Evolutionary Scale Modeling 2 (ESM-2) (9), a pre-trained masked protein language model (pLM), have been particularly impactful due to their ability to capture valuable evolutionary relationships of protein sequences. The pLM models learn about protein properties during pre-training, where the models are trained to predict residues that have been masked in their input sequences (10, 11). The pre-training process causes the models to encode knowledge about protein biochemistry and protein evolution in the models' internal representations, known as embeddings, which encapsulate the biochemical characteristics of individual amino acids as well as complex higher-order interactions that reflect both local and global structural and functional properties of proteins (12–14). The capacity of pLMs to learn rich representations of proteins from vast amounts of unlabeled sequence data and subsequently apply this knowledge to secondary supervised tasks, a process known as transfer learning (15), has enabled their application in various downstream tasks including functional annotation, mutational effect analysis, and the design of novel proteins and peptides (8, 16–21). Finally, pLM embeddings can also be used for homology search, in particular when sequences are short or highly diverged (22, 23).

There has been a trend towards increasing the size of pLMs, following similar advancements in natural language processing, where model scaling laws predict that model performance systematically increases with increasing model size and commensurate increase in pre-training data (24, 25). Models such as the largest ESM-2 variant, with 15 billion parameters (13), and more recently ESM3, with a staggering 98 billion parameters (7), have demonstrated that scaling the model size can enhance performance by capturing more complex relationships in protein sequences. Although models such ESM3 can deliver notable gains in accuracy, their high computational cost hinders their development and broad utilization. For example, fine-tuning these large models is highly computationally demanding (26), limiting their use primarily to private industry or highly resourced laboratories. In fact, a recent study (27) questioned whether larger models are the best option, or if alternative training methods, such as improving data quality, increasing data quantity, or extending the number of training steps,

Author affiliations: [a]Department of Integrative Biology, The University of Texas at Austin, Austin, TX, United States of America

Please provide details of author contributions here.

The authors declare no competing interest.

[1]To whom correspondence should be addressed. E-mail: wilke@utexas.edu

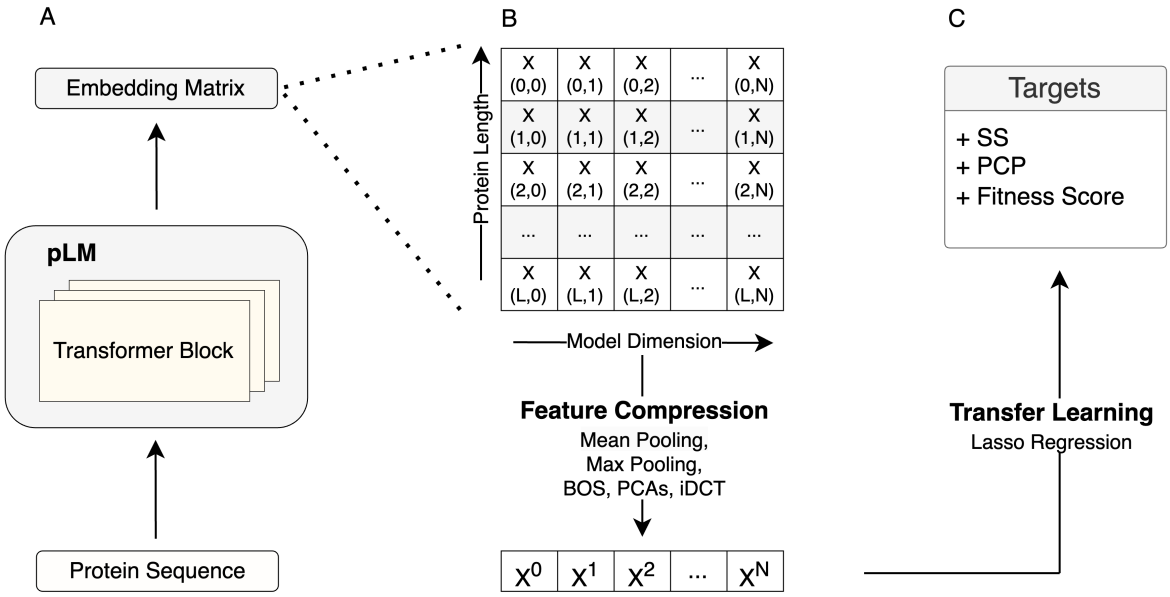

Fig. 1. Schematic view of the transfer learning approach used throughout this work. A) Protein representation (embeddings matrix) extraction using the pLMs. B) Embeddings matrix compression using various methods including mean pooling, max pooling, BOS, PCAs, and iDCT. C) Transfer learning to predict downstream tasks, such as secondary structure (SS), fitness, and physico-chemical properties (PCPs).

could boost smaller models' performance. In general, the field of AI applications in biology appears to be increasingly recognizing the benefits of smaller models. This trend is exemplified by the recent launch of the protein language model ESM-Cambrian (ESM C) (28), which has demonstrated impressive performance in protein contact prediction, even outperforming the much larger ESM-2 15B. This shift in thinking could provide a more accessible and cost-effective alternative for the broader research community, as smaller models might offer sufficient performance for many scientific applications at a lower cost.

Here, we systematically evaluate the impact of model size on transfer learning, considering models with parameter counts ranging from 8 million to 15 billion across various biological datasets. These datasets include 12 different metrics calculated from proteins in the PISCES dataset (29) and 41 deep mutational scanning (DMS) datasets (30). Our findings suggest that the effectiveness of transfer learning is influenced by the choice of embedding compression method, dataset size, protein length, and model size. Larger models perform better with larger datasets, underscoring the need for sufficient data to maximize their utility. However, when data is limited, medium-sized models perform comparably to, and in some cases outperform, larger models. In addition, while the impact of compression methods varies depending on dataset type, mean embeddings generally yield excellent results across tasks. These findings emphasize the importance of selecting model size and embedding method based on the unique characteristics of the data and task to enhance both the efficiency and accuracy of transfer learning in protein applications.

Results

Mean embeddings outperform all other compression methods. Despite the general success of transformer model embeddings, their high dimensionality poses challenges in practical applications, particularly in transfer learning scenarios (31, 32). Embeddings typically need to be compressed before any downstream prediction tasks are performed. The most commonly used compression strategy is simply averaging embeddings over all sites (mean pooling) (5, 17, 33–35), though other methods, such as inverse Discrete Cosine Transform (iDCT), have also been proposed (22). Because mean pooling averages over the contributions of all sites in a sequence, this strategy may not retain all critical information, in particular in deep mutational scanning (DMS) applications, where one or a few amino acid changes can significantly affect a protein's thermodynamic stability or enzymatic activity (36).

To assess the effect of compression method on transfer learning performance, we systematically explored a wide range of different compression methods. We benchmarked the compression methods on 40 DMS datasets (which tend to be datasets of single or few point mutations relative to a reference sequence) and on a set of diverse protein sequences taken from PISCES (A Protein Sequence Culling Server), for which we calculated a variety of target variables including physicochemical properties (PCPs), instability index, amino acid frequencies, and secondary structure frequencies. In our prediction pipeline, for a given protein sequence, we extracted embeddings from the last hidden layer of the ESM-2 150M model, compressed them via one of the methods, and then used the compressed embeddings as input features in

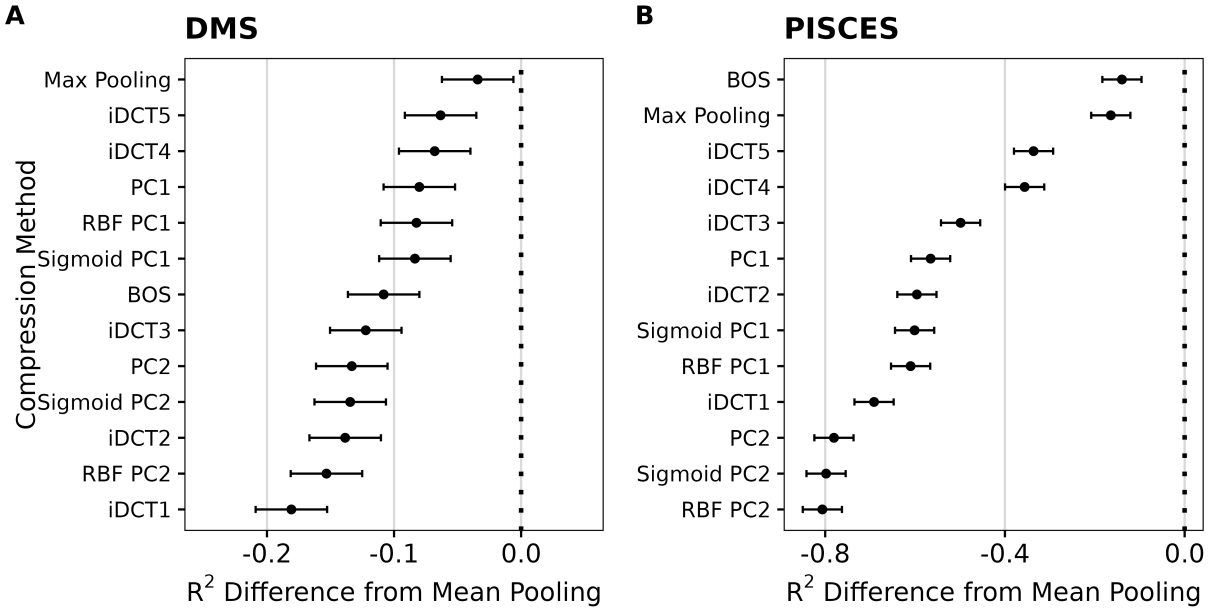

Fig. 2. Mean reduction in R^2 when embeddings are compressed with methods other than mean pooling. A) Results for DMS data. B) Results for diverse protein sequences (PISCES data). In all cases, the y-axis represents different compression methods and the x-axis shows the resulting difference in R^2. Dots represent the fixed effects estimates from mixed-effects modeling, and error bars represent 95% confidence intervals.

regularized regression models (LassoCV) to predict the target of interest (Figure 1).

For both the DMS datasets and the PISCES sequences, we generally found that mean pooling outperformed all other compression methods we considered (SI Appendix, Fig. S1 and Fig. S2). However, the extent to which mean pooling performed better depended on the type of data. For diverse protein sequences, mean pooling was strictly superior in all cases, and in many cases by a wide margin (Fig. S2). By contrast, for DMS data, some alternative compression methods, including max pooling, iDCT, and PCA, were slightly better than mean pooling on some datasets, even if on average mean pooling was superior (Fig. S1). These findings suggest that mean pooling does to some extent average out effects from individual sites, yet on the whole it still tends to perform better than the alternatives.

To summarize these findings more systematically, we fit linear mixed-effects models to the results of all compression methods and either all DMS datasets or all PISCES prediction targets, respectively. In this setup, we treated the type of compression as a fixed effect and the dataset/prediction target as a random effect, to isolate the impact of the compression method while accounting for the variability between datasets. This analysis showed that mean pooling was, on average, significantly better than all other alternatives we considered, in both types of datasets (Figure 2). For DMS data, mean pooling led to an increase in variance explained (measured by R^2 from the regularized regression, calculated on a hold-out test set) between 5 and 20 percentage points (Figure 2A). For diverse protein sequences, the difference was even more stark, where mean pooling led to an increase in variance explained between 20 and 80 percentage points (Figure 2B).

In aggregate, these results show that mean pooling is strictly superior in transfer-learning applications where the input sequences are widely diverged, and it performs well also with DMS data. Therefore, for our subsequent assessment of model size on transfer learning, we only considered mean pooling throughout.

Moderately sized models perform well in transfer learning. We next turned to the effect of model size. We considered all 6 ESM-2 models, ranging from 8 million parameters (EMS2 8M) to 15 billion parameters (ESM-2 15B). We also included the older model ESM1v with 650 million parameters, which was developed specifically for variant effect prediction. Importantly, because ESM1v only accepts sequences up to 1,022 residues, we excluded any DMS datasets with longer proteins from the main analysis, which reduced our DMS data to 36 distinct datasets. Results on the remaining datasets but excluding EMS1v are presented in SI Appendix, Fig. S3. In addition to ESM1v, we also evaluated several more recently published models, including AMPLIFY (120M, and 350M) and ESM C (300M and 600M, 6B).

We found that, overall, larger models tended to yield better results for transfer learning tasks across both types of datasets, as expected (Figure 3). However, the performance improvements for very large models (with 3, 6, or 15 billion parameters) were moderate or small. For example, the medium-sized model ESM-2 150M consistently performed well across many targets—particularly in DMS analyses, where it slightly outperformed ESM-2 3B despite being twenty times smaller (Figure 3A). Additionally, all ESM C models outperformed ESM-2 15B, and scaling ESM C up to 6B offered only a minor advantage over the 600M parameter variant of the same model family. On the PISCES dataset, all models with 150M parameters or more delivered comparably strong performance, with the large ones reaching a close to optimal results (Figure 3B). These findings indicate that, for the datasets evaluated, the ESM C 600M model strikes an ideal balance between computational efficiency and

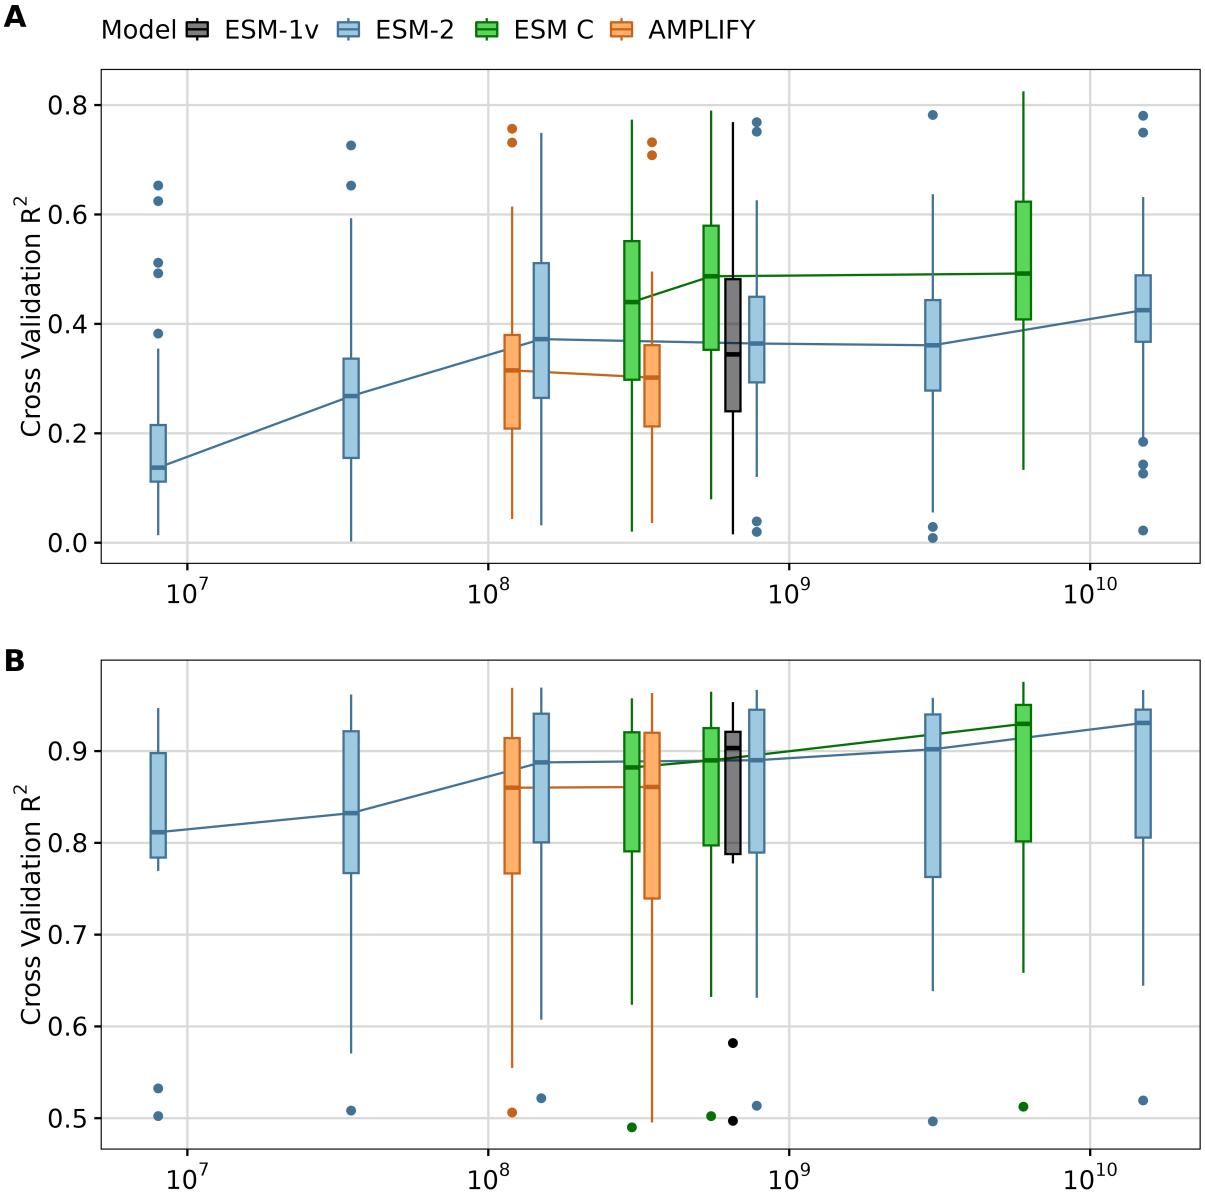

Fig. 3. Impact model size on transfer learning. A) LassoCV regression results using mean embeddings from different pLMs for 36 DMS datasets. B) LassoCV regression results using mean embeddings from different pLMs for 12 targets originated from the PISCES dataset. The y-axis displays the average R^2 scores from 5-fold cross-validation for each dataset or target. The x-axis represents the number of parameters in each model, with different models grouped by color: ESM1v (black), ESM-2 models (blue), ESM C (green), and AMPLIFY (orange).

predictive power, offering robust results in the context of transfer learning.

Although ESM1v was specifically designed and trained on UniRef90 (37) for variant effect prediction, its limitation to sequences no longer than 1,022 residues restricts its applicability. To evaluate sequences exceeding this length, we turned to the ESM-2, ESM C, and AMPLIFY models. While these models were also pretrained with maximum sequence lengths of 1,022 for ESM-2 and 2,048 for ESM C and AMPLIFY, their use of a rotary attention mechanism allows them, in principle, to handle longer sequences. However, this capability remains underexplored. This part of our analysis was inconclusive, as none of the larger sequences (four in total) achieved particularly high predictability scores (SI Appendix, Fig. S3). Nonetheless, it is worth noting that ESM C models seem to display better performance overall for larger sequences, but more data is required to truly evaluate this question. At the very least, we can conclude that there is no strong evidence suggesting that ESM-2 and ESM C cannot be used with sequences longer than the maximum sequence length they were trained on.

Sample size drives transfer learning performance in pLMs. Next, we investigated whether sample size impacts transfer learning performance. To explore this question, we selected three DMS datasets with good model performance ($R^2 > 0.6$ for the larger models), large size (> 1000 observations), and

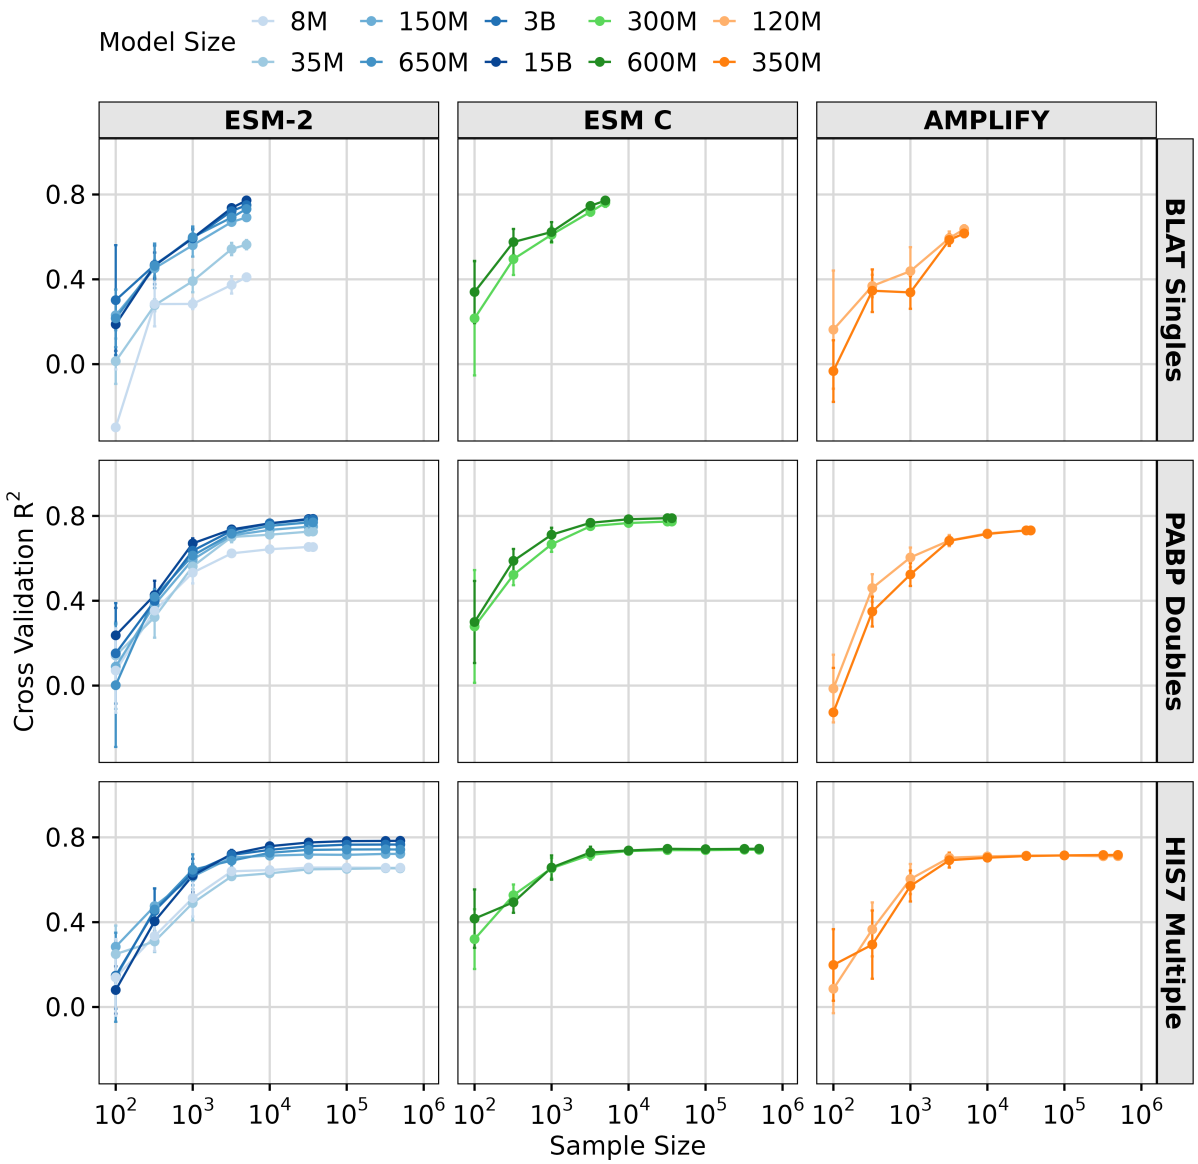

Fig. 4. Effect of sample size on transfer learning. Results of LassoCV regression on three DMS datasets, using downsampled subsets ranging from 100 to the maximum number of samples in the dataset. The y-axis represents the averaged 5-fold cross-validation R^2 scores, and error bars represent the standard deviation. The x-axis shows the sample sizes tested. The colored lines represent different models and model sizes: ESM-2 8M, 35M, 150M, 650M, 3B, and 15B (Blues), ESM C 300M, and 600M (Greens) and AMPLIFY 120M, and 350M (Oranges).

differing complexity (single, double, and multiple mutations). We then progressively downsampled each dataset into subsets, ranging from 100 observations to the full dataset size. For each subset, we ran LassoCV regression with five-fold cross-validation. We found that smaller datasets negatively affected transfer learning performance, resulting in reduced accuracy for all three datasets and all three models (Figure 4). Notably, aside from the smallest two ESM-2 models (8M and 35M parameters) and AMPLIFY (120M and 350M parameter), all models performed comparably when sufficient data was available ($\gtrsim 10^4$ observations).

We hypothesized that meaningful protein features are distributed across more dimensions in larger models, leading to improved performance as the number of retained features increases. To test this hypothesis, we plotted the number of features retained by LassoCV for the three DMS datasets and models. We found that the ESM-2 15B model starts outperforming the smaller models when it starts utilizing a greater number of features (SI Appendix, Figure S4). Interestingly, the HIS7 dataset was the only one that showed clear saturation in the number of features being used across all models, for sample sizes in excess of $\sim 10^4$ (SI Appendix, Figure S4). For all models, the number of features used plateaued well below the models' respective embedding dimensions. By contrast, in the other two datasets, for the largest models, the number of retained features continued to increase until the maximum dataset size was reached, suggesting that with a larger sample size performance could improve further. On the other hand, several medium-sized models showed this saturation with fewer features and a smaller sample size.

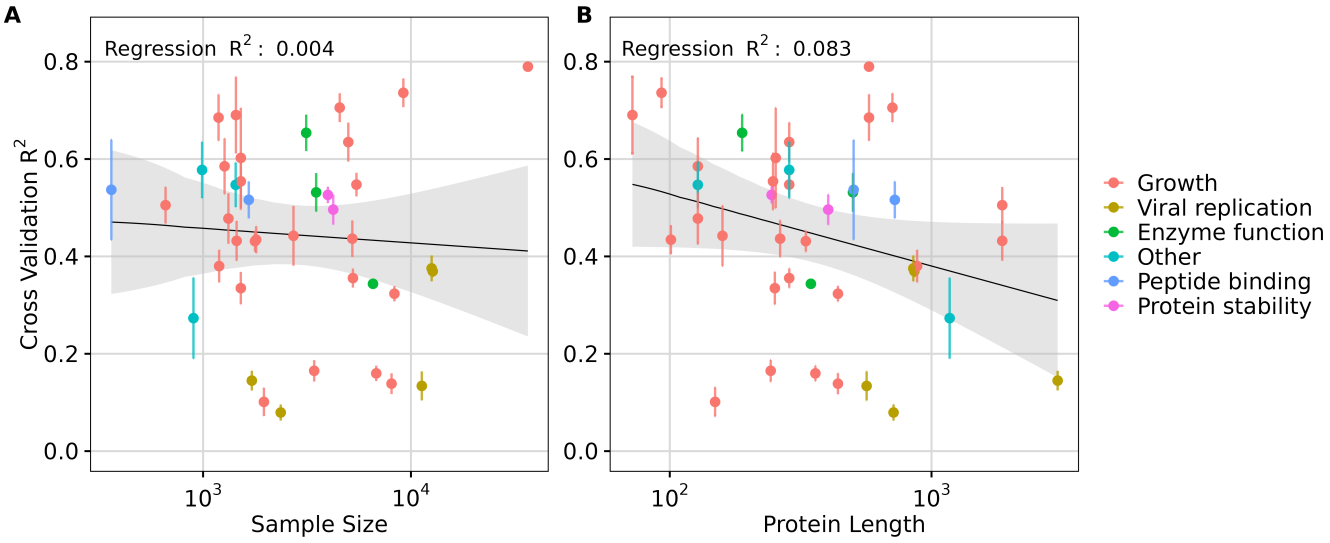

Fig. 5. Effect of sample size and protein length in transfer learning. A) Effect of **sample size** on transfer learning using ESM C 600M model embeddings quantified by the regression results (R^2 score). B) Effect of **protein length** on transfer learning using ESM C 600M model embeddings quantified by the regression results (R^2 score). Each dot represents a dataset, and its color indicates the corresponding data measurement type. In each subplot, the y-axis represents the averaged 5-fold cross-validation R^2 scores, and error bars represent the standard deviation.

Interestingly, the number of features selected by LassoCV when using embeddings from ESM-2 650M and ESM C (300M and 600M) required only half of the number of features utilized by the larger ESM-2 15B model, and yet achieved comparable performance (SI Appendix, Figure S4). This result highlights the superior (or less sparse) protein representation encoded by the medium size pLMs, and in particular ESM C. Furthermore, utilizing fewer features offers significant benefits in transfer learning by reducing the risk of overfitting and by improving generalization. These results suggest that the larger embedding space provided by larger models cannot be fully taken advantage of by moderately sized datasets of fewer than 10^4 observations, and this may be one of the main reasons moderately sized models are frequently sufficient and perform just as well as the largest models.

Transfer learning is limited by data quality. Since the preceding analysis was performed by downsampling, where we expected prediction performance to systematically decline for more extreme ranges of downsampling, we next asked whether a similar relationship between dataset size and prediction accuracy could be observed across all DMS datasets, which inherently had differences in sample size. We additionally explored the effect of protein length and the type of data measured on model performance. We fit separate linear regression models, using transfer learning model performance (R^2 score) as the dependent variable and either DMS dataset size or protein length as the independent variable. Our analysis revealed that the percent variation explained by sample size was minimal (for example, $R^2 = 0.004$ for ESM C 600M, Figure 5A), indicating that while performance may improve with larger datasets, sample size is not the primary driver of transfer learning effectiveness in the 40 DMS datasets analyzed. Notably, most datasets had a sample size of over a thousand (SI Appendix, Fig. S5), a number close to optimal for transfer learning, as suggested by our downsampling results (Figure 4). However, a small sample size appears to have a more pronounced negative impact on the performance of the AMPLIFY family of models, which however also did not demonstrate strong overall performance (SI Appendix, Fig. S8).

By contrast, protein length accounted for a meaningful proportion of variation observed (for example, $R^2 = 0.083$ for ESM C 600M), suggesting a modest influence of protein length on model performance. Specifically, longer protein sequences, particularly those exceeding the length of 1,022 residues, were associated with decreased model performance in transfer learning (Figure 5B and SI Appendix, Fig. S3). This trend was broadly consistent across different models (SI Appendix, Figs. S6, S7, and S8).

We also broke down the datasets by dataset type (growth, viral replication, peptide binding, etc.) but did not see a strong trend (Figure 5). One exception were viral proteins, which tended to be among the longest sequences in the DMS data and also exhibited some of the lowest R^2 scores. This observation suggests that either pLMs embeddings may struggle to capture effective representations of viral proteins specifically, or that the embeddings generally lack representational power for larger proteins. We note that the vast majority of DMS datasets (all but four) used sequencing as the primary measurement method, and frequency measurements via sequencing are subject to numerous sources of noise, including sampling errors, PCR amplification errors, and genetic drift during competition (38). These factors may partially explain the limited predictive power we observed for several of the datasets. When datasets are inherently noisy no machine learning model can possibly make good predictions on them.

Discussion

We have evaluated the performance of ESM-2, ESM1v, ESM C and AMPLIFY embeddings across various model sizes (from 8 million to 15 billion parameters) in transfer learning tasks on a wide range of different biological datasets. We have found that larger models do not consistently outperform smaller ones, especially when data is limited. In fact, medium-size models, such as ESM-2 150M, ESM-2 650M or ESM C 600M, have demonstrated consistently good performance. This observation suggests that model size should be carefully aligned with dataset size and data type to optimize transfer learning performance. Additionally, we have evaluated different methods for compressing per-token embedding matrices, and we have found that, on average, mean embeddings consistently outperform other commonly used compression methods. However, for fairly homogeneous DMS datasets, alternative compression methods including max pooling and iDCT can on occasion achieve better performance than mean embeddings.

Scaling large language models (LLMs) to ever larger numbers of parameters has become a trend since Kaplan et al. (39) showed that there is a power-law relationship between model size and performance. This scaling law has driven the development of increasingly larger models (24, 25). Since the introduction of Generative Pre-Trained Transformers (GPTs) (40) in 2018, LLMs have scaled to trillions of parameters (41). This tendency has also extended to biology, with the latest protein language models (pLMs) reaching an astonishing 98 billion parameters (7). Although larger pLMs have demonstrated enhanced performance by capturing more complex relationships in protein sequences (13), a recent study has shown that as models scale, they become more prone to overfitting, often preferring to predict the wild-type residue in masked tasks (27). This observation aligns with the known LLM scaling laws, which suggest that performance improvements from scaling are only achieved if the amount of training data is scaled as well (39). As ESM-2 models were not trained with data scaled across different model sizes, their performance gains may have been limited, in particular for the larger models.

While fine-tuning pLMs can yield good predictions across various tasks (42), this approach necessitates a sufficient amount of training data (43, 44). In cases where data is limited, transfer learning is often a better alternative (45). Lasso regression using embeddings as input is particularly useful in this context, as it prioritizes simplicity (fewer features) while fitting the training data well (low error) (46). However, this approach, by construction, will not necessarily take advantage of the full embeddings from the largest models when transfer learning on a smaller dataset. Fundamentally, this limitation arises from the protection against overfitting that Lasso regression provides, in particular when the regularization parameter is chosen via cross validation as we have done here.

Smaller models offer many advantages, such as reduced inference costs and improved compatibility with compact hardware, providing substantial benefits to protein studies. In the context of regular (i.e., non-biological) large language models, Hoffmann et al. (47) showed that training a smaller model with more data (70 billion parameters and 1.4 trillion tokens) outperformed larger models with four times more parameters when pre-trained on less data. Our findings show similarly that the more recently developed smaller ESM C models in many cases outperform the older ESM-2 models, even for the model variants that have over an order of magnitude more parameters. While the model architecture of ESM C models is slightly different from ESM-2 models, we suspect the main benefit of the newer models comes from an improved training schedule. Overall, our results align with other studies indicating that pLMs beyond 650 million parameters have limited advantages in many application areas, only showing true benefits when used for tasks focused on structure prediction (48–50).

In addition to evaluating model size, we also systematically evaluated various feature compression methods. Compression methods are required to reduce the position-dependent embedding matrix generated by pLMs into a comparatively much lower-dimensional feature vector whose dimension is independent of the protein length. The compression method iDCT in particular has been proposed as a better alternative to mean pooling (22). However, we could not recapitulate this finding here, except for some specific DMS datasets. As a general and simple approach, mean pooling seems to be the best option. We caution however that our results do not imply no better compression methods exist. For any particular dataset, it is generally possible to find a better dataset-specific compression method, by training the compression method on the data (51–53). Our recommendation here is to use mean pooling as a reasonable starting point, and then potentially consider dataset-specific compression methods in case mean pooling does not yield adequate results.

Along the same lines, we used the final embedding layer of each model for transfer learning but other choices are possible. In particular, there is some evidence in the literature that using intermediate layers or multiple layers can yield better transfer learning performance than using the last layer (22, 49, 54). However, all available evidence indicates that any optimizations with regards to optimal layer choice will be dataset dependent. For any specific application task, it may always be worthwhile to perform optimizations with respect to which pLM, compression method, or embedding layer(s) are used, as long as good care is taken that a final test dataset is set aside and has never influenced any of these modeling choices. In fact, the comprehensive prediction benchmark ProteinGym has previously suggested that model selection should be guided by the dataset itself, as performance can vary significantly across assays (55).

We note that while working on this project, we encountered numerous reproducibility issues with computed embeddings that we think the field needs to be aware of. These issues arise because calculating embeddings for even a single protein sequence requires billions of floating-point operations and small rounding errors or other inaccuracies can accumulate to the point where the final result has widely diverged from what it should be. The only strategy that avoids this issue entirely is to calculate embeddings on the exact same hardware architecture and GPU drivers that were used to train a given model. In practice, this will frequently not be possible. However, we found that models can be re-used across architectures, even from different vendors (NVIDIA vs AMD vs Apple Silicon), as long as the user is aware of the most common pitfalls and ensures they are avoided.

We can group reproducibility issues into three distinct categories: (i) issues with software libraries/drivers; (ii) issues due to numerical types; and (iii) issues due to batch processing of multiple sequences at once. We encountered the first issue when trying to run ESM-2 models on Apple Silicon GPUs via MPS (Metal Performance Shaders). We found that calculated embeddings had accumulated rounding errors to the point that even the first significant digit in most embedding dimensions differed from the correct value. Issues filed on the PyTorch GitHub repository (e.g., issue #84936) suggest that these problems were due to the specific floating point math implementation used by Apple in their drivers. We note that these issues seem to have been resolved since the release of MacOS Sequoia and PyTorch version 2.5 in the fall of 2024 but may still be present on many Apple devices not yet updated to the latest software.

We encountered problems with numerical types in the ESM C models, which by default use reduced accuracy floating point types and autocasting to speed up performance during inference. Autocasting in particular dynamically alters numerical precision, potentially leading to inconsistencies across runs or hardware configurations. These precision variations were more prominent when passing an input array with multiple sequences, and they were overall sufficiently severe that we felt we could not generate reproducible embeddings from ESM C models in the default configuration. However, we were able to circumvent these issues by using the `float32` type consistently throughout and disabling autocasting.

Finally, we encountered issues with padding with the AMPLIFY models. In all models, when processing a batch of multiple sequences with different lengths, shorter sequences need to be padded to the maximum length. This padding should not affect computed embeddings, but if a transformer model does not properly mask padded sites when calculating attention then the padding can influence output embeddings, which will result in poor reproducibility. In models with this issue, the embeddings for one sequence will depend on the lengths of other sequences processed in the same batch. We were able to circumvent this issue by not using any batch processing and always processing each sequence individually without any padding.

Fortunately, we have found that it is relatively straightforward to check for the presence of any of these reproducibility issues. Embeddings calculated on a CPU using the `float32` type, on individual sequences without batching, seem to always be reliable, for all models. Therefore, we recommend to always compare GPU results to this gold standard and to only proceed if any observed differences between CPU and GPU results are minor.

In summary, our study challenges the general preference for larger protein language models in biology. In particular, it suggests that in common transfer learning applications, most available datasets are too small for the largest available pLMs. While larger models can achieve good performance in principle, smaller models frequently perform similarly or deliver stable performance with smaller samples, making them more accessible and practical in the more common scenarios working biologists may experience in practice. Additionally, mean embeddings consistently perform well across datasets, suggesting that alternative compression methods will rarely be needed. In general, mean embeddings calculated from models of intermediate size, such as ESM C 600M, are likely sufficient for transfer learning on the vast majority of data sets encountered in biological research.

Materials and Methods

Data collection. We used two types of datasets in our analysis. First, we obtained 41 deep mutational scanning (DMS) datasets previously curated in Ref. (30). A complete list of all datasets included is available at: https://github.com/ziul-bio/SWAT/blob/main/data/DMS_metadata.csv. Based on the collected data, we created FASTA and CSV files containing the mutated sequences for each protein, along with the corresponding target values of the mutations. The code to reconstruct sequences is available at: https://github.com/ziul-bio/SWAT/blob/main/DMS_pre_processing.ipynb.

Second, we downloaded the PISCES dataset (29), which contains a collection of proteins with known structures and at most 50% pairwise sequence similarity. We filtered this dataset to include only proteins with lengths between 64 and 1022 residues, resulting in a total of 23,487 sequences. The minimum length of 64 residues was determined by a requirement in our implementation of the inverse Discrete Cosine Transform (iDCT, see below), while the maximum length of 1022 residues corresponds to the upper limit supported by ESM-2 models. The final dataset is available at: https://github.com/ziul-bio/SWAT/blob/main/data/PISCES/pisces_len64-1022.fasta. From this dataset, we generated 12 target variables, including a selection of physicochemical properties, amino acid frequencies, and secondary structure frequencies (SI Appendix, Table S11). For physicochemical properties, we selected six interpretable features from those available in the Peptides python package (version 0.3.2), accessible on PyPI. For the selection of amino acid frequencies, we arbitrarily chose alanine, cysteine, and leucine. For secondary structure targets, we selected the three most common ones (alpha helix, beta strand, and coil), since all others were rare in most sequences. Secondary structure targets were computed using the DSSP module provided by Biopython. The code required to generate these targets is available at: https://github.com/ziul-bio/SWAT/blob/main/PISCES_pre_processing.ipynb

Computing protein embeddings from ESM-2 and ESM1v. We calculated protein embeddings from the ESM-2 family of protein language models (13) and from the older ESM1v model (56) using the same approach. For all model variants, the model's internal representation is a matrix of embeddings in $\mathbf{R}^{n \times d}$ dimensions, where n represents the protein sequence length and d represents the embedding dimension, which differs for different model variants and generally increases for models with more parameters. For each protein sequence in our datasets, we fed the sequence into the respective model and obtained the corresponding embeddings from the final hidden layer. For all models, we obtained embeddings using the `extract.py` script available at the ESM-2 GitHub repository (https://github.com/facebookresearch/esm/blob/main/scripts/extract.py). This script allows us to define three output representations: mean representation (the embeddings averaged across the sequence length), BOS representation (the CLS token or beginning of the sequence), and per token representation (the full embedding matrix). We used this script to obtain all three representations. The mean and BOS representations were used as it is, and the per token representation was further processed as discussed below.

To extract ESM-2 embeddings of sequences longer than 1022 residues, we modified the script by increasing the max truncation length of sequences to 5000. See the script available at: https://github.com/ziul-bio/SWAT/blob/main/scripts/extract_esm2.py

Computing protein embeddings from ESM C. To extract protein mean embeddings from the ESM C models, we followed the instructions provided on the Evolutionary Scale GitHub page: https://github.com/evolutionaryscale/esm. Based on these instructions, we developed a custom extraction script, available at: https://github.com/ziul-bio/SWAT/blob/main/scripts/extract_ESMC.py. Before running any of the ESM C models, we had to make a few

modifications to ensure the reproducibility of results. First, we changed the data type from `bfloat16` to `float32` in the `esmc.py` file. The `bfloat16` data type is optimized for efficiency and reduced memory usage during training, but due to its half-precision nature, it can sometimes introduce numerical precision issues. By switching to `float32`, we ensured higher precision, which is particularly important when generating embeddings for downstream tasks or analyses. This adjustment also ensured consistency in results independent of the hardware being used. Second, we disabled autocasting, by setting `enabled=FALSE` and `dtype=torch.bfloat32` in the `torch.autocast()` function. Autocasting dynamically adjusts the precision during computations to optimize performance, but this can lead to inconsistencies in results across different runs or hardware setups. Disabling autocasting and explicitly setting the data type to `float32` ensures that all calculations are performed with consistent precision, making the embeddings more reliable and reproducible.

As of this writing, the ESM C 6B model is not publicly available. Instead, access requires enrollment in the beta access program and the creation of an access token for API calls. Due to this limited access, we were unable to obtain results for the DMS HIS dataset, which contains 500,000 sequences. Another limitation concerns the maximum sequence length: While ESM C 300M and 600M can handle sequences longer than 1022 residues out of the box without any modifications, ESM C 6B, accessed through the API, has a preset maximum sequence length of 2048. This restriction excluded the dataset 'POLG Sun2014' which the protein has 3033 residues.

Computing protein embeddings from AMPLIFY. To extract embeddings from the AMPLIFY model, we built a custom Python script, available at: https://github.com/ziul-bio/SWAT/blob/main/scripts/extract_AMPLIFY.py. We then downloaded the model through Hugging Face, and after having the model locally, we made a few modifications to ensure reproducibility. First, we normalized the last hidden layer embeddings. The ESM-2 and ESM C models have the output of the last hidden layer normalized with `pytorch.nn.LayerNorm`. To ensure that results were comparable between these models and the AMPLIFY model we did the same for the AMPLIFY embeddings, which are not normalized by default. Second, AMPLIFY has a maximum sequence truncation defined as 2048. To be able to handle longer sequences, we modified the file `config.json` downloaded with the model and changed the parameter `max_length` to 5000.

Per-token embeddings compression. To compress the per-token representation matrix, we employed several techniques from scikit-learn (57), such as `MinMaxScaler`, PCA (Principal Component Analysis), and Kernel PCA. Initially, the embeddings matrix was scaled using `MinMaxScaler(feature_range=(-1, 1))` applied separately to each feature dimension along the protein sequence. To compress the embeddings, we applied PCA to the scaled embeddings matrix, treating the sequence length as features and the model dimensions as samples and retaining the first two PCs. Similarly, we used kernel PCA with either radial basis functions (RBF) or sigmoid kernels, using the scikit lern function `KernelPCA()`. Following the transformations to extract a single vector of fixed length, the first and second components of the PCAs were used and subsequently named PC1/PC2, RBF PC1/RBF PC2, and Sigmoid PC1/Sigmoid PC2, respectively.

Lastly, we employed an inverse Discrete Cosine Transform (iDCT) quantization method as a method of embeddings compression, as described in the Protein Ortholog Search Tool (PROST) (22). iDCT is particularly adept at capturing fine-grained details in data while reducing dimensionality by discarding high-frequency components, enhancing the signal-to-noise ratio in the embeddings (58, 59). We applied the Discrete Cosine Transform (DCT) to the embeddings, retained the top rows and columns, and then performed the inverse DCT (iDCT) to reconstruct the embeddings. The reconstructed matrix was reshaped into a 1D vector, with final length determined by the number of retained rows and columns. We experimented with various configurations, retaining subsets of rows and columns such as 5x44, 10x64, 10x128, 10x512, and 10x640. These configurations correspond to fixed-length embeddings of 220, 640, 1280, 5120, and 6400, respectively, which we designated as iDCT1, iDCT2, iDCT3, iDCT4, and iDCT5.

Predicting Protein Fitness and Properties. To evaluate the predictive performance of compressed protein embeddings, we employed LassoCV, a cross-validation-based version of the Lasso algorithm (46). LassoCV retains the feature selection properties of Lasso by reducing the coefficients of less relevant features to zero while selecting the optimal regularization parameter through cross-validation, thereby identifying the most important predictors.

All modeling was conducted in Python using methods from scikit-learn, including LassoCV, cross-validation (KFold), and performance metrics (R^2, MAE, RMSE, and Spearman's ρ). Before any modeling, the dataset features were scaled using `MinMaxScaler(feature_range=(-1, 1))`. Using a 5-fold cross-validation strategy, the data was then divided into five parts, with four parts used for training and one for testing during each round of cross-validation.

To determine the optimal regularization parameter α within each round of cross-validation, the LassoCV algorithm as implemented by scikit-learn generates a sequence of 100 α values, spanning three orders of magnitude in range. For each α, Lasso regression is performed on the training data, and a 3-fold cross-validation evaluates model performance. The optimal α is selected based on the lowest cross-validation error, and a final Lasso model is fit using the best α and the entire training dataset. Using this final Lasso model, we then calculated performance metrics for both the training and the test set.

This process was repeated five times and the average and standard deviation of the performance metrics on the cross-validation test sets are reported.

Statistical Analysis. To assess the statistical significance of differences among compression methods, we employed linear mixed-effects models using the lme4 package in R (60). This method is well-suited for analyzing data with hierarchical or nested structures. In our analysis, the fixed effect was the compression method (e.g., Mean, BOS, PCA, kernel PCA, and iDCT), while the random effect accounted for variability across the different datasets. To evaluate the significance of the fixed effects estimates we used the package multcomp (61) in R to calculate 95% confidence intervals around each estimate, taking into account that we were testing multiple hypotheses at once (one for each compression method).

Code availability. All the code used to generate the metadata, embeddings, compressed embeddings and models are available at https://github.com/ziul-bio/SWAT.

ACKNOWLEDGMENTS. This work was supported by National Institutes of Health grants R01 AI169462, R01 AI148419, and R01 GM088344. The Texas Advanced Computing Center (TACC) provided high-performance computing support. C.O.W. also acknowledges support from the Blumberg Centennial Professorship in Molecular Evolution. The authors thank Aaron Feller for support and discussions on the topic of large language models.

References

1. J Jumper, et al., Highly accurate protein structure prediction with alphafold. *Nature* **596**, 583–589 (2021).
2. M Baek, et al., Accurate prediction of protein structures and interactions using a three-track neural network. *Science* **373**, 871–876 (2021).
3. JL Watson, et al., De novo design of protein structure and function with rfdiffusion. *Nature* **620**, 1089–1100 (2023).
4. V Zambaldi, et al., De novo design of high-affinity protein binders with alphaproteo (arXiv preprint arXiv:2409.08022) (2024).
5. A Villegas-Morcillo, et al., Unsupervised protein embeddings outperform hand-crafted sequence and structure features at predicting molecular function. *Bioinformatics* **37**, 162–170 (2021).
6. J Laimer, H Hofer, M Fritz, S Wegenkittl, P Lackner, Maestro-multi agent stability prediction upon point mutations. *BMC bioinformatics* **16**, 1–13 (2015).
7. T Hayes, et al., Simulating 500 million years of evolution with a language model. *Science* p. eads0018 (2025).
8. X Pan, et al., Identification of protein subcellular localization with network and functional embeddings. *Front. Genet.* **11**, 626500 (2021).
9. A Rives, et al., Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *PNAS* **118**, e2016239118 (2021).
10. T Bepler, B Berger, Learning the protein language: Evolution, structure, and function. *Cell Syst.* **12**, 654–669 (2019).
11. A Chandra, L Tünnermann, T Löfstedt, R Gratz, Transformer-based deep learning for predicting protein properties in the life sciences. *Elife* **12**, e82819 (2023).
12. M Heinzinger, et al., Modeling aspects of the language of life through transfer-learning protein sequences. *BMC Bioinforma.* **20**, 1–17 (2019).
13. Z Lin, et al., Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science* **379**, 1123–1130 (2023).
14. A Elnaggar, et al., Prottrans: Toward understanding the language of life through self-supervised learning. *IEEE Transactions on Pattern Analysis Mach. Intell.* **44**, 7112–7127 (2021).
15. T Bepler, B Berger, Learning protein sequence embeddings using information from structure (arXiv preprint arXiv:1902.08661) (2019).
16. TN Dong, G Brogden, G Gerold, M Khosla, A multitask transfer learning framework for the prediction of virus-human protein–protein interactions. *BMC Bioinforma.* **22**, 1–24 (2021).
17. JR Randall, LC Vieira, CO Wilke, BW Davies, Deep mutational scanning and machine learning for the analysis of antimicrobial-peptide features driving membrane selectivity. *Nat. Biomed. Eng* **8**, 1–12 (2024).
18. EC Alley, G Khimulya, S Biswas, M AlQuraishi, GM Church, Unified rational protein engineering with sequence-based deep representation learning. *Nat. Methods* **16**, 1315–1322 (2019).
19. DJ Diaz, AV Kulikova, AD Ellington, CO Wilke, Using machine learning to predict the effects and consequences of mutations in proteins. *Curr. Opin. Struct. Biol.* **78**, 102–518 (2023).
20. B Allman, LC Vieira, DJ Diaz, CO Wilke, A systematic evaluation of the language-of-viral-escape model using multiple machine learning frameworks (bioRxiv, doi:10.1101/2024.09.04.611278) (2024).
21. JM Taft, et al., Deep mutational learning predicts ace2 binding and antibody escape to combinatorial mutations in the sars-cov-2 receptor-binding domain. *Cell* **185**, 4008–4022 (2022).
22. M Kilinc, K Jia, RL Jernigan, Improved global protein homolog detection with major gains in function identification. *Proc. Nat. Acad. Sci.* **120**, e2211823120 (2023).
23. AV Kulikova, JK Parker, BW Davies, CO Wilke, Semantic search using protein large language models detects class II microcins in bacterial genomes. *mSystems* **9**, e01044–24 (2024).
24. TB Brown, et al., Language models are few-shot learners. *Adv. Neural Inf. Process. Syst.* **33**, 1877–1901 (2020).
25. H Touvron, et al., Llama: Open and efficient foundation language models (arXiv preprint arXiv:2302.13971) (2023).
26. A Singh, N Pandey, A Shirgaonkar, P Manoj, V Aski, A study of optimizations for fine-tuning large language models (2024).
27. Q Fournier, et al., Protein language models: Is scaling necessary? (bioRxiv, doi:10.1101/2024.09.23.614603) (2024).
28. ESM-Team, Esm cambrian: Revealing the mysteries of proteins with unsupervised learning. (EvolutionaryScale Website) (2024) https://evolutionaryscale.ai/blog/esm-cambrian.
29. G Wang, RLD Jr., Pisces: a protein sequence culling server. *Bioinformatics* **19**, 1589–1591 (2003).
30. AJ Riesselman, JB Ingraham, DS Marks, Deep generative models of genetic variation capture the effects of mutations. *Nat. Methods* **15**, 816–822 (2018).
31. E Fenoy, AA Edera, G Stegmayer, Transfer learning in proteins: evaluating novel protein learned representations for bioinformatics tasks. *Briefings Bioinforma.* **23**, bbac232 (2022).
32. G Friedland, *The Curse of Training and the Blessing of High Dimensionality.* (Springer International Publishing, Cham), pp. 209–219 (2024).
33. JR Randall, et al., Designing and identifying β-hairpin peptide macrocycles with antibiotic potential. *Sci. Adv.* **9**, eade0008 (2023).
34. V Nallapareddy, et al., Cathe: detection of remote homologues for cath superfamilies using embeddings from protein language models. *Bioinformatics* **39**, btad029 (2023).
35. M Littmann, M Heinzinger, C Dallago, T Olenyi, B Rost, Embeddings from deep learning transfer GO annotations beyond homology. *Sci. Reports* **11**, 1160 (2021).
36. AM Freeman, BM Mole, RE Silversmith, RB Bourret, Action at a distance: amino acid substitutions that affect binding of the phosphorylated chey response regulator and catalysis of dephosphorylation can be far from the chez phosphatase active site. *J. Bacteriol.* **193**, 4709–4718 (2011).
37. BE Suzek, H Huang, P McGarvey, R Mazumder, CH Wu, Uniref: comprehensive and non-redundant uniprot reference clusters. *Bioinformatics* **23**, 1282–1288 (2007).
38. VO Pokusaeva, et al., Experimental assay of a fitness landscape on a macroevolutionary scale (bioRxiv, doi:10.1101/222778) (2018).
39. J Kaplan, et al., Scaling laws for neural language models (arXiv preprint arXiv:2001.08361) (2020).
40. A Radford, K Narasimhan, T Salimans, I Sutskever, Improving language understanding by generative pre-training (OpenAI) (2018).
41. NC Hudson, et al., Trillion parameter ai serving infrastructure for scientific discovery: A survey and vision in *Proceedings of the IEEE/ACM 10th International Conference on Big Data Computing, Applications and Technologies*. pp. 1–10 (2023).
42. R Schmirler, M Heinzinger, B Rost, Fine-tuning protein language models boosts predictions across diverse tasks. *Nat. Commun.* **15**, 7407 (2024).
43. JA Leite, DF Silva, K Bontcheva, C Scarton, Toxic language detection in social media for brazilian portuguese: New dataset and multilingual analysis (arXiv preprint arXiv:2010.04543) (2020).
44. B Zhang, Z Liu, C Cherry, O Firat, When scaling meets llm finetuning: The effect of data, model and finetuning method (arXiv preprint arXiv:2402.17193) (2024).
45. K Weiss, TM Khoshgoftaar, D Wang, A survey of transfer learning. *J. Big data* **3**, 1–40 (2016).
46. R Tibshirani, Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Method* **58**, 267–288 (1996).
47. J Hoffmann, et al., Training compute-optimal large language models (arXiv preprint arXiv:2203.15556) (2022).
48. E Nijkamp, JA Ruffolo, EN Weinstein, N Naik, A Madani, Progen2: exploring the boundaries of protein language models. *Cell systems* **14**, 968–978 (2023).
49. F Li, AP Amini, Y Yue, KK Yang, AX Lu, Feature reuse and scaling: Understanding transfer learning with protein language models (bioRxiv, doi:10.1101/2024.02.05.578959) (2024).
50. KE Wu, H Chang, J Zou, Proteinclip: enhancing protein language models with natural language (bioRxiv, doi:10.1101/2024.05.14.594226) (2024).
51. H Stärk, C Dallago, M Heinzinger, B Rost, Light attention predicts protein location from the language of life. *Bioinform. Adv.* **1**, vbab035 (2021).
52. NS Detlefsen, S Hauberg, W Boomsma, Learning meaningful representations of protein sequences. *Nat. Commun.* **13**, 1914 (2022).
53. N NaderiAlizadeh, R Singh, Aggregating residue-level protein language model embeddings with optimal transport (bioRxiv, doi:10.1101/2024.01.29.577794) (2024).
54. T Bikias, E Stamkopoulos, ST Reddy, PLMFit: Benchmarking transfer learning with protein language models for protein engineering (bioRxiv, doi: 10.1101/2025.01.15.633186) (2025).
55. P Notin, et al., ProteinGym: Large-scale benchmarks for protein fitness prediction and design in *Advances in Neural Information Processing Systems*. Vol. 36, pp. 64331–64379 (2023).
56. J Meier, et al., Language models enable zero-shot prediction of the effects of mutations on protein function in *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*. (2021).
57. F Pedregosa, et al., Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
58. J Yang, X Liao, Exploiting fine-grained dct representations for hiding image-level messages within jpeg images in *Proceedings of the 31st ACM International Conference on Multimedia*. pp. 7373–7382 (2023).
59. R Oktem, NN Ponomarenko, Image filtering based on discrete cosine transform. *Telecommun. Radio Eng.* **66**, 1685–1701 (2007).
60. D Bates, M Mächler, B Bolker, S Walker, Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* **67**, 1–48 (2015).
61. T Hothorn, F Bretz, P Westfall, Simultaneous inference in general parametric models. *Biom. Journal: J. Math. Methods Biosci.* **50**, 346–363 (2008).