
Is Transfer Learning Necessary for Protein Landscape Prediction?

Amir Shانهsazzadeh

Harvard University

Cambridge, MA 02138

amirshanehsazzadeh@college.harvard.edu

David Belanger

Google Research

dbelanger@google.com

David Dohan

Google Research

ddohan@google.com

Abstract

Recently, there has been great interest in learning how to best represent proteins, specifically with fixed-length embeddings. Deep learning has become a popular tool for protein representation learning as a model’s hidden layers produce potentially useful vector embeddings. TAPE [12] introduced a number of benchmark tasks and showed that semi-supervised learning, via pretraining language models on a large protein corpus, improved performance on downstream tasks. Two of the tasks (fluorescence prediction and stability prediction) involve learning fitness landscapes. In this paper, we show that CNN models trained solely using supervised learning both compete with and sometimes outperform the best models from TAPE that leverage expensive pretraining on large protein datasets. These CNN models are sufficiently simple and small that they can be trained using a Google Colab notebook¹. We also find for the fluorescence task that linear regression outperforms our models and the TAPE models. The benchmarking tasks proposed by TAPE are excellent measures of a model’s ability to predict protein function and should be used going forward. However, we believe it is important to add baselines from simple models to put the performance of the semi-supervised models that have been reported so far into perspective.

1 Introduction

Recent work has advocated for building protein models using semi-supervised learning [1, 2, 3, 9, 12, 13, 17]. An important question regarding these models is the effect of transfer learning: do semi-supervised models, models pretrained on a large corpus of unlabeled protein sequences, outperform supervised models? The aforementioned works all point to the affirmative. In particular, TAPE [12] benchmarks a number of language models (attentive, recurrent, convolutional, UniRep multiplicative LSTM [1], Bepler bidirectional LSTM [2]) on 5 downstream protein prediction tasks: secondary structure, contacts, remote homology, fluorescence landscape, and stability landscape. They show that pretraining a language model on a large protein corpus, specifically the Pfam database [4, 5], provides a significant improvement of the model on their proposed downstream tasks. Their downstream value prediction architecture, used for modeling the fluorescence and stability landscapes, learns fixed-length embeddings via an attention-weighted mean-pooling.

¹https://github.com/googleinterns/protein-embedding-retrieval/blob/master/cnn_protein_landscapes.ipynb

Motivated by the success of semi-supervised models, we asked if purely supervised models would be effective for modeling protein fitness landscapes. We benchmark relatively small CNN models on two of the TAPE tasks: fluorescence landscape prediction and stability landscape prediction. See Table 1 for a summary of our results and existing baselines. We find that relatively shallow CNN encoders (1-layer for fluorescence, 3-layer for stability) can compete with and even outperform the models benchmarked in TAPE. For the fluorescence task, in particular, a simple linear regression model trained on full one-hot encodings outperforms our models and the TAPE models. Additionally, 2-layer CNN models offer competitive performance with Rives et al.’s ESM (evolutionary scale modeling) transformer models [13] on β -lactamase variant activity prediction. While TAPE’s benchmarking argued that pretraining improves the performance of language models on downstream landscape prediction tasks, our results show that small supervised models can, in a fraction of the time and compute required for semi-supervised models, achieve competitive performance on the same tasks.

Table 1: Best model performances (Spearman ρ) on test sets for downstream tasks: fluorescence prediction, stability prediction, β -lactamase activity prediction. Supervised models are competitive with semi-supervised models. Note the strong performance of linear regression on one-hot encodings for the fluorescence task.

	Fluorescence	Stability	β -lactamase
Linear Regression	0.69	0.49	0.70
CNN	0.69	0.79	0.87
TAPE (Non-pretrained)	0.22	0.61	-
TAPE (Pretrained)	0.68	0.73	-
CPCProt (Best)	0.68	0.68	-
ESM (Best)	0.68	0.71	0.89
ESM (Pretrained)	-	-	0.80
ESM (Fine-tuned)	-	-	0.89

2 Background

2.1 Proteins

We consider proteins using only their primary structure, that is their amino acid sequence, with a 21-letter alphabet that includes the 20 standard amino acids [10] as well as a pad index. A length ℓ protein $a = a_1 a_2 \dots a_\ell$ is thus modeled as a discrete sequence $x = (x_1, x_2, \dots, x_\ell)$ with $x_i \in \{0, 1, \dots, 19\}$ and potentially padded to $(x_1, x_2, \dots, x_\ell, 20, 20, \dots, 20)$.

2.2 Model Embeddings

In general, a model Φ (one-hot encoding, CNN, RNN, transformer, ...) maps the encoded protein x of length ℓ to a sequence-length-dependent array representation: $\Phi(x) \in \mathbb{R}^{n \times \ell}$ for some constant n . For learning fixed-length embeddings, the goal is to learn a mapping Ψ such that $z = \Psi(\Phi(x)) \in \mathbb{R}^m$ with m being independent of ℓ . We also want z to be a “useful” protein embedding, in the sense of having strong signal for some downstream feature(s).

Our approach to computing and learning embeddings from a model Φ is to apply a potentially learnable mapping from the encoded representations, which are arrays of the shape (n, ℓ) , to fixed-length arrays of the shape $(m,)$. The simplest mapping we consider is pooling, which is non-learnable and involves either taking the average (MeanPool) or the maximum (MaxPool) over the length-dependent axis of $\Phi(x) \in \mathbb{R}^{n \times \ell}$. A learnable, but simple, mapping involves a single dense layer with weight matrix $W \in \mathbb{R}^{m \times n}$ and bias vector $b \in \mathbb{R}^m$ and a non-linear activation ϕ . Then we consider the array of transformed amino acid level embeddings

$$\left[\phi(W\Phi(x)^{(i)} + b) \right]_{i=1}^{\ell} \in \mathbb{R}^{m \times \ell}$$

and apply a pooling operation. If we apply MaxPool we call this operation LinearMaxPool and likewise for MeanPool. Note that we only use the ReLU activation for ϕ . At a high-level our approach is motivated by a similar technique from NLP known as contextual lenses [8].

2.3 Protein Landscape Prediction Tasks

Two of the TAPE tasks, fluorescence prediction and stability prediction, are scalar value prediction problems [12]. Their prediction model involves taking the sequence-length-dependent array representation outputted by a language model, applying an attention-weighted mean-pooling operation to get a sequence-length-independent vector representation, applying a dense hidden layer of size 512 followed by ReLU activation, and finally performing value prediction using a dense output layer. The task-specific learning produces fixed-length embeddings designed to be classified by scalar value (fluorescence or stability) using a small MLP trained with mean squared error (MSE) loss.

We also consider a variant prediction task from Rives et al [13]. The publicly available implementation² uses precomputed pretrained transformer protein embeddings and their paper includes the results from a transformer model that is both pretrained and fine-tuned.

It is worth noting that the tasks we consider are fine-grained landscape prediction, where we attempt to predict the output of a real-valued functional assay (e.g. brightness). Works such as Biswas et. al. [3] consider more coarse-grained tasks where the goal is to predict the presence or lack of function.

We give a brief overview of the prediction tasks and the datasets used. For more detail, particularly from the biological perspective, refer to [6, 14, 15].

2.3.1 Fluorescence Landscape Prediction

This regression task involves mapping a protein to its log-fluorescence, which is a real-valued label. The experimental data is from Sarkisyan et al. [15] and the curated dataset is from TAPE³. The data consists of mutated variants of a wild-type GFP protein with edit distance up to 14. The train set contains all variants within edit distance 3 of the wildtype (at most 3 mutations away) and the test set contains all variants at least 4 mutations away from the wildtype. This split by edit distance allows for testing the generalizability of a model trained on a small (local) neighborhood of the wildtype to a larger (global) neighborhood. Note that because of this split the train set consists of 82% bright proteins (log-fluorescence greater than 2.5) and 18% dark proteins (log-fluorescence less than 2.5) whereas the test set consists of only 32% bright proteins and 68% dark proteins. This class imbalance makes it difficult for models to generalize from the low mutation train data to the high mutation test data, as evidenced by the TAPE results and our results. The primary metric of interest is Spearman’s rank correlation coefficient ρ on the test set. MSE on the test set is also reported. Additionally, due to the class imbalance ρ and MSE on the bright and dark subsets of the test set are reported.

2.3.2 Stability Landscape Prediction

This regression task involves mapping a protein to its stability, a real-valued label, which is a measurement of the log-concentration of protease required for $1/2$ of the cells cultured with said protein to satisfy a collection threshold, normalized against the corresponding predicted log-concentration. The experimental data is from Rocklin et al. [14] and the curated dataset is from TAPE. The train set contains proteins from experimental design rounds whereas the test set contains proteins 1 mutation away from top candidates. This task thus tests for the ability of a model to learn local properties from global information, the opposite of the fluorescence task. The primary metric of interest is again Spearman’s rank correlation coefficient ρ on the test set. Additionally, an accuracy metric is reported. This accuracy measures the degree to which binarized model predictions agree with binarized true values. Specifically, a prediction is said to be accurate in 2 of 4 cases:

- (i) The variant protein stability is greater than the parent protein stability and the predicted variant stability is greater than the predicted parent stability.
- (ii) The variant protein stability is less than the parent protein stability and the predicted variant stability is less than the predicted parent stability.

²https://github.com/facebookresearch/esm/blob/master/examples/variant_prediction.ipynb

³<https://github.com/songlab-cal/tape>

The proteins primarily consist of 4 fold topologies: $\{\alpha\alpha\alpha, \alpha\beta\beta\alpha, \beta\alpha\beta\beta, \beta\beta\alpha\beta\beta\}$ and so ρ and the accuracy metric are reported on the full test set as well as the test set restricted to each of these topologies.

2.3.3 β -lactamase Variant Prediction

The data for this task comes from Envision (Gray et. al.) [6] and is curated by Rives et al. It consists of mutated β -lactamase sequences as inputs and scaled mutation effects as targets. We report mean and standard deviation of Spearman’s ρ on a partition of size 5 of the generated test set. The train set is also varied and generated by random splits of 1%, 10%, 30%, 50%, and 80% of the data. FAIR’s publicly available implementation measures the performance of pretrained, but not fine-tuned, transformer embeddings with an 80-20 data split. They measure performance on the whole test set. To compare with their full paper we adapt their implementation by using train sets ranging from 1% to 80% of the data and create a partition of size 5 of the test set.

3 Related Work

The most directly related work is TAPE [12]. which curates the datasets for and formulates the specific fluorescence and stability landscape prediction problems we consider. The experimental data for the fluorescence task is from Sarkisyan et al. [15]. They model the GFP landscape using linear regression on single mutation effects, multiple regression which is similar to their linear regression but models fluorescence as a sigmoidal function of mutations, and MLPs where the inputs are binarized according to the presence of mutations in the protein genotype. The experimental data for the stability task comes from Rocklin et al. [14]. Their library design is computational, but uses energy-based and molecular dynamics approaches as opposed to primary sequence modeling.

Lu et al. [9] develop CPCProt, an autoregressive model pretrained on the same protein corpus as TAPE using mutual information maximization. Their model is substantially smaller (~ 1.7 M parameters) than the TAPE models, which range from ~ 19 M-182M parameters, but its self-supervised embeddings provide comparable downstream performance using TAPE’s attention-weighted mean architecture.

Rives et al. [13] does similar work but trains varying size transformer language models on 250 million UniRef [16] sequences. Their transformer embeddings achieve state-of-the-art performance on a number of downstream tasks such as contact, secondary structure, and remote homology. They also show that larger model size, which results in better protein language modeling performance, improves downstream performance. We consider their variant prediction task, specifically for β -lactamase data from Envision [6].

We are unaware of other directly comparable works that use the same datasets, but machine learning techniques have been used in a similar fashion to learn useful protein embeddings in [1, 2, 13, 17, 17, 18].

4 Models and Training Procedure

We use linear regression, specifically scikit-learn’s [11] ridge regression, on the full one-hot protein encoding as a baseline for all tasks. It is important to note that the one-hot baseline in TAPE [12] does not use the full one-hot encoding but instead uses a bag-of-words representation (MeanPool applied to the full one-hot encoding), which measures the frequency of amino acids. We refer to this as One-hot (AA counts). This pooling results in a substantial loss of information, especially for the fluorescence task which is trained on proteins that are very close in edit distance.

Our supervised models only rely on 1-D convolution layers, dense layers, and ReLU activations. Training is done using the Adam optimizer [7] with variable learning rates and weight decays per architecture component, but with no learning rate warmup. Mean squared error loss is used. Sequences are padded to maximum length and the corresponding padded components are zeroed out before pooling.

4.1 Fluorescence Models

For modeling the fluorescence landscape we use a 1-layer 1-D CNN with feature size 1024 and kernel size 5 as an encoder. ReLU is applied after the convolution layer. To create embeddings we use either MaxPool (embedding dimension 1024) or LinearMaxPool with an embedding dimension of 2048. For the predictor a dense layer is used. Our fluorescence CNN models have $\sim 110\text{k}$ - 220k parameters, which is ~ 1 order of magnitude less than CPCProt [9] and ~ 2 - 3 orders of magnitude less than the TAPE models. We also use a linear regression model with $\sim 5\text{k}$ parameters as a baseline. We report the hyperparameters in Table 2. Note that Ep is epochs, BS is batch size, Dim is embedding dimension, and the learning rates and weight decays are reported as tuples with the values corresponding to the encoder, mapping, and predictor in that order.

4.2 Stability Models

For modeling the stability landscape we use a 3-layer 1-D CNN with feature sizes all 1024 and kernel sizes all 5 as an encoder. ReLU is applied between and after convolution layers. To create embeddings we use either MaxPool (embedding dimension 1024) or LinearMaxPool with an embedding dimension of 2048. We also try a Dilated CNN + MaxPool model with kernel dilation of 2 in the 2nd layer. For the predictor a dense layer is used. Our stability CNN models have $\sim 10.6\text{M}$ - 12.7M parameters, which is ~ 1 order of magnitude greater than CPCProt and ~ 0 - 1 orders of magnitude less than the TAPE models. Additionally, we ensemble the above three models by averaging their predictions to create an “Ensemble of CNNs” model. The motivation behind this averaging is the variance in individual model performance across different fold topologies. We also use a linear regression model with $\sim 1\text{k}$ parameters as a baseline. We report the hyperparameters in Table 2.

4.3 β -lactamase Variant Models

For modeling the β -lactamase variant landscape we use a 2-layer 1-D CNN with feature sizes 1024 and 512 (layer 1 and layer 2) and kernel sizes 9 and 7 (layer 1 and layer 2) as an encoder. ReLU is applied between and after convolution layers. To create embeddings we use MaxPool (embedding dimension 512). We also try a Dilated CNN + MaxPool model with kernel dilation of 2 in the 1st layer (we at times refer to this as “D 2-Layer + MaxPool”). For the predictor a dense layer is used. Our variant CNN models have $\sim 38.7\text{M}$ parameters. Additionally, we ensemble the above two models by averaging their predictions to create an Ensemble of CNNs model. We also use a linear regression model with $\sim 6\text{k}$ parameters as a baseline. We report the hyperparameters in Table 2.

Table 2: Model hyperparameters

Task	Model	Ep	BS	Dim	Learning Rates	Weight Decays
GFP	1-Layer CNN + MaxPool	50	256	1024	(1e-3, 0, 5e-6)	(0, 0, 0.05)
	1-Layer CNN + LinearMaxPool	50	256	2048	(1e-3, 5e-5, 5e-6)	(0, 0.05, 0.05)
Stability	3-Layer CNN + MaxPool	5	16	1024	(5e-4, 0, 5e-5)	(0.025, 0, 0.025)
	3-Layer Dilated CNN + MaxPool	5	16	1024	(5e-4, 0, 5e-5)	(0.025, 0, 0.025)
	3-Layer CNN + LinearMaxPool	10	256	2048	(1e-5, 5e-5, 5e-6)	(0, 0, 0)
Variant	2-Layer CNN + MaxPool	500	32	512	(1e-3, 0, 1e-3)	(0.1, 0, 0.05)
	2-Layer Dilated CNN + MaxPool	500	32	512	(1e-3, 0, 1e-3)	(0.1, 0, 0.05)

5 Results

For the TAPE landscape prediction tasks, we report results in a similar fashion to TAPE [12] by recreating an analog of their tables with our results and the best of their results over all models. We also include the results provided in CPCProt [9]. For the β -lactamase prediction task, we recreate an analog of Table S7 in Rives et al. [13].

5.1 Fluorescence Results

Table 3 contains mean squared error (MSE) and Spearman’s ρ for the full GFP test set as well as the bright subset (log-fluorescence greater than 2.5) and dark subset (log-fluorescence less than 2.5). For the TAPE results we present the best reported metric across a class of models. Note the significant difference between our one-hot linear regression baseline and the one-hot baseline of TAPE. This is because we learn a regressor on the full one-hot vector whereas TAPE uses a bag-of-words representation, which is the full one-hot vector with MeanPool applied. For the CPCProt results, we provide the best reported metric.

Table 3: Fluorescence prediction results. ρ is Spearman’s rank correlation coefficient.

		Full Test Set		Bright Mode Only		Dark Mode Only	
		MSE	ρ	MSE	ρ	MSE	ρ
Baseline	Linear Regression	0.35	0.69	0.09	0.68	0.33	0.05
CNN	1-Layer + MaxPool	0.26	0.69	0.09	0.65	0.29	0.05
	1-Layer + LinearMaxPool	0.23	0.69	0.12	0.66	0.28	0.05
TAPE	One-hot (AA counts)	2.69	0.14	0.08	0.03	3.95	0.0
	Non-pretrained models	2.35	0.22	0.07	0.08	3.43	0.0
	Pretrained models	0.19	0.68	0.08	0.63	0.22	0.05
	Best of all models	0.19	0.68	0.07	0.63	0.22	0.05
CPCProt	Best of all models	0.81	0.68	-	-	-	-
ESM	Best of all models	-	0.68	-	-	-	-

Figure 1 shows the top 2 principal components of the model embeddings colored by log-fluorescence. We separate the sequences by the train and test split and fit the PCA only to the train set. For models we use a trained CNN + MaxPool model and an identical, but randomly initialized, CNN + MaxPool model. Notice that the trained model separates sequences by log-fluorescence whereas the randomly initialized model does not.

We see from Table 3 that the CNN architectures outperform the TAPE models and CPCProt [9] in the more relevant Spearman ρ metric. As shown in Figure 1, the 1-layer CNN models are able to learn fixed-length embeddings that separate the proteins based on log-fluorescence. It is worth noting that due to the bimodality of the data, MSE is not a particularly good metric. For example, the TAPE model that achieves the lowest bright mode MSE of 0.07 has predictions that are negatively rank-correlated with the true fluorescence values. Furthermore, a simple linear regression model outperforms every other reported model.

5.2 Stability Results

Table 4 shows overall performance on the stability test set, and Table 5 shows performance on specific fold topologies. In addition to reporting Spearman’s ρ , we report the previously defined accuracy measurement that uses the parent protein as a decision boundary and labels mutations as beneficial if predicted stability is greater than predicted parent stability and deleterious if the opposite is true. We again present the best reported metric across a class of models for the TAPE results and the best reported metric for CPCProt.

Figure 2 shows plots of the predicted stabilities from the linear regression model and from the Ensemble of CNNs.

We see from Table 4 that the CNN ensemble model outperforms the TAPE models and CPCProt on the full dataset. Table 5 shows that our 3-layer CNN models are all competitive with the TAPE models and in fact outperform on 3 out of 8 metrics. In particular, our models achieve better binarized accuracies indicating that they are able to effectively predict whether or not a mutation will be beneficial or deleterious. This is perhaps more relevant than rank-correlation, since in general for protein engineering the aim is to predict whether or not a mutation will be beneficial or deleterious

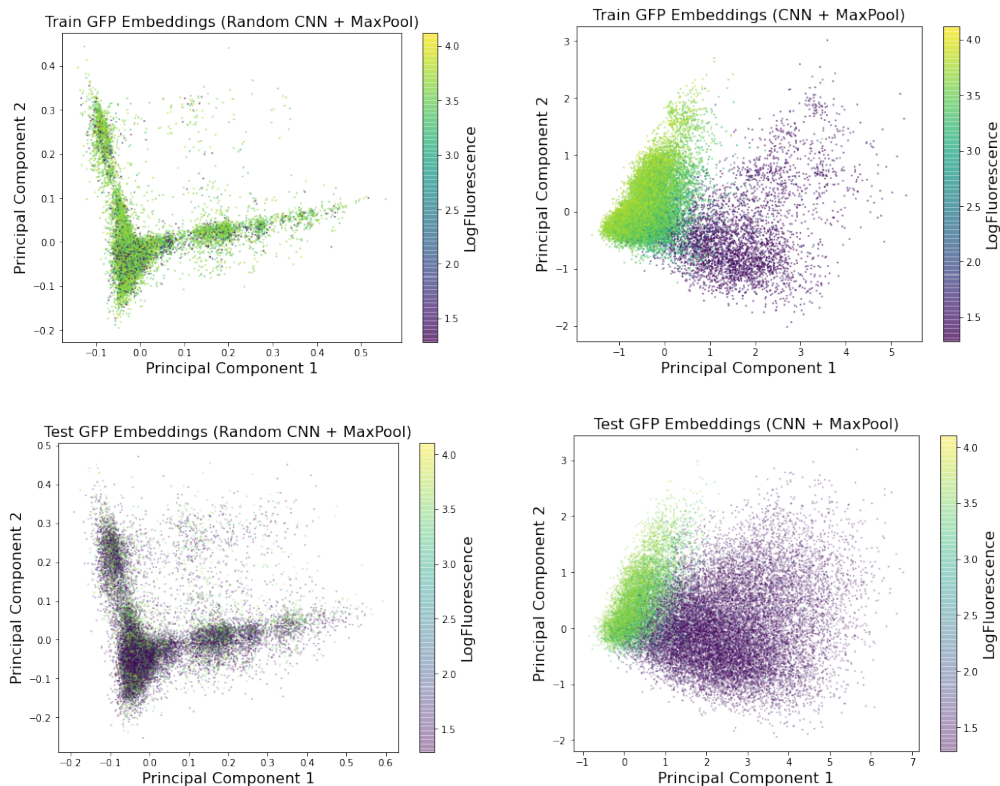


Figure 1: PCA of GFP train and test set embeddings colored by log-fluorescence for randomly initialized CNN model and trained CNN model. Supervised training produces embeddings that separate both the train and test sets according to log-fluorescence.

Table 4: Overall stability prediction results

		Spearman's ρ	Accuracy
Baseline	Linear Regression	0.49	0.60
	CNN		
	3-Layer + MaxPool	0.76	0.75
	Dilated 3-Layer + MaxPool	0.75	0.73
	3-Layer + LinearMaxPool	0.71	0.77
	Ensemble of CNNs	0.79	0.77
TAPE	One-hot (AA counts)	0.19	0.58
	Non-pretrained models	0.61	0.68
	Pretrained models	0.73	0.70
	Best of all models	0.73	0.70
CPCProt	Best of all models	0.68	-
ESM	Best of all models	0.71	-

as opposed to *how* beneficial or deleterious said mutation will be. Interestingly, for this task, linear regression performs quite poorly. This is likely because the task requires the model to learn from broader, more global data and localize information which is the opposite of what the fluorescence task requires.

Table 5: Stability prediction results broken down by protein topology

		$\alpha\alpha\alpha$		$\alpha\beta\beta\alpha$		$\beta\alpha\beta\beta$		$\beta\beta\alpha\beta\beta$	
		ρ	Acc	ρ	Acc	ρ	Acc	ρ	Acc
Baseline	Linear Regression	0.21	0.66	-0.03	0.6	0.51	0.64	0.38	0.61
	3-Layer + MaxPool	0.69	0.71	0.37	0.70	0.50	0.72	0.60	0.68
	Dilated 3-Layer + MaxPool	0.67	0.69	0.49	0.69	0.61	0.70	0.53	0.64
	3-Layer + LinearMaxPool	0.59	0.69	0.52	0.77	0.55	0.73	0.6	0.70
	Ensemble of CNNs	0.67	0.71	0.53	0.75	0.65	0.74	0.60	0.70
TAPE	One-hot (AA counts)	0.58	0.59	0.04	0.58	-0.05	0.58	0.54	0.58
	Non-pretrained models	0.64	0.69	0.39	0.70	0.63	0.67	0.65	0.67
	Pretrained models	0.72	0.70	0.48	0.79	0.68	0.71	0.67	0.70
	Best of all models	0.72	0.70	0.48	0.79	0.68	0.71	0.67	0.70

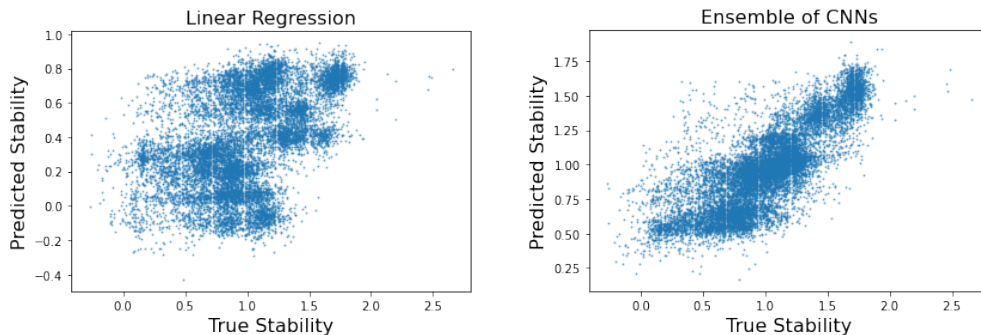


Figure 2: Predicted vs. true stabilities for linear regression model and ensemble of CNNs.

5.3 β -lactamase Variant Results

Table 6 compares the performance of our CNN models with the ESM transformer models [13]. We report mean and standard deviation of Spearman ρ across a size 5 partition of the test sets and present results for the fine-tuned ESM model from the corresponding paper.

Table 6: Comparison across models of mean and standard deviation β -lactamase variant prediction performance, measured by Spearman ρ , on a partition of size 5 of the test set.

		1% data	10% data	30% data	50% data	80% data	
Baseline	Linear Regression	0.19 \pm 0.10	0.52 \pm 0.01	0.67 \pm 0.02	0.70 \pm 0.02	0.70 \pm 0.03	
	2-Layer + MaxPool	0.24 \pm 0.03	0.59 \pm 0.02	0.77 \pm 0.01	0.83 \pm 0.02	0.86 \pm 0.01	
	CNN	D 2-Layer + MaxPool	0.15 \pm 0.03	0.62 \pm 0.02	0.76 \pm 0.02	0.83 \pm 0.02	0.86 \pm 0.01
		Ensemble of CNNs	0.23 \pm 0.03	0.62 \pm 0.02	0.78 \pm 0.01	0.84 \pm 0.02	0.87 \pm 0.01
ESM	Pretrained	0.51 \pm 0.03	0.68 \pm 0.02	0.75 \pm 0.01	0.78 \pm 0.01	0.80 \pm 0.02	
	Fine-tuned	0.39 \pm 0.03	0.69 \pm 0.01	0.84 \pm 0.01	0.88 \pm 0.01	0.89 \pm 0.01	

Figure 3 shows a 2-D t-SNE of the model embeddings colored by activity. We separate the sequences by the 80-20 train and test split. For models, we use a trained ensemble of CNN + MaxPool models and an identical, but randomly initialized, ensemble of CNN + MaxPool models. Again, notice that the trained ensemble of models separates sequences by activity whereas the randomly initialized ensemble does not.

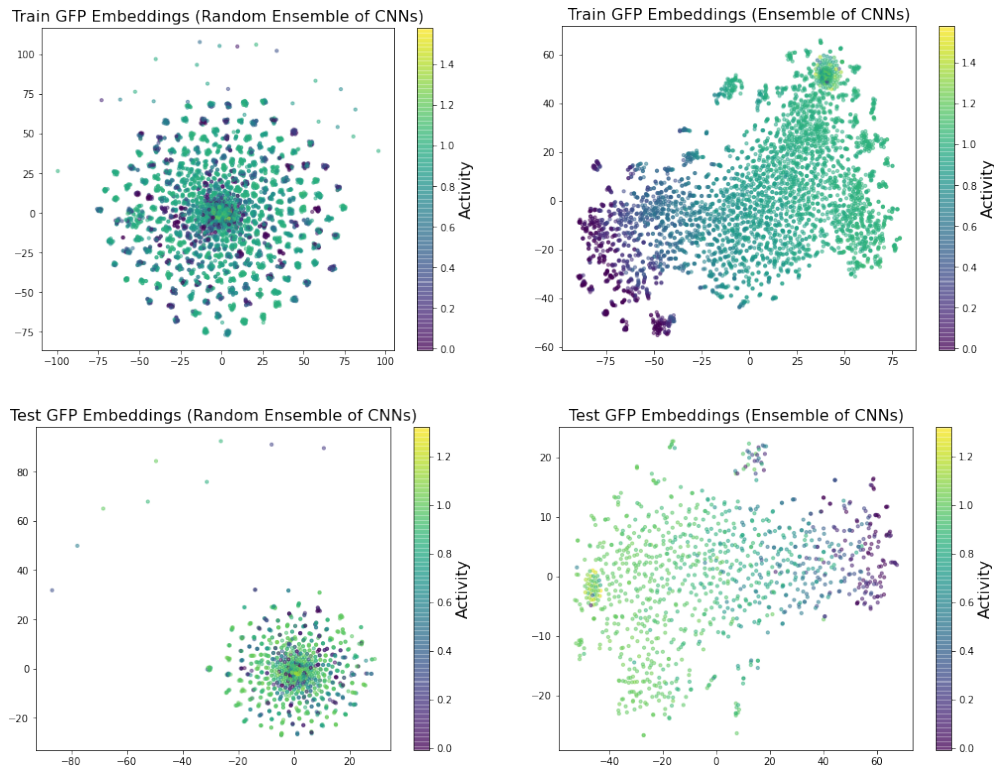


Figure 3: t-SNE of β -lactamase train and test set embeddings colored by activity for randomly initialized ensemble of CNN models and trained ensemble of CNN models. Supervised training produces embeddings that separate both the train and test sets according to activity.

We see in Table 6 that a 2-layer CNN with an 80-20 train test split outperforms the pretrained ESM transformer [13] and is competitive with the pretrained and fine-tuned ESM transformer. The t-SNE plots in Figure 3 shows that the Ensemble of CNNs embeddings, when trained, separate the variants by activity. All deep models outperform linear regression by a noticeable margin. It is worth noting, however, that with lower percentages of the data used for training, the ESM models considerably outperform our CNN models. Pretraining the ESM transformer resulted in embeddings that split the data modestly by activity without any downstream supervision, which we believe explains this advantage with less data.

6 Discussion

We see that relatively simple and small CNN models trained entirely with supervised learning for fluorescence or stability prediction compete with and outperform the semi-supervised models benchmarked in TAPE [12], despite requiring substantially less time and compute. While the benefit of transfer learning via pretraining is evident, the performance of the pretrained language models does not appear to justify the cost of pretraining for these protein landscape modeling tasks. Furthermore, it appears that with a sufficient amount of data, CNNs trained only with supervised learning can compete with the pretrained (and fine-tuned) transformers from Rives et. al. [13] on β -lactamase variant prediction. However, the semi-supervised models outperform considerably on this task when training data is limited. This leads us to conjecture that pretraining may be most useful for downstream tasks with limited supervised data.

References

- [1] 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, October 2019.
- [2] Tristan Bepler and Bonnie Berger. Learning protein sequence embeddings using information from structure. *International Conference on Learning Representations, 2019*, 2019.
 - [3] Surojit Biswas, Grigory Khimulya, Ethan C. Alley, Kevin M. Esvelt, and George M. Church. Low-n protein engineering with data-efficient deep learning. *bioRxiv preprint*, January 2020.
 - [4] Sara El-Gebali, Jaina Mistry, Alex Bateman, Sean R Eddy, Aurélien Luciani, Simon C Potter, Matloob Qureshi, Lorna J Richardson, Gustavo A Salazar, Alfredo Smart, Erik L L Sonnhammer, Layla Hirsh, Lisanna Paladin, Damiano Piovesan, Silvio C E Tosatto, and Robert D Finn. The pfam protein families database in 2019. *Nucleic Acids Research*, 47(D1):D427–D432, October 2018.
 - [5] Robert D. Finn, Alex Bateman, Jody Clements, Penelope Coggill, Ruth Y. Eberhardt, Sean R. Eddy, Andreas Heger, Kirstie Hetherington, Liisa Holm, Jaina Mistry, Erik L. L. Sonnhammer, John Tate, and Marco Punta. Pfam: the protein families database. *Nucleic Acids Research*, 42(D1):D222–D230, November 2013.
 - [6] Vanessa E. Gray, Ronald J. Hause, Jens Luebeck, Jay Shendure, and Douglas M. Fowler. Quantitative missense variant effect prediction using large-scale mutagenesis data. *Cell Systems*, 6(1):116–124.e3, January 2018.
 - [7] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014.
 - [8] Jamie Kiros. Contextual lensing of universal sentence representations, 2020.
 - [9] Amy X. Lu, Haoran Zhang, Marzyeh Ghassemi, and Alan Moses. Self-supervised contrastive learning of protein representations by mutual information maximization. *bioRxiv preprint*, September 2020.
 - [10] International Union of Pure and Applied Chemistry. Nomenclature and symbolism for amino acids and peptides (recommendations 1983). *Pure and Applied Chemistry*, 56(5):595–624, January 1984.
 - [11] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
 - [12] Roshan Rao, Nicholas Bhattacharya, Neil Thomas, Yan Duan, Xi Chen, John Canny, Pieter Abbeel, and Yun S Song. Evaluating protein transfer learning with tape. In *Advances in Neural Information Processing Systems*, 2019.
 - [13] Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, 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. *bioRxiv*, 2019.
 - [14] Gabriel J Rocklin, Tamuka M Chidyausiku, Inna Goreschnik, Alex Ford, Scott Houliston, Alexander Lemak, Lauren Carter, Rashmi Ravichandran, Vikram K Mulligan, Aaron Chevalier, et al. Global analysis of protein folding using massively parallel design, synthesis, and testing. *Science*, 357(6347):168–175, 2017.
 - [15] Karen S Sarkisyan, Dmitry A Bolotin, Margarita V Meer, Dinara R Usmanova, Alexander S Mishin, George V Sharonov, Dmitry N Ivankov, Nina G Bozhanova, Mikhail S Baranov, Onuralp Soylemez, et al. Local fitness landscape of the green fluorescent protein. *Nature*, 533(7603):397, 2016.
 - [16] B. E. Suzek, Y. Wang, H. Huang, P. B. McGarvey, and C. H. Wu and. UniRef clusters: a comprehensive and scalable alternative for improving sequence similarity searches. *Bioinformatics*, 31(6):926–932, November 2014.
 - [17] Jesse Vig, Ali Madani, Lav R. Varshney, Caiming Xiong, Richard Socher, and Nazneen Fatema Rajani. Bertology meets biology: Interpreting attention in protein language models, 2020.
 - [18] Kevin K Yang, Zachary Wu, Claire N Bedbrook, and Frances H Arnold. Learned protein embeddings for machine learning. *Bioinformatics*, 34(15):2642–2648, March 2018.