

**This item is the archived peer-reviewed author-version of:**

PEvoLM : Protein Sequence Evolutionary Information Language Model

**Reference:**

Arab Issar.- PEvoLM : Protein Sequence Evolutionary Information Language Model  
2023 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology (CIBCB), 29-31 August 2023, Eindhoven, Netherlands - ISBN 979-83-503-1018-4 - IEEE Xplore, (2023), p. 1-8  
Full text (Publisher's DOI): <https://doi.org/10.1109/CIBCB56990.2023.10264890>  
To cite this reference: <https://hdl.handle.net/10067/2000470151162165141>

# PEvoLM: Protein Sequence Evolutionary Information Language Model

Issar Arab

Department of Computer Science, University of Antwerp, 2020 Antwerp, Belgium  
Biomedical Informatics Network Antwerpen (biomina), 2020 Antwerp, Belgium

**Abstract**— With the exponential increase of the protein sequence databases overtime, multiple-sequence alignment (MSA) methods, like PSI-BLAST, perform exhaustive and time-consuming database search to retrieve evolutionary information. The resulting position-specific scoring matrices (PSSMs) of such search engines represent a crucial input to many machine learning (ML) models in the field of bioinformatics and computational biology. A protein sequence is a collection of contiguous tokens or characters called amino acids (AAs). The analogy to natural language allowed us to exploit the recent advancements in the field of Natural Language Processing (NLP) and therefore transfer NLP state-of-the-art algorithms to bioinformatics. This research presents an Embedding Language Model (ELMo), converting a protein sequence to a numerical vector representation. While the original ELMo trained a 2-layer bidirectional Long Short-Term Memory (LSTMs) network following a two-path architecture, one for the forward and the second for the backward pass, by merging the idea of PSSMs with the concept of transfer-learning, this work introduces a novel bidirectional language model (bi-LM) with four times less free parameters and using rather a single path for both passes. The model was trained not only on predicting the next AA but also on the probability distribution of the next AA derived from similar, yet different sequences as summarized in a PSSM, simultaneously (multi-task learning), hence learning evolutionary information of protein sequences as well. The network architecture and the pre-trained model are made available as open source under the permissive MIT license on GitHub at <https://github.com/issararab/PEvoLM>.

**Keywords**— *Protein Sequences, Evolutionary Information, PSSM, Variational Inference, Deep Learning, Multi-task learning, ELMo*

## I. INTRODUCTION

In contemporary computational biology and, more specifically, in the area of proteins and the prediction of their properties, sequence alignment forms the de-facto standard input to nearly all machine learning methods [1]. Multiple sequence alignment methods, or MSAMs, are a set of algorithmic solutions for the alignment of evolutionarily related sequences. They can be applied to DNA, RNA, or protein sequences. Those algorithms are designed to take into account evolutionary events such as mutations, insertions, deletions, and rearrangements under certain conditions [2]. The main functionality of these alignment techniques is to search for homologs of a query sequence in a database of protein sequences, as they tend to share structure and function. For the past two decades, training machine learning models with evolutionary information representations, generated by multiple sequence alignments, has revolutionized the prediction power of AI methods. Multiple aspects of protein function and structure were studied and investigated following the same approach and achieved significant results in the prediction performance. Such downstream-specific tasks

include protein secondary structure [3, 4, 5, & 6], transmembrane protein regions [7, 8, & 9], inter-residue contacts [10], and sub-cellular localization predictions [11, 12] as well as protein to protein interactions [13, 14, & 15]. However, this increase in performance is becoming costly in recent years, with the continuous exponential growth of bio-sequence data pools. UniProt is one example of such data stores, in which the entries keep doubling every couple of years [16]. To cope with such tremendous growth, alternative approaches are researched among the community. One prominent solution that can compete with conventional methods is the direction of Embedding Language Models (ELMo) [17], a state-of-the-art technology borrowed from the NLP field.

In the NLP setting, pre-trained word representations are a central component to several natural language comprehension models [18, 19]. However, learning high-quality representations is a difficult task. Ideally, these representations have to model both the dynamic features of word usage, like semantics and morphology, and how they differ across linguistic domains, like polysemy modeling. Pre-trained word vectors [18, 19, & 20] learned from a large corpus of unlabelled content have the ability to model these syntactic and semantic word representations. They represent the core of many state-of-the-art NLP architectures out there, such as semantic role labeling [21], question answering [22], and textual entailment [23].

An ELMo is trained on a large corpus of unlabelled natural text, Wikipedia as an example [2], to predict the next most probable word in a sentence given all the previously seen tokens. However, in a bi-directional language model, during training we learn the probability distribution of the next word in the sentence from both directions, i.e. predicting a pivot word given all the previous tokens from a forward pass and from a backward pass of a sentence. This bi-directional autoregressive [17, 24] paradigm has revolutionized NLP allowing the model to develop a syntactic and a semantic self-learning of the word in a sentence, a.k.a. the context. This means that, for a particular word, the model will provide different contextualized embeddings, depending on the sentence it is used in.

Given the close nature of protein sequences to natural language sentences, the same approach was adopted to train SeqVec [25] on UniRef50, a corpus of 9.5 billion amino acids, which is around 10 folds larger than Wikipedia in terms of tokens(words). In their research work, Heinzinger et al. [25] proposed a novel embedding tool of protein sequences that replaces the explicit search for evolutionary-related proteins in a database. The model was trained on predicting the next amino acid in the sequence. The new approach can be described as an implicit transfer of biophysical and biochemical information learned during the training of a bi-

directional language model embedder [17], inspired from NLP, on a large unlabelled set of sequences.

The predictive power of the embeddings was then tested on downstream tasks categorized under two levels: per-residue and per-protein predictions. The results showed that the models were able to reach a good performance, but did not outperform the state-of-the-art MSA-based tools. This paper’s idea was then to train a novel bi-language model on PSI-BLAST’s output with transfer learning, which would eventually encode evolutionary information of the proteins within its embedding representations, with the goal of boosting the final embedding power and potentially reduce the size of residue embedding while maintaining the amount of information encoded.

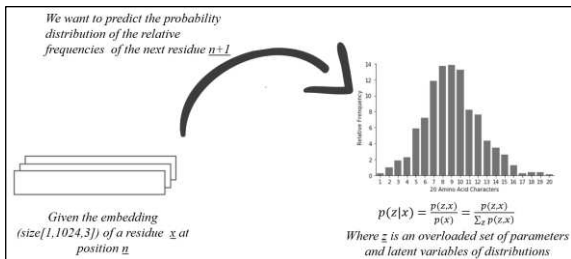
To train the new embedder, a large curated dataset of sequences was compiled with their corresponding PSSMs of size 1.83 Million proteins (~0.8 billion amino acids). The dataset of proteins is reduced to 40% sequence identity, with respect to the validation/test sets, and contains sequences ranging between 18 and 9858 residues in length. The next sections will include the research question, data and methods, followed by the experiments, results and discussion section.

All results presented in this paper were conducted on a remote Linux VM granted by Google with a system memory size (RAM) of 120GB and 32 Intel® Xeon® CPUs with a maximum speed of 2.3 GHz. The machine also contained a cluster of 8 Tesla V100 SXM2 GPUs with a dedicated memory of 16GB each, from which 2 GPUs were used to train the novel embedding language model.

## II. RESEARCH QUESTION

In probabilistic machine learning, a probabilistic model is a joint distribution of hidden variables, referred to as  $\underline{z}$ , and observed variables, referred to as  $\underline{x}$ . This statement can be written in probabilistic notations as  $p(x, z)$ .

In this setting, inference about the unknowns  $\underline{z}$  is done through what we call the *posterior* distribution. It is a conditional distribution of the hidden variables given the observations. This statement is translated in probabilistic notations as  $p(z | x)$ .



**Figure 1.** Problem statement visualization of the research goal on learning representations of protein evolutionary information

The goal of this research work is to go directly from the amino acid sequence of a given protein to a Position-Specific Scoring Matrix, more specifically to the matrix of relative frequencies. In other words, we want to translate a vector of characters of length  $L$ , here the protein sequence, to a matrix of dimensions  $L \times 20$ , which is the Position-Specific Scoring Matrix in our context. In this PSSM map, each row represents the relative frequencies of a residue in the sequence with

respect to all the 20 known AAs. These row values represent a discrete probability distribution, where the statistics in each row sum up to 1 or 100 depending on the scale. Our approach aims at learning those representations directly from the sequence following the same paradigm applied in autoregressive models [26]: From the observations of previous time steps given as an input, we want to make predictions of the next time step.

Figure 1 visualizes the main problem statement in this research work, which aims to learn protein evolutionary information. In a protein sequence of length  $L$ , given the embedding of a residue  $\underline{x}$  at position  $\underline{n}$  as well as the previous residues information passed on, we want to predict the probability distribution of the relative frequencies of the next residue  $\underline{n+1}$ . In mathematical notations, this is equivalent to a conditional probability distribution  $p(z | x)$ , where  $\underline{z}$  is an overloaded set of parameters and latent variables of the distribution. Since the evolutionary information embedded within the relative frequencies matrix at the residue level is a discrete probability distribution and the evidence can be written as the marginalization of the joint distribution over  $\underline{z}$ , our probabilistic model can be reformulated as follows:

$$p(z|x) = \frac{p(z,x)}{\sum_z p(z,x)}$$

To solve this model and since the evidence is not tractable, we appeal to approximating our posterior inference through Variational Inference (VI) [27] using Kullback-Leibler divergence (KLL-Div). The main idea is to: (1) find a tractable distribution  $q(z|v)$  that is similar to  $p(z|x)$ , and then (2) use  $q(z|v)$  to answer the questions about  $p(z|x)$  that we care about. Here, we just have to find the optimal parameters  $v^*$  that minimize the KLL-divergence.

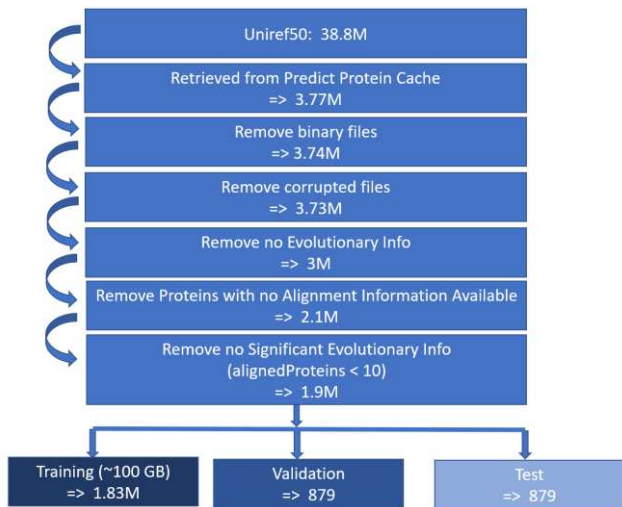
Following the universal approximation theorem [28], we opted for a deep neural network as a parametric complex function that can learn any distribution given enough neurons in a two-layer network. Therefore, we converted our inference to an optimization problem. The final loss function that we need to minimize is:

$$v^* = \underset{v}{\operatorname{argmin}} \left[ \sum_x p(z|x) \log(p(z|x)) - \sum_x p(z|x) \log(q(z|v)) \right]$$

## III. DATA AND METHODS


Running PSI-BLAST on a large set of protein sequences from UniProt [29], for example Reference Cluster with 50% sequence identity, was unrealistic as the search is highly exhaustive and might take months, depending on the available computing resources, for a couple of millions of protein sequences. Therefore, I opted for collecting the cached PSSMs from Predict-Protein [30], an Internet service for sequence analysis as well as prediction of protein structure and function. Figure 2 depicts the high-level integration and transformation steps applied in the pipeline used for data gathering and pre-processing. Protein sequences from UniProt [29] Reference Cluster with 50% sequence identity (uniref50 2019\_12), ~38.8 Million proteins, were checked one by one, and for the matched hits we retrieved PSSMs with evolutionary information along with their corresponding alignment files. All data was cleaned and pre-processed for high-quality

PSSMs to use during training. The final set of 1.83 Million protein sequences (~0.8 billion amino acids) and their corresponding PSSMs, along with validation, and test sets have been made public to the community and deposited to Zenodo at (<https://zenodo.org/record/4300971>).




**Figure 2.** Simplified visualization of the gathering and pre-processing pipeline of the whole Uniref50 dataset to obtain the final training validation and test sets

To analyze our development set (~2.1M sequences) in terms of the amount of evolutionary information encoded in the relative frequencies’ matrices, we examined the distribution of protein sequences with respect to the family size of the number of aligned proteins for each sequence in the set. The analysis showed that the majority of protein PSSM files are built on family sizes of more than 1000 aligned sequences while ~200k PSSMs are built on families of size less than 10 aligned sequences. Further analysis was conducted by computing the complement of the cumulative distribution at different threshold values (i.e. family sizes). Mathematically speaking, we calculated the distribution of  $1-F(x)$ , where  $F(x)$  is the CDF of the number of aligned proteins to each sequence in our development set. The analysis showed that 9% of the sequences have less than 10 aligned proteins, leaving us with 91% of our total development set with significant encoded evolutionary information. Therefore, the threshold of 10 was picked to discard such sequences with no significant evolutionary information.



**Figure 3.** Proteins sequence length distribution in the final training set (~1.83 M protein sequences)

In ML-based NLP algorithms, two major pieces of information are helpful in the model architecture design and the hyperparameters selection, especially when it comes to the size of mini-batch training enforced by the GPU memory constraint. The first represents the distribution of proteins in our final training set by their sequence length (Figure 3). The histogram displays the length distribution we are dealing with in the training set, and it characterizes a right-skewed distribution with a mean of around 300 residues. The second important piece of information is the distribution of vocabulary in our training set. Figure 4 shows the amino acids composition of our dataset, which is a rather right-skewed distribution than a uniform one. This information will be used to compute the cross-entropy (CE) random baseline for predicting the next amino acid in the sequence.



**Figure 4.** Distribution of the amino acid composition in the final training set (~0.8 billion amino acids)

#### IV. PEVOLM ARCHITECTURE

PEvoLM core technology is LSTM, which is the same machinery used in the initial bi-directional autoregressive approach presented by Peters et al. [17]. ELMos are known to require both large datasets and significant training time and resources to reach convergence. Therefore, we opted for transfer learning using SeqVec/ELMo [36], a sequence embedder trained on a corpus of 9.5 billion amino acids to predict the next residue in a sequence. The architecture provides a per-residue embedding of size  $(3 \times 1024)$ .

Initial experiments showed that a simple bi-directional LSTM succeeds in perfectly learning a mapping function on very small training sets, i.e. a couple of hundreds, but struggles to improve the performance on a large dataset. To boost the capacity of a network, one can either go deeper by increasing the number of layers, go wider by increasing the number of neurons, or both. A joint research work between a faculty member at the University of Toronto and a scientist at Microsoft Research [31] has shown empirically that shallow feed-forward networks can learn the complex functions formerly learned by deep neural networks and obtain performances that were previously only possible with deep architectures. This conclusion was convincing enough to drop the idea of going deeper, in terms of layers, but rather look towards the direction of increasing the number of hidden units.

An LSTM architecture contains several components named *memory blocks*. Such blocks are called gates controlling the information flow in a network, including the input, output and forget gates. The LSTM structure is uniquely defined by the number of its input and output units. To raise the capacity of an LSTM, one can just increase the hidden size units, hence increasing the capacity of the cell state to carry more information along to the next time steps for complex tasks. However, those numbers determine the computational complexity for training an LSTM network. For a moderate number of hidden input units, dimensionality, i.e. the complexity of the network, is largely dominated by the size of the output hidden units. Sak et al. [32] from Google labs suggested an alternative approach that addresses the complexity of large-capacity LSTMs. They proposed the addition of a linear projection layer applied right after the output and before the recurrent connection to the cell. Therefore, the recurrence is applied from a smaller projected hidden state, which ultimately reduces the number of computations within the whole LSTM cell, while maintaining the large capacity of the cell state. The later described module is referred to as LSTM with recurrent projection layer (LSTMP) [32]. This technology has been adopted by Jozefowicz et al. [33] in Google Brain labs to train different variations of language models on a very large corpus of 0.8 billion words and a vocabulary size of 793471 words [34]. Jozefowicz et al. [33] empirically proved that when trying to fit an LSTM network architecture on very large and complex datasets, the size of the LSTM heavily matters.

To improve the LM architecture, I customized the CUDA-optimized Long-Short Term Memory cell implemented with TorchScript<sup>1</sup> to include a recurrent linear projection layer. For the hidden and projection sizes, we decided to go with half the dimensions used in [17 & 36], i.e. 2048 units and 256-dimension projections.

Inspired by the original ELMo paper by Peters et al. [17], a residual skip connection was added from the first to the second LSTM layer, with the goal of boosting the training performance. To further increase the capacity of the network and have more control over the memory size allocated by the


LSTM layers for variable input sizes, we also added a fixed size (1024 hidden unit) non-linear input layer, to the LSTM cell, with LeakyReLU as an activation function. Concerning the weights initialization of the LSTM matrices, we initialized the input projection layer with a Kaiming [35] uniform initialization, the biases of the LSTM gates to a value of 1.0, as it was shown to perform well for long size dependencies [36], and the rest of the matrices were initialized with a uniform distribution of a standard deviation

$$\sigma = \frac{6.0}{\text{hidden\_size} + \text{projection\_size}}.$$

Combining all the LM state-of-the-art technologies discussed in this subsection, a complex network architecture was designed as visualized in figure 5. The model comprises 2 layers of LSTMs, with projections, stacked one after the other. The first layer takes as input SeqVec uncontextualized embeddings concatenated with a one-hot encoding vector of size 20, making it a total input size of 512+20 = 532 units. The second LSTM layer takes as input contextual representations concatenating the output of the first layer with the 2-layer context-aware embeddings from SeqVec, resulting in an input size of 512\*2+256 = 1280 units. As for the residual block, the two representations, of size 256, from both layers are then summed element-wise to serve as input to the two parallel linear layers, one for predicting the next amino acid and the second for inferring the next PSSM column distribution.

For this multi-task learning, two losses are then computed and summed according to eq. 9 to form the final objective function to be minimized. As our primary goal is to learn evolutionary information representations, we want to assign higher weight to the  $\mathbb{K}L$ -Divergence loss. Additionally, previous experiments showed that the CE loss is more dominant in the final objective function, being nearly 4 times larger than the  $\mathbb{K}L$ -divergence loss. Thus, we decided to bring both losses to the same scale and give more weight to the  $\mathbb{K}L$ -divergence loss. Tuning the loss coefficients, our final objective function is defined as:

$$L_{final} = 0.25 * CE_{loss} + 0.75 * KL_{loss}$$



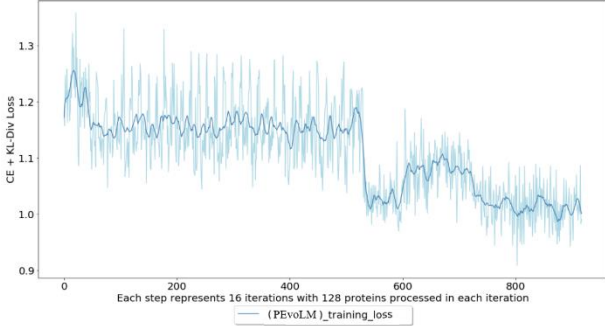
**Figure 5.** Conceptual visualization of our novel single path 2-layer bidirectional embedding language model architecture. The model is trained on predicting both a distribution of the next amino acid and its corresponding evolutionary information relative frequencies. The architecture makes use of LSTMs with projection and a residual connection

<sup>1</sup> <https://pytorch.org/blog/optimizing-cuda-rnn-with-torchscript/>



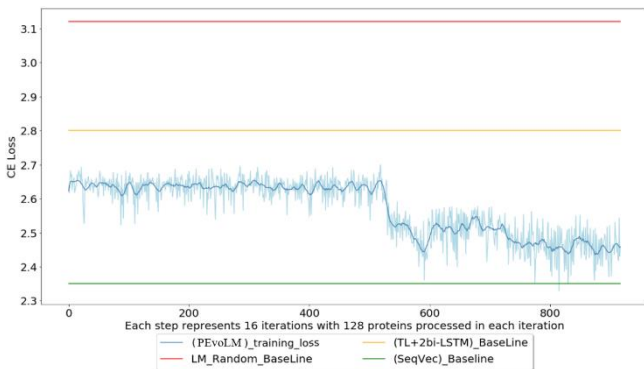
## V. RESULTS AND DISCUSSION

Hyperparameter tuning on the presented architecture revealed that the best-performing model is trained with truncated backpropagation through time of 100 timesteps and a batch size of 128 sequences.



**Figure 6.** Smoothed loss training curve of PEvoLM multi-task learning on a dataset of 1.83 million proteins using a maximum time step of 100 residues.


The model training took 1 week, and the converged training curve for the joint final loss is displayed in Figure 6. The curve shows how the high-capacity architecture has allowed the model to learn representations from our 1.83 million sequences. From the training curve, we clearly see that the loss function has made 3 improvement drops: The first around 800 iterations after processing 100K sequences, the second around 8300 iterations after processing 1 million sequences, and the third 12K iterations after processing 1.5 million proteins. To assess how well each of the separate tasks, predicting either the next residue or the next PSSM column did contribute to the final loss, we plot the training curves of each task separately and compare them to baselines.



**Figure 7.** PEvoLM smoothed loss training curve predicting the next amino acid in the sequence. The model was trained on a dataset of 1.83 Million proteins using a maximum time step of 100 residues. The training curve is compared with three baselines.

Figure 7 shows the training curve (blue) of the final self-learning task predicting the next amino acid in a sequence from all previously seen residues. Comparing the learning trend with the one using plain LSTMs (Orange baseline), we see how the large capacity architecture did improve the training significantly when compared to plain architectures. While the initial plain LSTM architecture’s training maintained a constant rate ( $\sim 2.8$ ) over time, the training loss

of the final architecture starts at a much lower value ( $\sim 2.65$ ), with a low range oscillation throughout the whole training and converged at a loss of  $\sim 2.4$ , which is close to SeqVec reported performance of  $\sim 2.35$  (green baseline).




**Figure 8.** PEvoLM smoothed loss training curve predicting the next PSSM column in the sequence. The model was trained on a dataset of 1.83 million proteins with corresponding relative frequencies using a maximum time step of 100 residues. The training curve is compared with the vanilla LSTM architecture baseline.

Similarly to the cross-entropy loss, the  $\mathbb{KL}$ -divergence showed a comparable behavior for the novel ELMo architecture. Figure 8 displays the next PSSM column smoothed training curve where the model converged around 0.5, more than  $1/3$  lower than the plain LSTM architecture (orange). Table 1 summarizes the performance improvement throughout this research, achieving a joint best loss of 1.0, with a cross-entropy converging around  $CE = 2.4$  and a  $\mathbb{KL}$ -Divergence loss value at  $\mathbb{KL}\text{-Div} = 0.5$ . The best model architecture is defined with a hidden size of 2048 and an output projection size of 256 units. The embedding predictive power was further evaluated on two categories of downstream tasks: the first task involves secondary structure, which is a per-residue type of predictions; the second comprises subcellular localization and soluble vs. membrane proteins as a per-protein level of predictions.

**Table 1.** The table compares the performance of the final ELMo architecture training 2 stacked bidirectional LSTMs with projection via transfer learning (TL). The table shows the distinct performance metrics of the CE and  $\mathbb{KL}$ -Divergence losses as well as the joint final loss. Besides, the table displays the values of the random, plain architecture, and SeqVec reported baselines.

	Cross-Entropy loss (AA)	$\mathbb{KL}$ -Divergence loss (PSSM)	Joint loss (AA + PSSM)	Hidden size / Output size
Random baseline	3.12	-	-	-
SeqVec reported baseline	<b>2.35</b>	-	-	4096/512
TL+2bi-LSTM	2.8	0.8	1.8	256/256
<b>PEvoLM</b>	2.4	<b>0.5</b>	<b>1.0</b>	2048/256



**Figure 9.** (a) 3-state secondary structure prediction comparison between MSA-based and ELMo-based inducers (b) Subcellular localization and membrane vs soluble protein prediction comparison between MSA-based and ELMo-based inducers

Regarding the first evaluation, a model was trained using the new embeddings as input to predict the three states of a protein secondary structure: helix (H), strand (E), and coil (C). The embeddings were evaluated on three test sets:

- $TS_{115}$  [37]: a set of 115 sequences derived from high-quality protein structures (i.e.  $< 3 \text{ \AA}$ ) with no more than 30% PIDE to any protein of known structure in the PDB [38] in 2015.
- $CB_{513}$  [39]: a set of 513 non-redundant sequences compiled after a Structure Integration with Function, Taxonomy and Sequence (SIFTS) [86] mapping.
- $CASP_{12}$  [40]: a set of 21 protein sequences retrieved in 2018 from the  $CASP_{12}$  free-modelling targets after a SIFTS mapping.

Figure 9-a displays the performance results of PEvoLM embedder compared to SeqVec, another ELMo-based model, and ReProf [87], an MSA-based model. ReProf is built on PSSM matrices input generated by MSA methods. The latter is still considered as one of the state-of-the-art methods on this task. From the histogram plot, we see that SeqVec is performing quite well without the need of evolutionary information to make predictions. We also observe that SeqVec and PEvoLM are performing in a quite comparable way where SeqVec is still doing better. However, we should mention that SeqVec is relying on embeddings of size  $3 \times 1024$  for each residue in the sequence to reach this performance, whereas PEvoLM requires embeddings with only half the size of its rival ( $3 \times 512$ ).

The second evaluation was conducted on predicting the membrane-bound proteins from the water-soluble ones. The two-state predictions were tested on a set of 846 proteins retrieved from DeepLoc [41] published supplementary data set. DeepLoc is a state-of-the-art tool relying on MSA profiles output to build its models. It also includes a model predicting the 10 states of subcellular localization, corresponding to our third benchmarking evaluation. The same annotated test set was used to evaluate and benchmark

PEvoLM and SeqVec embeddings with DeepLoc (Figure 9-b). The orange bars correspond to the performance results of the  $Q_2$  membrane predictions. We see that PEvoLM is lower by 10% and 4% compared to DeepLoc and SeqVec, respectively.

While DeepLoc is still outperforming both of the ELMo-based methods ( $Q_2 = 93\%$ ), our suggested embedding performance is still significantly high hitting an 83% with an embedding size of only  $3 \times 512$  per residue. However, the  $Q_{10}$  localization predictions seem to perform poorly, with an overall accuracy of 50%, which is 28% lower than DeepLoc and 18% lower than SeqVec.

## VI. CONCLUSION

The current algorithms for generating evolutionarily related information of protein sequences is largely dominated by multiple sequence alignment methods. We have seen that this technique is one of the most widely used modeling approaches in biology. MSAMs are used to expose those restricted evolutionary regions within a sequence. The results of these methods represent an essential input to several downstream applications in the field of bioinformatics. The process is simply described as searching for homologs of a query sequence in a database of protein sequences, capturing the conservation patterns in the alignment, and storing this information as a matrix of numerical scores for each position in the alignment.

Even though, the evolutionary information representations generated by MSA methods have revolutionized the prediction power of AI methods for the past two decades, this increase in performance has become costly in recent years, with the continuous exponential growth of bio-sequence data pools. Thanks to the similarity of protein sequences to natural language, NLP state-of-the-art algorithms, such as ELMo, has successfully been applied in bioinformatics.

The main objective of this research was to deploy an ELMo tool incorporating evolutionary information in its representations and offering a better trade-off between computing resources and runtime. While pre-training is costly, the inference is cheaper. A novel bidirectional language model was trained following the autoregressive approach. The model was trained on a set of 1.83 million proteins (~0.8 billion residues), predicting both the next amino acid in a sequence and the probability distribution of the next residue derived from similar, yet different, sequences, as summarized in a PSSM. With 975MB required GPU memory to load the final pre-trained model, the average inference time takes around 1.03 second to embed a human protein sequence segment of size 512 residues. The embeddings' prediction power was evaluated on three downstream tasks: secondary structure, subcellular localization, and membrane vs. soluble protein predictions.

Even if we did not succeed in reaching our initial goal of outperforming SeqVec, we did achieve a performance close to SeqVec with half the embedding dimensions (i.e. 3 x 512). Additionally, the results have confirmed a finding by Jozefowicz et al. [33], who did empirically prove that, when trying to fit an LSTM network architecture on very large and complex datasets, the size of the LSTM matters. In general, and from the downstream tasks perspective, there was no evidence that adding PSSM input will improve the knowledge that LMs could learn from a sequence. Different hypotheses can be drawn here: SeqVec was able to learn certain patterns already without PSSMs, our dataset was too small for training an LM, we should have trained the whole architecture from scratch (i.e. without transfer learning), we should have used larger capacity LSTMs, or we should have simply trained longer.

#### ACKNOWLEDGMENT

The author would like to thank Michael Heinzinger, Violetta Cavalli-Sforza, and Burkhard Rost for their insights and feedback on this piece of work.

#### REFERENCES

- [1] Barton J. Geoffrey, Protein Sequence Alignment Techniques, European Molecular Biology Laboratory Outstation, The European Bioinformatics Institute, Wellcome Trust Genome Campus, Hinxton, Cambridge CB10 1SD, England. *Acta Cryst.* (1998), D54,1139-1146
- [2] Maria Chatzou, Cedrik Magis, Jia-Ming Chang, Carsten Kemena, Giovanni Bussotti, Ionas Erb, Cedric Notredame, Multiple sequence alignment modeling: methods and applications, *Briefings in Bioinformatics*, Volume 17, Issue 6, November 2016, Pages 1009–1023, doi: <https://doi.org/10.1093/bib/bbv099>
- [3] Rost B., Sander C. (1993) Prediction of protein secondary structure at better than 70% accuracy. *J Mol Biol.* 232:584–99.
- [4] Rost B, Sander C. Improved prediction of protein secondary structure by use of sequence profiles and neural networks. *Proc Natl Acad Sci.* 1993;90:7558–62.
- [5] Barton GJ. Protein secondary structure prediction. *Curr Opin Struct Biol.* 1995;5:372–6.
- [6] Chandonia J-M, Karplus M. Neural networks for secondary structure and structural class predictions. *Protein Sci.* 1995;4:275–85.
- [7] Bigelow H, Petrey D, Liu J, Przybylski D, Rost B. Predicting transmembrane beta-barrels in proteomes. *Nucleic Acids Res.* 2004;32:2566–77.
- [8] Rost B, Casadio R, Fariselli P. Topology prediction for helical transmembrane proteins at 86% accuracy. *Protein Sci.* 1996;5:1704–18.

- [9] Rost B, Casadio R, Fariselli P, Sander C. Transmembrane helix prediction at 95% accuracy. *Protein Sci.* 1995;4:521–33.
- [10] Punta M, Rost B. PROFcon: novel prediction of long-range contacts. *Bioinform.* 2005;21(13):2960–8.
- [11] Nair R, Rost B. Better prediction of sub-cellular localization by combining evolutionary and structural information. *Proteins.* 2003;53(4):917–30.
- [12] Nair R, Rost B. Mimicking cellular sorting improves prediction of subcellular localization. *J Mol Biol.* 2005;348(1):85–100.
- [13] Marino Buslje C, Teppa E, Di Domenico T, Delfino JM, Nielsen M. Networks of high mutual information define the structural proximity of catalytic sites: implications for catalytic residue identification. *PLoS Comput Biol.* 2010;6(11):e1000978.
- [14] Ofra Y, Rost B. ISIS: interaction sites identified from sequence. *Bioinform.* 2007;23(2):e13–6.
- [15] Ofra Y, Rost B. Protein-protein interaction hot spots carved into sequences. *PLoS Comput Biol.* 2007;3(7):e119.
- [16] Suzek BE, Wang Y, Huang H, McGarvey PB, Wu CH, UniProt C. UniRef clusters: a comprehensive and scalable alternative for improving sequence similarity searches. *Bioinform.* 2015;31(6):926–32.
- [17] Peters M., Neumann M., Iyyer M., Gardner M., Clark C., Lee K., Zettlemoyer L. (2018). Deep contextualized word representations. arXiv:1802.05365
- [18] Mikolov T., Sutskever I., Chen K., Corrado G. S., and Dean J. (2013). Distributed representations of words and phrases and their compositionality. In NIPS.
- [19] Pennington J., Socher R., and Manning C. D. (2014). Glove: Global vectors for word representation. In EMNLP
- [20] Turian J. P., Ratinov L., and Bengio Y. (2010). Word representations: A simple and general method for semi-supervised learning. In ACL
- [21] He L., Lee K., Lewis M., and Zettlemoyer L. S. (2017). Deep semantic role labeling: What works and what's next. In ACL
- [22] Liu X., Shen Y., Duh K., and Gao. J. (2017). Stochastic answer networks for machine reading comprehension. arXiv preprint arXiv:1712.03556.
- [23] Chen Q., Zhu X., Ling Z., Wei S., Jiang H., and Inkpen D. (2017). Enhanced lstm for natural language inference. In ACL.
- [24] Akaike H. (1998) Autoregressive Model Fitting for Control. In: Parzen E., Tanabe K., Kitagawa G. (eds) Selected Papers of Hirotugu Akaike. Springer Series in Statistics (Perspectives in Statistics). Springer, New York, NY
- [25] Heinzinger, M., Elnaggar, A., Wang, Y. et al. Modeling aspects of the language of life through transfer-learning protein sequences. *BMC Bioinformatics* 20, 723 (2019). <https://doi.org/10.1186/s12859-019-3220-8>
- [26] Akaike H. (1998) Autoregressive Model Fitting for Control. In: Parzen E., Tanabe K., Kitagawa G. (eds) Selected Papers of Hirotugu Akaike. Springer Series in Statistics (Perspectives in Statistics). Springer, New York, NY
- [27] C. Zhang, J. Bütepage, H. Kjellström and S. Mandt, "Advances in Variational Inference," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 8, pp. 2008–2026, 1 Aug. 2019, doi: 10.1109/TPAMI.2018.2889774.
- [28] Lu Y. & Lu J. (2020). A Universal Approximation Theorem of Deep Neural Networks for Expressing Distributions. arxiv:2004.08867v2
- [29] The UniProt Consortium, UniProt: a worldwide hub of protein knowledge, *Nucleic Acids Research*, Volume 47, Issue D1, 08 January 2019, Pages D506–D515, <https://doi.org/10.1093/nar/gky1049>
- [30] Rost, B., Yachdav, G., & Liu, J. (2004). The PredictProtein server. *Nucleic acids research*, 32(Web Server issue), W321–W326. doi: 10.1093/nar/gkh377
- [31] Ba J., Caruana R. (2014) Do Deep Nets Really Need to be Deep? doi: arXiv:1312.6184v7
- [32] Sak H., Senior A., Beaufays F. (2014) Long Short-Term Memory Based Recurrent Neural Network Architectures for Large Vocabulary Speech Recognition. doi: arXiv:1402.1128v1
- [33] Jozefowicz R., Vinyals O., Schuster M., Shazeer N, and Wu Y. (2016) Exploring the limits of language modeling. *CoRR abs/1602.02410*.
- [34] Chelba C., Mikolov T., Schuster M., Ge Q., Brants T., Koehn P., and Robinson T. (2013) One billion word benchmark for measuring



progress in statistical language modeling. arXiv preprint arXiv:1312.3005

- [35] He K., Xiangyu Z., Shaoqing R., and Jian S. (2015) Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In International Conference on Computer Vision (ICCV), 2015. URL <http://arxiv.org/abs/1502.01852>.
- [36] Jozefowicz R., Zaremba W., and Sutskever I. (2015) An empirical exploration of recurrent network architectures. In Proceedings of the 32nd International Conference on Machine Learning (ICML-15), pp. 2342–2350, 2015.
- [37] Yang Y, Gao J, Wang J, Heffernan R, Hanson J, Paliwal K, Zhou Y. Sixty-five years of the long march in protein secondary structure prediction: the final stretch? *Brief Bioinform.* 2016;19(3):482–94.
- [38] Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, Shindyalov IN, Bourne PE. The protein data bank. *Nucleic Acids Res.* 2000;28(1):235–42.
- [39] Cuff JA, Barton GJ. Evaluation and improvement of multiple sequence methods for protein secondary structure prediction. *Proteins Struct Funct Genet.* 1999;34(4):508–19.
- [40] Abriata LA, Tamò GE, Monastyrskyy B, Kryshtafovych A, Dal Peraro M. Assessment of hard target modeling in CASP12 reveals an emerging role of alignment-based contact prediction methods. *Proteins.* 2018;86:97–112
- [41] Almagro Armenteros JJ, Sonderby CK, Sonderby SK, Nielsen H, Winther O. DeepLoc: prediction of protein subcellular localization using deep learning. *Bioinform.* 2017;33(24):4049.