

# MODELING THE LANGUAGE OF LIFE – DEEP LEARNING PROTEIN SEQUENCES

Michael Heinzinger<sup>1,2,\*</sup>, Ahmed Elnaggar<sup>1,2,\*</sup>, Yu Wang<sup>3</sup>, Christian  
Dallago<sup>1,2</sup>, Dmitrii Nechaev<sup>1,2</sup>, Florian Matthes<sup>4</sup> & Burkhard Rost<sup>1,5</sup>

1 TUM (Technical University of Munich) Department of Informatics, Bioinformatics & Computational Biology - i12, Boltzmannstr. 3, 85748 Garching/Munich, Germany

2 TUM Graduate School, Center of Doctoral Studies in Informatics and its Applications (CeDoSIA), Boltzmannstr. 11, 85748 Garching, Germany

3 Leibniz Supercomputing Centre, Boltzmannstr. 1, 85748 Garching/Munich, Germany

4 TUM Department of Informatics, Software Engineering and Business Information Systems, Boltzmannstr. 1, 85748 Garching/Munich, Germany

5 Institute for Advanced Study (TUM-IAS), Lichtenbergstr. 2a, 85748 Garching/Munich, Germany & TUM School of Life Sciences Weihenstephan (WZW), Alte Akademie 8, Freising, Germany & Columbia University, Department of Biochemistry and Molecular Biophysics & New York Consortium on Membrane Protein Structure (NYCOMPS), 701 West, 168<sup>th</sup> Street, New York, NY 10032, USA

\* These authors contributed equally to this work.

Corresponding author: mheinzinger[at]rostlab.org, tel: +49-289-17-811 (email rost: assistant[@]rostlab.org)

## Abstract

**Background:** One common task in Computational Biology is the prediction of aspects of protein function and structure from their amino acid sequence. For 26 years, most state-of-the-art approaches toward this end have been marrying machine learning and evolutionary information. The retrieval of related proteins from ever growing sequence databases is becoming so time-consuming that the analysis of entire proteomes becomes challenging. On top, evolutionary information is less powerful for small families, e.g. for proteins from the *Dark Proteome*.

**Results:** We introduce a novel way to represent protein sequences as continuous vectors (*embeddings*) by using the deep bi-directional model ELMo taken from natural language processing (NLP). The model has effectively captured the biophysical properties of protein sequences from unlabeled big data (UniRef50). After training, this knowledge is transferred to single protein sequences by predicting relevant sequence features. We refer to these new embeddings as *SeqVec* (*Sequence-to-Vector*) and demonstrate their effectiveness by training simple convolutional neural networks on existing data sets for two completely different prediction tasks. At the per-residue level, we significantly improved secondary structure (for NetSurfP-2.0 data set: Q3=79%±1, Q8=68%±1) and disorder predictions (MCC=0.59±0.03) over methods not using evolutionary information. At the per-protein level, we predicted subcellular localization in ten classes (for DeepLoc data set: Q10=68%±1) and distinguished membrane-bound from water-soluble proteins (Q2= 87%±1). All results built upon the embeddings gained from the new tool *SeqVec* neither explicitly nor implicitly using evolutionary information. Nevertheless, it improved over some methods using such information. Where the lightning-fast *HHblits* needed on average about two minutes to generate the evolutionary information for a target protein, *SeqVec* created the vector representation on average in 0.03 seconds.

1 **Conclusion:** We have shown that transfer learning can be used to capture biochemical or  
2 biophysical properties of protein sequences from large unlabeled sequence databases. The  
3 effectiveness of the proposed approach was showcased for different prediction tasks using  
4 only single protein sequences. *SeqVec* embeddings enable predictions that outperform even  
5 some methods using evolutionary information. Thus, they prove to condense the underlying  
6 principles of protein sequences. This might be the first step towards competitive predictions  
7 based only on single protein sequences.

8  
9 **Availability:** SeqVec: <https://github.com/mheininger/SeqVec> Prediction server: <https://embed.protein.properties>

10 **Key words:** Machine Learning, Language Modeling, Sequence Embedding, Secondary structure prediction,  
11 Localization prediction, Transfer Learning, Deep Learning.

12 **Abbreviations:** **1D**, one-dimensional – information representable in a string such as secondary structure or solvent  
13 accessibility; **3D**, three-dimensional; **3D structure**, three-dimensional coordinates of protein structure; **MCC**,  
14 Matthews-Correlation-Coefficient; **RSA**, relative solvent accessibility;

## 16 Background

17 Over two decades ago, the combination of evolutionary information (from Multiple Sequence  
18 Alignments – MSA) and machine learning (standard feed-forward artificial neural networks –  
19 ANN) completely changed protein secondary structure prediction [1-3]. The concept was  
20 quickly taken up [4-8], and it was shown how the improvement in prediction increased with  
21 even larger families including more diverse evolutionary information [9, 10]. The idea was  
22 applied to other tasks, including the prediction of transmembrane regions [11-13], solvent  
23 accessibility [14], residue flexibility (B-values) [15, 16], inter-residue contacts [17] and protein  
24 disorder [15, 18-20]. Later, methods predicting aspects of protein function improved through  
25 the combination of evolutionary information and machine learning, including predictions of sub-  
26 cellular localization (aka cellular location [21, 22]), protein interaction sites [23-25], and the  
27 effects of sequence variation upon function [26, 27]. Arguably, the most important  
28 breakthrough for protein structure prediction over the decade was a more efficient way of using  
29 evolutionary couplings [28-31].

30 Although evolutionary information has become increasingly crucial, it is also becoming  
31 increasingly costly. Firstly, bio-sequence databases grow faster than computers making it  
32 expensive to find and align related proteins. For instance, the number of UniProt entries is now  
33 more than doubling every two years [32]. Consequently, methods as fast as PSI-BLAST [33]  
34 have to be replaced by faster solutions such as HHblits [34]. Even its latest version HHblits3  
35 [35] still needs several minutes to search UniRef50 (subset of UniProt; ~20% of UniProt  
36 release 2019\_02). The next step up in speed such as MMSeqs2 [36] appear to cope with the  
37 challenge at the expense of increasing hardware requirements while databases keep growing.  
38 However, even these solutions might eventually lose the battle against the speedup of  
39 sequencing. Analyzing data sets involving millions of proteins, i.e. samples of the human gut  
40 microbiota or metagenomic samples, will be one of the major challenges for computational  
41 biology. Secondly, evolutionary information is missing for some proteins, e.g. for proteins with  
42 substantial intrinsically disordered regions [15, 37, 38], or the entire *Dark Proteome* [39] full of  
43 proteins that are less-well studied but important for function [40].

1

2 Here, we propose a novel embedding of protein sequences that replaces the explicit search  
3 for evolutionary related proteins by an implicit transfer of biophysical information derived from  
4 large, unlabeled sequence data (here UniRef50). Towards this end, we have adopted a  
5 method that has been revolutionizing Natural Language Processing (NLP), namely the bi-  
6 directional language model ELMo (Embeddings from Language Models) [41]. In NLP, ELMo  
7 is trained on unlabeled text-corpora such as Wikipedia to predict the most probable next word  
8 in a sentence, given all previous words in this sentence. By learning a probability distribution  
9 for sentences, these models develop autonomously a notion for syntax and semantics of  
10 language. The trained vector representations (embeddings) are contextualized, i.e.  
11 embeddings of a given word depend on its context. This has the advantage that two identical  
12 words can have different embeddings, depending on the words surrounding them.

13 We hypothesized that the ELMo concept could be applied to learn aspects of what  
14 makes up the language of life distilled in protein sequences. Three main challenges arose. (1)  
15 Proteins range from about 30 to 33,000 residues, a much larger range than for the average  
16 English sentence extending over 15-30 words [42], and even more extreme than notable  
17 literary exceptions such as James Joyce's *Ulysses* (1922) with almost 4,000 words in a  
18 sentence. Longer proteins require more GPU memory and the underlying models (so called  
19 Long Short-Term Memory networks (LSTMs) [43]) have only a limited capability to remember  
20 long-range dependencies. (2) Proteins mostly use 20 standard amino acids, 100,000 times  
21 less than in the English language. Smaller vocabularies might be problematic if protein  
22 sequences encode a similar complexity as sentences. (3) We found UniRef50 to contain  
23 almost ten times more tokens (9.5 billion amino acids) than the largest existing NLP corpus (1  
24 billion words). Simply put: Wikipedia is roughly ten times larger than Webster's Third New  
25 International Dictionary and the entire UniProt is over ten times larger than Wikipedia. As a  
26 result, larger models might be required to absorb the information provided.

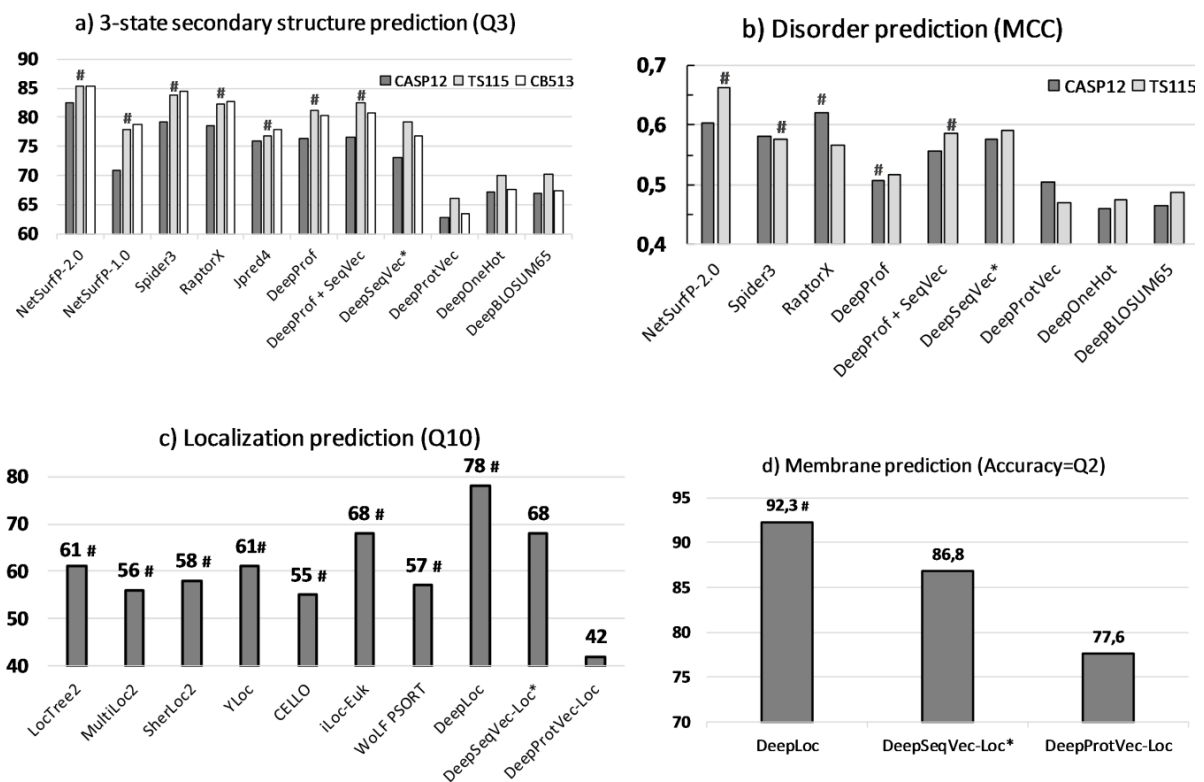
27 We trained the bi-directional language model ELMo on UniRef50. Then we assessed  
28 the predictive power of the embeddings by application to tasks on two levels: per-residue  
29 (word-level) and per-protein (sentence-level). For the per-residue prediction task, we predicted  
30 secondary structure in three (helix, strand, other) and eight states (all DSSP [44]), as well as  
31 long intrinsic disorder in two states. For the per-protein prediction task, we implemented the  
32 predictions of protein subcellular localization in ten classes and a binary classification into  
33 membrane-bound and water-soluble proteins. We used publicly available data sets from two  
34 recent methods that achieved break-through performance through Deep Learning, namely  
35 NetSurfP-2.0 (secondary structure [45]) and DeepLoc (localization [46]). We compared the  
36 performance of the *SeqVec* embeddings to state-of-the-art methods, and also to a popular  
37 embedding tool for protein sequences, namely *ProtVec* [47]. *ProtVec* assumes that every  
38 token or word consists of three consecutive residues (amino acid 3-mers). During training,  
39 each protein sequence in *SwissProt* [48] is split into overlapping 3-mers and another tool from  
40 NLP, namely the *word2vec* (skip-gram) model [49] is used to predict adjacent words, given the  
41 word at the center. After training, protein sequences can be split into overlapping 3-mers which  
42 are mapped onto a 100-dimensional latent space. While this approach can capture local  
43 information, it loses information on sequence ordering and the resulting embeddings are  
44 insensitive to their context (non-contextualized), i.e. the same word results in the same  
45 embedding regardless of the specific context.

## Results

**SeqVec embeddings compress UniProt50.** *SeqVec*, our ELMo-like implementation was trained for three weeks on 5 Nvidia Titan GPUs with 12 GB memory each. The model was trained until its *perplexity* (uncertainty when predicting the next token) converged at around 10.5 (Fig. SOM\_1). Training and testing were not split due to technical limitations (incl. shortage of CPU/GPU). ELMo was designed to reduce the risk of overfitting by sharing weights between forward and backward LSTMs and by using dropout. The model had about 93M (mega/million) free parameters which was about 100-times smaller than the 9.6G (giga/billion) tokens to predict. Similar approaches have shown that even today's largest models (750M free parameters), are not able to overfit on a large corpus (250M protein sequences) [50].

**Per-residue performance high, not highest.** NetSurfP-2.0 uses HHblits profiles along with advanced combinations of Deep Learning architectures [45]. NetSurfP-2.0 could be today's best method for protein secondary structure prediction, reaching a three-state per-residue accuracy Q3 of 82-85% (lower value: small very non-redundant CASP12 set, upper value: larger slightly more redundant TS115 and CB513 sets; Table 1, Fig. 1; note that several contenders such as *Spider3* and *RaptorX* appear to differ by fewer than three standard errors). All six methods developed by us (DeepProf, DeepSeqVec, DeepProf+SeqVec, DeepProtVec, DeepOneHot, DeepBLOSUM65) fell short of reaching this mark (Fig. 1A, Table 1). When comparing methods that use only single protein sequences as input (DeepSeqVec, DeepProtVec, DeepOneHot, DeepBLOSUM65; all white in Table 1), the proposed *SeqVec* outperformed others by 5-10 (Q3), 5-13 (Q8) and 0.07-0.12 (MCC) percentage points. The evolutionary information (DeepProf with HHblits profiles) remained about 4-6 percentage points below NetSurfP-2.0 (Q3=76-81%, Fig. 1, Table 1). Depending on the test set, using *SeqVec* embeddings instead of evolutionary information (DeepSeqVec: Fig. 1A, Table 1) remained 2-3 percentage points below that mark (Q3=73-79%, Fig. 1A, Table 1). Using both evolutionary information and *SeqVec* embeddings (DeepProf+SeqVec) improved over both, but still did not reach the top (Q3=77-82%). In fact, the embedding alone (DeepSeqVec) did not surpass any of the existing methods using evolutionary information (Fig. 1A).

For the prediction of intrinsic disorder, we observed the same: NetSurfP-2.0 performed best; our implementation of evolutionary information (DeepProf) performed worse (Fig. 1B, Table 1). However, for this task the embeddings alone (DeepSeqVec) performed relatively well, exceeding our in-house implementation of a model using evolutionary information (DeepSeqVec MCC=0.575-0.591 vs. DeepProf MCC=0.506-0.516, Table 1). The combination of evolutionary information and embeddings (DeepProf+SeqVec) improved over using evolutionary information alone but did not improve over the *SeqVec* embeddings for disorder. Compared to other methods, the embeddings alone reached similar values (Fig. 1B).



**Figure 1: Performance comparisons.** The predictive power of the ELMo-like SeqVec embeddings was assessed for per-residue (upper row) and per-protein (lower row). Methods using evolutionary information (mostly in the form of alignments) are highlighted by a '#' above the bar(s) of the methods. Approaches using only the proposed SeqVec embeddings are highlighted by a '\*'. Panel A compared three-state secondary structure prediction of the proposed SeqVec to other embeddings based on single protein sequences. Panel B compared predictions of intrinsically disordered residues. Panel C compared per-protein predictions for subcellular localization between top methods (numbers taken from DeepLoc [46] and embeddings based on single sequences (ProtVec [47] and our SeqVec). Panel D: the same data set was used to assess the predictive power of SeqVec for the classification of a protein into membrane-bound and water-soluble.



1 **Table 1: Per-residue predictions: secondary structure and disorder** <sup>◇</sup>

Data	Prediction task	Secondary structure		Disorder	
	Method	Q3 (%)	Q8 (%)	MCC	FPR
CASP12	<i>NetSurfP-2.0 (hhblits)*</i>	<b>82.4</b>	<b>71.1</b>	0.604	<b>0.011</b>
	<i>NetSurfP-1.0*</i>	70.9	-	-	-
	<i>Spider3*</i>	79.1	-	0.582	0.026
	<i>RaptorX*</i>	78.6	66.1	<b>0.621</b>	0.045
	<i>Jpred4*</i>	76.0	-	-	-
	<i>DeepSeqVec</i>	73.1 ± 1.3	61.2 ± 1.6	0.575 ± 0.075	0.026 ± 0.008
	<i>DeepProf</i>	76.4 ± 2.0	62.7 ± 2.2	0.506 ± 0.057	0.022 ± 0.009
	<i>DeepProf + SeqVec</i>	76.5 ± 1.5	64.1 ± 1.5	0.556 ± 0.080	0.022 ± 0.008
	<i>DeepProtVec</i>	62.8 ± 1.7	50.5 ± 2.4	0.505 ± 0.064	0.016 ± 0.006
	<i>DeepOneHot</i>	67.1 ± 1.6	54.2 ± 2.1	0.461 ± 0.064	0.012 ± 0.005
	<i>DeepBLOSUM65</i>	67.0 ± 1.6	54.5 ± 2.0	0.465 ± 0.065	0.012 ± 0.005
TS115	<i>NetSurfP-2.0 (hhblits)*</i>	<b>85.3</b>	<b>74.4</b>	<b>0.663</b>	<b>0.006</b>
	<i>NetSurfP-1.0*</i>	77.9	-	-	-
	<i>Spider3*</i>	83.9	-	0.575	0.008
	<i>RaptorX*</i>	82.2	71.6	0.567	0.027
	<i>Jpred4*</i>	76.7	-	-	-
	<i>DeepSeqVec</i>	79.1 ± 0.8	67.6 ± 1.0	0.591 ± 0.028	0.012 ± 0.001
	<i>DeepProf</i>	81.1 ± 0.6	68.3 ± 0.9	0.516 ± 0.028	0.012 ± 0.002
	<i>DeepProf + SeqVec</i>	82.4 ± 0.7	70.3 ± 1.0	0.585 ± 0.029	0.013 ± 0.003
	<i>DeepProtVec</i>	66.0 ± 1.0	54.4 ± 1.3	0.470 ± 0.028	0.011 ± 0.002
	<i>DeepOneHot</i>	70.1 ± 0.8	58.5 ± 1.1	0.476 ± 0.028	0.008 ± 0.001
	<i>Deep BLOSUM65</i>	70.3 ± 0.8	58.1 ± 1.1	0.488 ± 0.029	0.007 ± 0.001
CB513	<i>NetSurfP-2.0 (hhblits)*</i>	<b>85.3</b>	<b>72.0</b>	-	-
	<i>NetSurfP-1.0*</i>	78.8	-	-	-
	<i>Spider3*</i>	84.5	-	-	-
	<i>RaptorX*</i>	82.7	70.6	-	-
	<i>Jpred4*</i>	77.9	-	-	-
	<i>DeepSeqVec</i>	76.9 ± 0.5	62.5 ± 0.6	-	-
	<i>DeepProf</i>	80.2 ± 0.4	64.9 ± 0.5	-	-
	<i>DeepProf + SeqVec</i>	80.7 ± 0.5	66.0 ± 0.5	-	-
	<i>DeepProtVec</i>	63.5 ± 0.4	48.9 ± 0.5	-	-
	<i>DeepOneHot</i>	67.5 ± 0.4	52.9 ± 0.5	-	-
	<i>DeepBLOSUM65</i>	67.4 ± 0.4	53.0 ± 0.5	-	-

2  
3 \* Performance comparison for 3- and 8-class secondary structure prediction as well as disorder  
4 prediction for the CASP12, TS115 and CB513 data sets. Accuracies (Q3, Q10) are given in  
5 percentage. Results marked with \* are taken from NetSurfP-2.0 [45]; the authors did not provide  
6 standard errors. DeepSeqVec, DeepProtVec, DeepOneHot and DeepBLOSUM65 use only  
7 information from single protein sequences. All methods using evolutionary information are marked  
8 by gray boxes, these performed best (best in each set marked by bold numbers) for all tasks and  
9 measures.

1 **Per-protein performance close to best.** For the task of predicting subcellular localization in  
 2 ten classes, DeepLoc [46] appears to be today's best tool with Q10=78% (Fig. 1C, Table 2).  
 3 Our simple embeddings model DeepSeqVec-Loc reached second best performance together  
 4 with iLoc-Euk [51] with Q10=68% (Fig. 1C, Table 2). For this application the SeqVec  
 5 embeddings outperformed several other methods that use evolutionary information by up to  
 6 Q10=13%.

7  
 8 **Table 2: Per-protein predictions: localization and membrane/globular** <sup>◇</sup>

Method	Localization		Membrane/globular	
	Q10 (%)	Gorodkin (MCC)	Q2	MCC
<i>LocTree2</i> *	61	0.53		
<i>MultiLoc2</i> *	56	0.49		
<i>CELLO</i> *	55	0.45		
<i>WoLF PSORT</i> *	57	0.48		
<i>YLoc</i> *	61	0.53		
<i>SherLoc2</i> *	58	0.51		
<i>iLoc-Euk</i> *	68	0.64		
<i>DeepLoc</i> *	<b>78</b>	<b>0.73</b>	<b>92.3</b>	<b>0.844</b>
<i>DeepSeqVec-Loc</i>	68 ± 1	0.61 ± 0.01	86.8 ± 1.0	0.725 ± 0.021
<i>DeepProtVec-Loc</i>	42 ± 1	0.19 ± 0.01	77.6 ± 1.3	0.531 ± 0.026

9  
 10 <sup>◇</sup> Performance for per-protein prediction of subcellular localization and the classification of proteins  
 11 into membrane-bound and water-soluble. Results marked with \* were taken from DeepLoc [46];  
 12 the authors did not provide standard errors. The results reported for SeqVec and ProtVec were  
 13 based on single protein sequences, i.e. methods NOT using evolutionary information (neither  
 14 during training nor testing). All methods using evolutionary information are marked by gray boxes,  
 15 these performed best (best in each set marked by bold numbers) for all tasks and measures.

16  
 17 Performance for the classification into membrane-bound and water-soluble proteins followed  
 18 a similar trend (Fig. 1D, Table 2): while DeepLoc still performed best (Q2=92.3, MCC=0.844),  
 19 DeepSeqVec-Loc reached just a few percentage points lower (Q2=86.8±1.0,  
 20 MCC=0.725±0.021; full confusion matrix Figure SOM\_2). In contrast to this, ProtVec, another  
 21 method using only single sequences, performed substantially worse (Q2=77.6±1.3,  
 22 MCC=0.531±0.026).

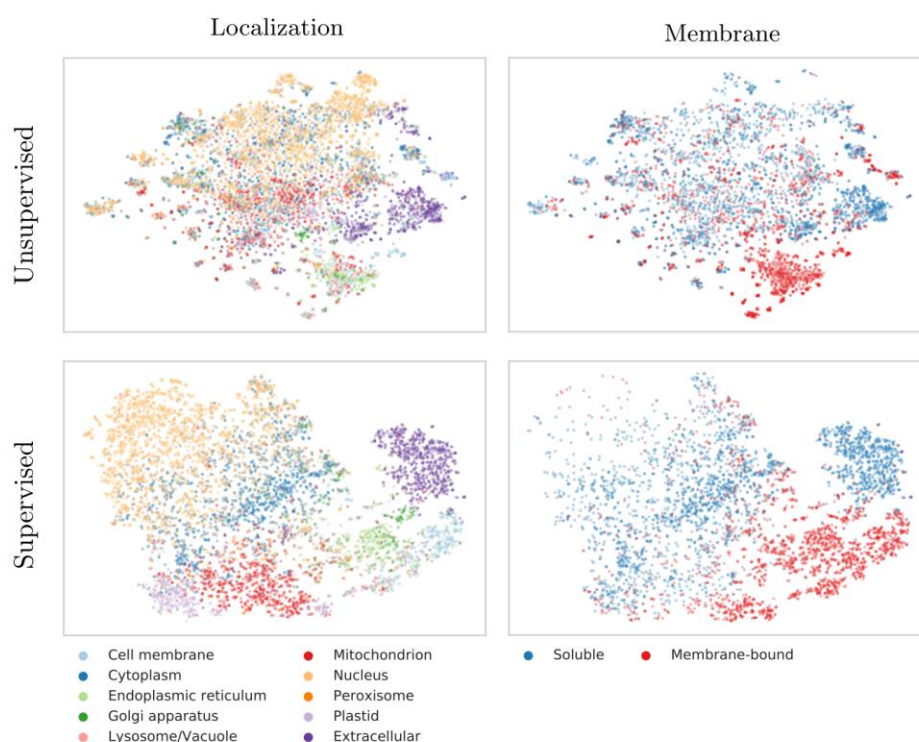
23  
 24 **Visualizing results.** Lack of insight often triggers the misunderstanding that machine learning  
 25 methods are black box solutions barring understanding. In order to interpret the *SeqVec*  
 26 embeddings, we have projected the protein-embeddings of the per-protein prediction data  
 27 upon two dimensions using t-SNE [52]. We performed this analysis once for the raw  
 28 embeddings (*SeqVec*, Fig. 2 upper row) and once for the hidden layer representation of the  
 29 per-protein network (*DeepSeqVec-Loc*) after training (Fig. 2 lower row). All t-SNE  
 30 representations in Fig. 2 were created using 3,000 iterations and the cosine distance as metric.  
 31 The two analyses differed only in that the perplexity was set to 20 for one (*SeqVec*) and 15 for  
 32 the other (*DeepSeqVec-Loc*). The t-SNE representations were colored either according to their

1 localization within the cell (left column of Fig. 2) or according to whether they are membrane-  
2 bound or water-soluble (right column).

3  
4 Despite never provided during training, the raw embeddings appeared to capture some  
5 signal for classifying proteins by localization (Fig. 2, upper row, left column). The most  
6 consistent signal was visible for extra-cellular proteins. Proteins attached to the cell membrane  
7 or located in the endoplasmic reticulum also formed well-defined clusters. In contrast, the raw  
8 embeddings neither captured an ambiguous signal for nuclear nor for mitochondrial proteins.  
9 Through training, the network improved the signal to reliably classify mitochondrial and plastid  
10 proteins. However, proteins in the nucleus and cell membrane continued to be poorly  
11 distinguished via t-SNE.

12 Coloring the t-SNE representations for membrane-bound or water-soluble proteins  
13 (Fig. 2, right column), revealed that the raw embeddings already provided well-defined clusters  
14 although never trained on membrane prediction (Fig. 2, upper row). After training, the  
15 classification was even better (Fig. 2, lower row).

16



17

18 **Figure 2: t-SNE representations of SeqVec.** t-SNE projections from embedded space onto a 2D  
19 representation; upper row: unsupervised 1024-dimensional ELMO embeddings, averaged over all  
20 residues in a protein; lower row: supervised 32 dimensional ELMO embeddings, reduced via per-protein  
21 machine learning predictions. The redundancy reduced DeepLoc data set was used for this figure.  
22 Proteins were colored according to their localization (left column) or whether they are membrane-bound  
23 or water-soluble (right column). Not all proteins in the data set have annotations for the classification  
24 into membrane-bound/water-soluble and have thus been excluded in the panels on the right.



1 The plots in the upper row suggest that the SeqVec embeddings already capture various aspects of  
2 proteins prior even to supervised learning. After supervised training (lower row), this information is  
3 transferred to, and further distilled by networks with simple architectures to distinguish aspects of protein  
4 function (sometimes drastically, as in the case of membrane boundness, as suggested by the coloured  
5 clusters in the bottom right panel).

6 After training machine learning devices (lower row) on the supervised tasks of localization (left) and  
7 membrane boundness (right) prediction, the power of SeqVec embeddings to distinguish aspects of  
8 function and structure become even more pronounced, sometimes drastically so, as suggested by the  
9 almost fully separable clusters in the lower right panel.

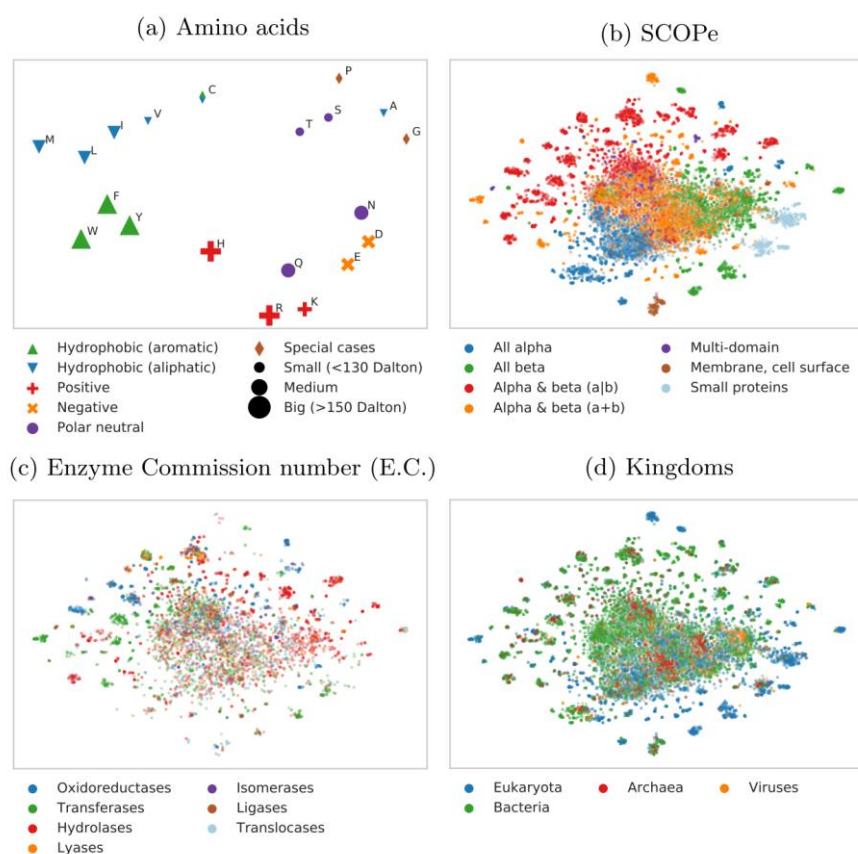
10

11 Analogously, we used t-SNE projections to analyze SeqVec embeddings on different  
12 levels of complexity inherent to proteins (Fig. 3), ranging from the building blocks (amino acids,  
13 Fig. 3a), to secondary structure defined protein classes (Fig. 3b), over functional features (Fig.  
14 3c), and onto the macroscopic level of the kingdoms of life and viruses (Fig. 3d; classifications  
15 in panels 3b-3d based on SCOPe [53]). Similar to the results described in [50], our projection  
16 of the embedding space confirmed that the model successfully captured bio-chemical and bio-  
17 physical properties on the most fine-grained level, i.e. the 20 standard amino acids (Fig. 3a).  
18 For example, aromatic amino acids (W, F, Y) are well separated from aliphatic amino acids (A,  
19 I, L, M, V) and small amino acids (A, C, G, P, S, T) are well separated from large ones (F, H,  
20 R, W, Y). The projection of the letter indicating an unknown amino acid (X), clustered closest  
21 to the amino acids alanine (A) and glycine (G) (data not shown). The two amino acids with the  
22 smallest side chains might be least biased towards other biochemical features like charge and  
23 they are the 2<sup>nd</sup> (A) and 4<sup>th</sup> (G) most frequent amino acids in our training set (Table SOM\_1).  
24 Rare (O, U) and ambiguous amino acids (Z, B) were removed from the projection as their  
25 clustering showed that the model could not learn a reasonable embedding from the very small  
26 number of samples.

27 Crude structural classes as defined in SCOPe (Fig. 3b) were also captured by SeqVec  
28 embeddings. Although the embeddings were only trained to predict the next amino acid in a  
29 protein sequence, well separated clusters emerged from those embeddings in structure space.  
30 Especially, membrane proteins and small proteins formed distinct clusters (note: protein length  
31 did not affect the t-SNE grouping). Also, these results indicated that the embeddings captured  
32 complex relationships between proteins which are not directly observable from sequence  
33 similarity alone as SCOPe was redundancy reduced at 40% sequence identity. Therefore, the  
34 new embeddings could complement sequence-based structural classification as it was shown  
35 that the sequence similarity does not necessarily lead to structural similarity [54].

36 To further investigate the clusters which emerged from the SCOPe data set, we colored  
37 the same data set based on protein functions (Fig. 3c) and kingdoms (Fig. 3d). This analysis  
38 revealed that many of the small, distinct clusters emerged based on protein functions. For  
39 example, transferases and hydrolases form many small clusters. When increasing the level of  
40 abstraction by coloring the proteins according to their kingdoms, we observed certain clusters  
41 which are dominated by e.g. eukaryotes. Comparing the different views captured in panels 3b-  
42 3 revealed connections, e.g. that all-beta or small proteins dominate in eukaryotes (compare  
43 blue and orange islands in Fig. 3b with the same islands in Fig. 3d – colored blue to mark  
44 eukaryotes).

45



1

2 **Figure 3: Modeling the language of life.** 2D t-SNE projections of unsupervised SeqVec embeddings  
3 highlight different realities of proteins and their constituent parts, amino acids. Panels (b) to (d) are  
4 based on the same data set (Structural Classification of Proteins – extended (SCOPe) 2.07, redundancy  
5 reduced at 40%). For these plots, only subsets of SCOPe containing proteins with the annotation of  
6 interest (enzymatic activity (c) and kingdom (d)) may be displayed. Panel (a): the embedding space  
7 confirms: the 20 standard amino acids are clustered according to their biochemical and biophysical  
8 properties, i.e. hydrophobicity, charge or size. The unique role of Cysteine (C, mostly hydrophobic and  
9 polar) is conserved. Panel (b): SeqVec embeddings capture structural information as annotated in the  
10 main classes in SCOPe without ever having been explicitly trained on structural features. Panel (c):  
11 many small, local clusters share function as given by the main classes in the Enzyme Commission  
12 Number (E.C.). Panel (d): similarly, small, local clusters represent different kingdoms of life.

13

14 **CPU/GPU time used.** Due to the sequential nature of LSTMs, the time required to embed a  
15 protein grew linearly with protein length. Depending on the available main memory or GPU  
16 memory, this process could be massively parallelized. To optimally use available memory,  
17 batches are typically based on tokens rather than on sentences. Here, we sorted proteins  
18 according to their length and created batches of  $\leq 15K$  tokens that could still be handled by a  
19 single Nvidia GeForce GTX1080 with 8GB VRAM. On average the processing of a single  
20 protein took 0.027s when applying this batch-strategy to the NetSurfP-2.0 data set (average  
21 protein length: 256 residues, i.e. shorter than proteins for which 3D structure is not known).  
22 The batch with the shortest proteins (on average 38 residues, corresponding to 15% of the  
23 average protein length in the whole data set) required about one tenth (0.003s per protein, i.e.

1 11% of that for whole set). The batch containing the longest protein sequences in this data set  
2 (1578 residues on average, corresponding to 610% of average protein length in the whole data  
3 set), took about six times more (1.5s per protein, i.e. 556% of that for whole set). When creating  
4 SeqVec for the DeepLoc set (average length: 558 residues; as this set does not require a 3D  
5 structure, it provides a more realistic view on the distribution of protein lengths), the average  
6 processing time for a single protein was 0.08 with a minimum of 0.006 for the batch containing  
7 the shortest sequences (67 residues on average) and a maximum of 14.5s (9860 residues on  
8 average). Roughly, processing time was linear with respect to protein length. On a single Intel  
9 i7-6700 CPU with 64GB RAM, processing time increased by roughly 50% to 0.41s per protein,  
10 with a minimum and a maximum computation time of 0.06 and 15.3s, respectively. Compared  
11 to an average processing time of one hour for 1000 proteins when using evolutionary  
12 information directly [45], this implied an average speed up of 120-fold on a single GeForce  
13 GTX1080 and 9-fold on a single i7-6700 when predicting structural features; the inference time  
14 of DeepSeqVec for a single protein is on average 0.0028s.

## 16 Discussion

17 **ELMo alone did not suffice for top performance.** On the one hand, none of our  
18 implementations of ELMo reached anywhere near today's best (NetSurfP-2.0 for secondary  
19 structure and protein disorder and DeepLoc for localization and membrane protein  
20 classification; Fig. 1, Table 1, Table 2). Clearly, "just" using ELMo did not suffice to crack the  
21 challenges. On the other hand, some of our solutions appeared surprisingly competitive given  
22 the simplicity of the architectures. In particular for the per-protein predictions, for which SeqVec  
23 clearly outperformed the previously popular ProtVec [47] approach and even commonly used  
24 expert solutions (Fig. 1, Table 2: no method tested other than the top-of-the-line DeepLoc  
25 reached higher numerical values). For that comparison, we used the same data sets but could  
26 not rigorously compare standard errors which were unavailable for other methods. Estimating  
27 standard errors for our methods suggested that the differences were statistically significant as  
28 the difference was more than 7 sigmas for all methods (except DeepLoc (Q10=78) and iLoc-  
29 Euk(Q10=68)). The results for localization prediction implied that frequently used methods  
30 using evolutionary information (all marked with stars in Table 2) did not clearly outperform our  
31 simple ELMo-based tool (DeepSeqVec-Loc in Table 2). This was very different for the per-  
32 residue prediction tasks: here almost all top methods using evolutionary information  
33 numerically outperformed the simple model built on the ELMo embeddings (DeepSeqVec in  
34 Fig. 1 and Table 1). However, all models introduced in this work were deliberately designed to  
35 be relatively simple to demonstrate the predictive power of SeqVec. More sophisticated  
36 architectures using SeqVec are likely to outperform the approaches introduced here.

37 When we combined ELMo with evolutionary information for the per-residue predictions,  
38 the resulting tool still did not quite achieve top performance (Q3(NetSurfP-2.0)=85.3% vs.  
39 Q3(DeepProf + SeqVec)=82.4%, Table 1). This might suggest some limit for the usefulness of  
40 ELMo/SeqVec. However, it might also point to the more advanced solutions realized by  
41 NetSurfP-2.0 which applies two LSTMs of similar complexity as our entire system (including  
42 ELMo) on top of their last step leading to 35M (35 million) free parameters compared to about  
43 244K for DeepProf + SeqVec. Twenty times more free parameters might explain some fraction  
44 of the success, however due to limited GPU resources, we could not test how much.

1 Why did the ELMo-based approach improve more (relative to competition) for per-  
2 protein than for per-residue predictions? Per-protein data sets were over two orders of  
3 magnitude smaller than those for per-residue (simply because every protein constitutes one  
4 sample in the first and protein length samples for the second). It is likely that ELMo helped  
5 more for the smaller data sets because the unlabeled data is pre-processed so meaningful  
6 that less information needs to be learned by the ANN during per-protein prediction. This view  
7 was strongly supported by the t-SNE [52] results (Fig. 2): ELMo apparently had learned enough  
8 to realize a very rough clustering into localization and membrane/not.

9 We picked four particular tasks as proof-of-principle for our ELMo/SeqVec approach.  
10 These tasks were picked because recently developed methods implemented deep learning  
11 and the associated data sets for training and testing were made publicly available. We cannot  
12 imagine why our findings should not hold for other tasks of protein prediction and welcome the  
13 community to test our *SeqVec* for their particular tasks. We assume that our findings will be  
14 more relevant for small data sets than for large ones. For instance, we assume predictions of  
15 inter-residue contacts to improve less, and those for protein binding sites possibly more.

16  
17 **Good and fast predictions without using evolutionary information.** *SeqVec* predicted  
18 secondary structure and protein disorder over 100-times faster on a single 8GB GPU than the  
19 top-of-the-line prediction method NetSurfP-2.0 which requires to retrieve evolutionary  
20 information summarized in alignments. For some applications, the speedup might outweigh  
21 the reduction in performance. Therefore, embedding-based approaches such as *SeqVec*  
22 suggested a promising solution toward solving one of the biggest challenges for computational  
23 biology: How to efficiently handle the exponentially increasing number of sequences? Here,  
24 we showed that relevant information from large unannotated biological databases can be  
25 compressed into embeddings that condense and abstract the underlying biophysical  
26 principles. These embeddings, essentially the weights of a neural network, help as input to  
27 many problems for which smaller sets of annotated data are available (secondary structure,  
28 disorder, localization). Although the compression step needed to build the *SeqVec* model is  
29 very GPU-intensive, it can be performed in a centralized way using large clusters. After  
30 training, the model can be used as input by any consumer hardware. Such solutions are ideal  
31 to support researches without access to expensive cluster infrastructure.

32  
33 **Modeling the language of life?** Our ELMo implementation learned to model a probability  
34 distribution over a protein sequence. The sum over this probability distribution constituted a  
35 very informative input vector for any machine learning task. It also picked up context-  
36 dependent protein motifs without explicitly explaining what these motifs are relevant for. In  
37 contrast, tools such as ProtVec will always create the same vectors for a k-mer, regardless of  
38 the residues surrounding this k-mer in a particular protein sequence.

39 Our hypothesis had been that the ELMo embeddings learned from large databases of  
40 protein sequences (without annotations) could extract a *probabilistic model of the language of*  
41 *life* in the sense that the resulting system will extract aspects relevant both for per-residue and  
42 per-protein prediction tasks. All the results presented here have added independent evidence  
43 in full support of this hypothesis. For instance, the three state per-residue accuracy for  
44 secondary structure prediction improved by over eight percentage points through ELMo (Table  
45 1: e.g. for CB513: Q3(DeepSeqVec)=76.9% vs. Q3(DeepBLOSUM65)=67.4%, i.e. 9.4

percentage points corresponding to 14% of 67.4), the per-residue MCC for protein disorder prediction also rose (Table 1: e.g. TS115:  $MCC(\text{DeepSeqVec})=0.591$  vs.  $MCC(\text{DeepBLOSUM65})=0.488$ , corresponding to 18% of 0.488). On the per-protein level, the improvement over the previously popular tool extracting “meaning” from proteins, ProtVec, was even more substantial (Table 1: localization:  $Q10(\text{DeepSeqVec-Loc})=68\%$  vs.  $Q10(\text{DeepProtVec-Loc})=42\%$ , i.e. 62% rise over 42; membrane:  $Q2(\text{DeepSeqVec-Loc})=86.8\%$  vs.  $Q2(\text{DeepProtVec-Loc})=77.6\%$ , i.e. 19% rise over 77.6). We could demonstrate this reality even more directly using the t-SNE [52] results (Fig. 2 and Fig. 3): different levels of complexity ranging from single amino acids, over some localizations, structural features, functions and the classification of membrane/non-membrane had been implicitly learned by SeqVec without any training on such data. Clearly, our ELMo-driven implementation succeeded to model some aspects of the language of life as proxied for proteins.

## Conclusion

We have shown that it is possible to capture and transfer knowledge, e.g. biochemical or biophysical properties, from a large unlabeled data set of protein sequences to smaller, labelled data sets. In this first proof-of-principle, our comparably simple models have already reached promising performance for a variety of per-residue and per-protein prediction tasks obtainable from only single protein sequences as input, that is: without any direct evolutionary information, i.e. without alignments. This reduces the dependence on the time-consuming and computationally intensive calculation of protein profiles, allowing the prediction of per-residue and per-protein features of a whole proteome within less than an hour. For instance, on a single GeForce GTX 1080, the creation of embeddings and predictions of secondary structure and subcellular localization for the whole human proteome took about 32 minutes. Building more sophisticated architectures on top of the proposed SeqVec will increase sequence-based performance further.

Our new SeqVec embeddings may constitute an ideal starting point for many different applications in particular when labelled data are limited. The embeddings combined with evolutionary information might even improve over the best available methods, i.e. enable high-quality predictions. Alternatively, they might ease high-throughput predictions of whole proteomes when used as the only input feature. Alignment-free predictions bring speed and improvements for proteins for which alignments are not readily available or limited, such as for intrinsically disordered proteins, for the Dark Proteome, or for particular unique inventions of evolution. The trick was to tap into the potential of Deep Learning through transfer learning from large repositories of unlabeled data by modeling the language of life.

## Methods

**Data. UniRef50 training of SeqVec:** We trained ELMo on UniRef50 [32], a sequence redundancy-reduced subset of the UniProt database clustered at 50% pairwise sequence identity (PIDE). It contained 25 different letters (20 standard and 2 rare amino acids (U and O) plus 3 special cases describing either ambiguous (B, Z) or unknown amino acids (X); Table



1 SOM\_1) from 33M proteins with 9,577,889,953 residues. Each protein was treated as a  
2 sentence and each amino acid was interpreted as a single word. We referred to the resulting  
3 embedding as to *SeqVec* (Sequence-to-Vector).

4 Visualization of projection space: The current release of the “Structural Classification  
5 Of Proteins” (SCOPe, [53]) database (2.07) contains 14323 proteins at a redundancy level of  
6 40%. Functions encoded by the Enzyme Commission number (E.C., [55]) were retrieved via  
7 the “Structure Integration with Function, Taxonomy and Sequence” (SIFTS) mapping [56].  
8 SIFTS allows among other things a residue-level mapping between UniProt and PDB entries  
9 and a mapping from PDB identifiers to E.C.s. If no function annotation was available for a  
10 protein or if the same PDB identifier was assigned to multiple E.C.s, it was removed from Fig.  
11 3c. Taxonomic identifiers were used to map proteins to one of the 3 kingdoms of life or to  
12 viruses. Again, proteins were removed if no such information was available. The number of  
13 iterations for the t-SNE projections was set again to 3000 and the perplexity was adjusted  
14 (perplexity=5 for Fig. 3a and perplexity=30 for Fig. 3b-3d).

15 Per-residue level: secondary structure & intrinsic disorder (*NetSurfP-2.0*). To simplify  
16 compatibility, we used the data set published with a recent method seemingly achieving the  
17 top performance of the day in secondary structure prediction, namely *NetSurfP-2.0* [45].  
18 Performance values for the same data set exist also for other recent methods such as *Spider3*  
19 [57], *RaptorX* [58, 59] and *JPred4* [60]. The set contains 10,837 sequence-unique (at 25%  
20 PIDE) proteins of experimentally known 3D structures from the PDB [61] with a resolution of  
21 2.5 Å (0.25 nm) or better, collected by the PISCES server [62]. DSSP [44] assigned secondary  
22 structure and intrinsically disordered residues are flagged (residues without atomic  
23 coordinates, i.e. REMARK-465 in the PDB file). The original seven DSSP states (+ 1 for  
24 unknown) were mapped upon three states using the common convention: [G,H,I] → H (helix),  
25 [B,E] → E (strand), all others to O (other; often misleadingly referred to as *coil* or *loop*). As the  
26 authors of *NetSurfP-2.0* did not include the raw protein sequences in their public data set, we  
27 used the SIFTS file to obtain the original sequence. Only proteins with identical length in SIFTS  
28 and *NetSurfP-2.0* were used. This filtering step removed 56 sequences from the training set  
29 and three from the test sets (see below: two from CB513, one from CASP12 and none from  
30 TS115). We randomly selected 536 (~5%) proteins for early stopping (*cross-training*), leaving  
31 10,256 proteins for training. All published values referred to the following three test sets (also  
32 referred to as validation set): **TS115** [63]: 115 proteins from high-quality structures (<3Å)  
33 released after 2015 (and at most 30% PIDE to any protein of known structure in the PDB at  
34 the time); **CB513** [64]: 513 non-redundant sequences compiled 20 years ago (511 after SIFTS  
35 mapping); **CASP12** [65]: 21 proteins taken from the CASP12 free-modelling targets (20 after  
36 SIFTS mapping; all 21 fulfilled a stricter criterion toward non-redundancy than the two other  
37 sets; non-redundant with respect to all 3D structures known until May 2018 and all their  
38 relatives). Each of these sets covers different aspects of the secondary structure prediction  
39 problem: CB513 and TS115 only use structures determined by X-ray crystallography, applying  
40 similar cutoffs with respect to redundancy (30%) and resolution (2.5-3.0Å). While they serve  
41 as a good proxy for a baseline performance, CASP12 might better reflect the true  
42 generalization capability for unseen proteins including NMR and Cryo-EM as means of  
43 structure determination. Also, the strict redundancy reduction based on publication date  
44 reduces the bias towards well studied families. Nevertheless, toward our objective of  
45 establishing a proof-of-principle, these sets sufficed. All test sets had fewer than 25% PIDE to  
46 any protein used for training and cross-training (ascertained by the *NetSurfP-2.0* authors). To

1 compare methods using evolutionary information and those using our new word embeddings,  
2 we took the *HHblits* profiles published along with the NetSurfP-2.0 data set.

3 Per-protein level: localization & membrane proteins (DeepLoc). Localization prediction  
4 was trained and evaluated using the *DeepLoc* data set [46] for which performance was  
5 measured for several methods, namely: LocTree2 [66], MultiLoc2 [67], SherLoc2 [68], CELLO  
6 [69], iLoc-Euk [51], WoLF PSORT [70] and YLoc [71]. The data set contained proteins from  
7 UniProtKB/Swiss-Prot [48] (release: 2016\_04) with experimental annotation (code:  
8 ECO:0000269). The *DeepLoc* authors mapped these to ten classes, removing all proteins with  
9 multiple annotations. All these proteins were also classified into *water-soluble* or *membrane-*  
10 *bound* (or as *unknown* if the annotation was ambiguous). The resulting 13,858 proteins were  
11 clustered through PSI-CD-HIT [72, 73] (version 4.0; at 30% PIDE or  $\text{Eval} < 10^{-6}$ ). Adding the  
12 requirement that the alignment had to cover 80% of the shorter protein, yielded 8,464 clusters.  
13 This set was split into training and testing by using the same proteins for testing as the authors  
14 of *DeepLoc*. The training set was randomly sub-divided into 90% for training and 10% for  
15 determining early stopping (cross-training set).

16  
17 **ELMo terminology.** One-hot encoding (also known as *sparse encoding*) assigns each word  
18 (referred to as token in NLP) in the vocabulary an integer  $N$  used as the  $N$ th component of a  
19 vector with the dimension of the vocabulary size (number of different words). Each component  
20 is binary, i.e. either 0 if the word is not present in a sentence/text or 1 if it is. This encoding  
21 drove the first application of machine learning that clearly improved over all other methods in  
22 protein prediction [1-3]. TF-IDF represents tokens with the product of “frequency of token in  
23 data set” times “inverse frequency of token in document”. Thereby, rare tokens become more  
24 relevant than common words such as “the” (so called *stop words*). This concept resembles  
25 that of using  $k$ -mers for database searches [33], clustering [74], motifs [75, 76], and prediction  
26 methods [66, 70, 77-81]. Context-insensitive word embedding replaced expert features, such  
27 as TF-IDF, by algorithms that extracted such knowledge from unlabeled corpus such as  
28 Wikipedia, by either predicting the neighboring words, given the center word (skip-gram) or  
29 vice versa (CBOW). This became known in *Word2Vec* [49] and showcased for computational  
30 biology through *ProtVec* [49, 82]. More specialized implementations are *mut2vec* [83] learning  
31 mutations in cancer, and *phoscontext2vec* [84] identifying phosphorylation sites. The  
32 performance of context-insensitive approaches was pushed to its limits by adding sub-word  
33 information (FastText [85]) or global statistics on word co-occurrence (GloVe [86]). Context-  
34 sensitive word embedding started a new wave of word embedding techniques for NLP in 2018:  
35 the particular embedding renders the meaning of the phrase “*paper tiger*” dependent upon the  
36 context. Popular examples like ELMo [41] and Bert [87] have achieved state-of-the-art results  
37 in several NLP tasks. Both require substantial GPU computing power and time to be trained  
38 from scratch. However, in this work we focused on ELMo as it allows processing of sequences  
39 of variable length. This ELMo model consists of a single CharCNN [88] over the characters in  
40 a word and two layers of bidirectional LSTMs that introduce the context information of  
41 surrounding words. The CharCNN transforms all characters within a single word via an  
42 embedding layer into vector space and runs multiple CNNs of varying window size (here:  
43 ranging from 1 to 7) and number of filters (here: 32, 64, ..., 1024). In order to obtain a fixed-  
44 dimensional vector for each word, regardless of its length, the output of the CNNs is max-  
45 pooled and concatenated. The first bi-directional LSTM operates directly on the output of the  
46 CharCNN, while the second LSTM layer takes the output of the first LSTM as input. As

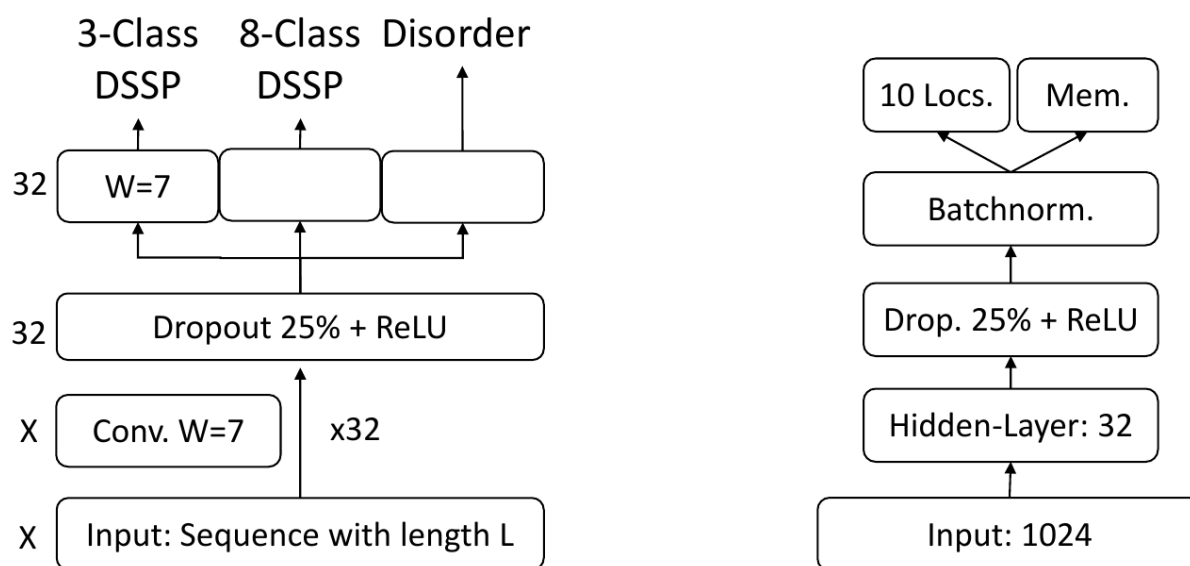
1 described in the original ELMo publication, the weights of the forward and the backward model  
2 are shared in order to reduce the memory overhead of the model and to combat overfitting.  
3 Even though, the risk of overfitting is small due to the high imbalance between number of  
4 trainable parameters (93M) versus number of tokens (9.3B), dropout at a rate of 10% was  
5 used to reduce the risk of overfitting. This model is trained to predict the next word given all  
6 previous words in a sentence. To the best of our knowledge, the context-sensitive ELMo has  
7 not been adapted to protein sequences, yet.

8  
9 **ELMo adaptation.** In order to allow more flexible models and easily integrate into existing  
10 solutions, we have used and generated ELMo as word embedding layers. No fine-tuning was  
11 performed on task-specific sequence sets. Thus, researchers could just replace their current  
12 embedding layer with our model to boost their task-specific performance. Furthermore, it will  
13 simplify the development of custom models that fit other use-cases. The embedding model  
14 takes a protein sequence of arbitrary length and returns 3076 features for each residue in the  
15 sequence. These 3076 features were derived by concatenating the outputs of the three internal  
16 layers of ELMo (1 CharCNN-layer, 2 LSTM-layers), each describing a token with a vector of  
17 length 1024. For simplicity, we summed the components of the three 1024-dimensional vectors  
18 to form a single 1024-dimensional feature vector describing each residue in a protein. In order  
19 to demonstrate the general applicability of *SeqVec*, we neither fine-tuned the model on a  
20 specific prediction task, nor optimized the combination of the three internal layers. Instead, we  
21 used the standard ELMo configuration [41] with the following changes: (i) reduction to 28  
22 tokens (20 standard and 2 rare (U,O) amino acids + 3 special tokens describing ambiguous  
23 (B,Z) or unknown (X) amino acids + 3 special tokens for ELMo indicating padded elements  
24 ('<MASK>') or the beginning ('<S>') or the end of a sequence ('</S>')), (ii) increase number of  
25 unroll steps to 100, (iii) decrease number of negative samples to 20, (iv) increase token number  
26 to 9,577,889,953.

27  
28 **SeqVec 2 prediction.** On the per-residue level, the predictive power of the new *SeqVec*  
29 embeddings was demonstrated by training a small two-layer Convolutional Neural Network  
30 (CNN) in PyTorch using a specific implementation [89] of the ADAM optimizer [90], cross-  
31 entropy loss, a learning rate of 0.001 and a batch size of 128 proteins. The first layer (in  
32 analogy to the sequence-to-structure network of earlier solutions [1, 2]) consisted of 32-filters  
33 each with a sliding window-size of  $w=7$ . The second layer (structure-to-structure [1, 2]) created  
34 the final predictions by applying again a CNN ( $w=7$ ) over the output of the first layer. These  
35 two layers were connected through a rectified linear unit (ReLU) and a dropout layer [91] with  
36 a dropout-rate of 25% (Fig. 4, left panel). This simple architecture was trained independently  
37 on six different types of input, each with a different number of free parameters. (i) DeepProf  
38 (14,000=14k free parameters): Each residue was described by a vector of size 50 which  
39 included a one-hot encoding (20 features), the profiles of evolutionary information (20 features)  
40 from HHblits as published previously [45], the state transition probabilities of the Hidden-  
41 Markov-Model (7 features) and 3 features describing the local alignment diversity. (ii)  
42 DeepSeqVec (232k free parameters): Each protein sequence is represented by the output  
43 from *SeqVec*. The resulting embedding described each residue as a 1024-dimensional vector.  
44 (iii) DeepProf+SeqVec (244k free parameters): This model simply concatenated the input  
45 vectors used in (i) and (ii). (iv) DeepProtVec (25k free parameters): Each sequence was split  
46 into overlapping 3-mers each represented by a 100-dimensional ProtVec [47]. (v) DeepOneHot

1 (7k free parameters): The 20 amino acids were encoded as one-hot vectors as described  
 2 above. Rare amino acids were mapped to vectors with all components set to 0. Consequently,  
 3 each protein residue was encoded as a 20-dimensional one-hot vector. (vi) DeepBLOSUM65  
 4 (8k free parameters): Each protein residue was encoded by its BLOSUM65 substitution matrix  
 5 [92]. In addition to the 20 standard amino acids, BLOSUM65 also contains substitution scores  
 6 for the special cases B, Z (ambiguous) and X (unknown), resulting in a feature vector of length  
 7 23 for each residue.

8 On the per-protein level, a simple feed-forward neural network was used to  
 9 demonstrate the power of the new embeddings. In order to ensure equal-sized input vectors  
 10 for all proteins, we averaged over the embeddings of all residues in a given protein resulting  
 11 in a 1024-dimensional vector representing any protein in the data set. ProtVec representations  
 12 were derived the same way, resulting in a 100-dimensional vector. These vectors (either 100-  
 13 or 1024 dimensional) were first compressed to 32 features, then dropout with a dropout rate  
 14 of 25%, batch normalization [93] and a rectified linear Unit (ReLU) were applied before the  
 15 final prediction (Fig. 4, right panel). In the following, we refer to the models trained on the two  
 16 different input types as (i) DeepSeqVec-Loc (33k free parameters): average over SeqVec  
 17 embedding of a protein as described above and (ii) DeepProtVec-Loc (320 free parameters):  
 18 average over ProtVec embedding of a protein. We used the following hyper-parameters:  
 19 learning rate: 0.001, Adam optimizer with cross-entropy loss, batch size: 64. The losses of the  
 20 individual tasks were summed before backpropagation. Due to the relatively small number of  
 21 free parameters in our models, the training of all networks completed on a single Nvidia  
 22 GeForce GTX1080 within a few minutes (11 seconds for DeepSeqVec-Loc, 15 minutes for  
 23 DeepSeqVec).



24  
 25 **Figure 4:** On the left the architecture of the model used for the per-residue level predictions  
 26 (secondary structure and disorder) is sketched, on the right that used for per-protein level predictions  
 27 (localization and membrane/not membrane). The 'X', on the left, indicates that different input features  
 28 corresponded to a difference in the number of input channels, e.g. 1024 for *SeqVec* or 50 for profile-  
 29 based input. The letter 'W' refers to the window size of the corresponding convolutional layer (W=7  
 30 implies a convolution of size 7x1).

1

2 **Evaluation measures.** To simplify comparisons, we ported the evaluation measures from the  
3 publications we derived our data sets from, i.e. those used to develop *NetSurfP-2.0* [45] and  
4 *DeepLoc* [46]. All numbers reported constituted averages over all proteins in the final test sets.  
5 This work aimed at a proof-of-principle that the *SeqVec* embedding contain predictive  
6 information. In the absence of any claim for state-of-the-art performance, we did not calculate  
7 any significance values for the reported values.

8 Per-residue performance: Toward this end, we used the standard three-state per-  
9 residue accuracy (Q3=percentage correctly predicted in either helix, strand, other [1]) along  
10 with its eight-state analog (Q8). Predictions of intrinsic disorder were evaluated through the  
11 Matthew's correlation coefficient (MCC [94]) and the False-Positive Rate (FPR) representative  
12 for tasks with high class imbalance. For completeness, we also provided the entire confusion  
13 matrices for both secondary structure prediction problems (Fig. SOM\_2). Standard errors were  
14 calculated over the distribution of each performance measure for all proteins.

15 Per-protein performance: The predictions whether a protein was membrane-bound or  
16 water-soluble were evaluated by calculating the two-state per set accuracy (Q2: percentage of  
17 proteins correctly predicted), and the MCC. A generalized MCC using the Gorodkin measure  
18 [95] for K (=10) categories as well as accuracy (Q10), was used to evaluate localization  
19 predictions. Standard errors were calculated using 1000 bootstrap samples, each chosen  
20 randomly by selecting a sub-set of the predicted test set that had the same size (draw with  
21 replacement).

22

## 23 **Availability**

24 The pre-trained ELMo-like *SeqVec* model and a description on how to implement the  
25 embeddings into existing methods can be found here: <https://github.com/Rostlab/SeqVec> .  
26 Accessed 10<sup>th</sup> September 2019.

27 Predictions on secondary structure, disorder and subcellular localization based on *SeqVec* can  
28 be accessed under: <https://embed.protein.properties> . Accessed 10<sup>th</sup> September 2019.

29 The *NetSurfP-2.0* data set [45] used for the evaluation of *SeqVec* on the task of secondary  
30 structure and disorder prediction are publicly available under:  
31 <http://www.cbs.dtu.dk/services/NetSurfP/> . Accessed 10<sup>th</sup> September 2019.

32 The *DeepLoc* data set [46] used for the evaluation of *SeqVec* on the task of subcellular  
33 localization prediction are publicly available under:  
34 <http://www.cbs.dtu.dk/services/DeepLoc/data.php> . Accessed 10<sup>th</sup> September 2019.

35

## 36 **Funding**

37 This work was supported by a grant from the Alexander von Humboldt foundation through the  
38 German Ministry for Research and Education (BMBF: Bundesministerium fuer Bildung und  
39 Forschung) as well as by a grant from Deutsche Forschungsgemeinschaft (DFG–GZ:



1 RO1320/4–1). We gratefully acknowledge the support of NVIDIA Corporation with the donation  
2 of two Titan GPU used for this research. We also want to thank the LRZ (Leibniz  
3 Rechenzentrum) for providing us access to DGX-V1.

4

5

## Authors contributions

6 AE and MH suggested to use ELMo for modeling protein sequences. AE adopted and trained  
7 ELMo. MH evaluated SeqVec embeddings on different data sets and tasks. YW helped with  
8 discussions about natural language processing. CD implemented the web-interface which  
9 allows to access and visualize the predictions. DN helped with various problems regarding the  
10 code. FM and BR helped with the design of the experiment and to critically improve the  
11 manuscript. MH and AE drafted the manuscript and the other authors provided feedback. All  
12 authors read and approved the final manuscript.

13

14

## Acknowledgements

15 The authors thank primarily Tim Karl for invaluable help with hardware and software and Inga  
16 Weise for support with many other aspects of this work. Last, not least, thanks to all those who  
17 deposit their experimental data in public databases, and to those who maintain these  
18 databases.

19

## References

- 1
- 2
- 3 1. Rost B, Sander C: **Prediction of protein secondary structure at better than 70% accuracy.** *Journal of Molecular Biology* 1993, **232**:584-599.
- 4
- 5 2. Rost B, Sander C: **Improved prediction of protein secondary structure by use of sequence profiles and neural networks.** *Proceedings of the National Academy of Sciences* 1993, **90**:7558-7562.
- 6
- 7
- 8 3. Rost B, Sander C: **Jury returns on structure prediction.** *Nature* 1992, **360**:540.
- 9 4. Rost B, Sander C: **Combining evolutionary information and neural networks to predict protein secondary structure.** *Proteins: Structure, Function, and Genetics* 1994, **19**:55-72.
- 10
- 11
- 12 5. Solovyev VV, Salamov AA: **Predicting a-helix and b-strand segments of globular proteins.** *Computer Applications in Biological Science* 1994, **10**:661-669.
- 13
- 14 6. Barton GJ: **Protein secondary structure prediction.** *Current Opinion in Structural Biology* 1995, **5**:372-376.
- 15
- 16 7. Chandonia J-M, Karplus M: **Neural networks for secondary structure and structural class predictions.** *Protein Science* 1995, **4**:275-285.
- 17
- 18 8. Mehta PK, Heringa J, Argos P: **A simple and fast approach to prediction of protein secondary structure from multiply aligned sequences with accuracy above 70%.** *Protein Science* 1995, **4**:2517-2525.
- 19
- 20
- 21 9. Frishman D, Argos P: **Knowledge-based protein secondary structure assignment.** *Proteins: Structure, Function, and Genetics* 1995, **23**:566-579.
- 22
- 23 10. Jones DT: **Protein secondary structure prediction based on position-specific scoring matrices.** *Journal of Molecular Biology* 1999, **292**(2):195-202.
- 24
- 25 11. Rost B, Casadio R, Fariselli P, Sander C: **Transmembrane helix prediction at 95% accuracy.** *Protein Science* 1995, **4**:521-533.
- 26
- 27 12. Rost B, Casadio R, Fariselli P: **Topology prediction for helical transmembrane proteins at 86% accuracy.** *Protein Science* 1996, **5**:1704-1718.
- 28
- 29 13. Bigelow H, Petrey D, Liu J, Przybylski D, Rost B: **Predicting transmembrane beta-barrels in proteomes.** *Nucleic Acids Research* 2004, **32**:2566-2577.
- 30
- 31 14. Rost B, Sander C: **Conservation and prediction of solvent accessibility in protein families.** *Proteins: Structure, Function, and Genetics* 1994, **20**(3):216-226.
- 32
- 33 15. Radivojac P, Obradovic Z, Smith DK, Zhu G, Vucetic S, Brown CJ, Lawson JD, Dunker AK: **Protein flexibility and intrinsic disorder.** *Protein Science* 2004, **13**:71-80.
- 34
- 35 16. Schlessinger A, Rost B: **Protein flexibility and rigidity predicted from sequence.** *Proteins: Structure, Function, and Bioinformatics* 2005, **61**(1):115-126.
- 36
- 37 17. Punta M, Rost B: **PROFcon: novel prediction of long-range contacts.** *Bioinformatics* 2005, **21**(13):2960-2968.
- 38
- 39 18. Peng K, Vucetic S, Radivojac P, Brown CJ, Dunker AK, Obradovic Z: **Optimizing long intrinsic disorder predictors with protein evolutionary information.** *J Bioinform Comput Biol* 2005, **3**(1):35-60.
- 40
- 41
- 42 19. Schlessinger A, Liu J, Rost B: **Natively unstructured loops differ from other loops.** *PLoS Computational Biology* 2007, **3**(7):e140.
- 43

- 1 20. Schlessinger A, Punta M, Rost B: **Natively unstructured regions in proteins**  
2 **identified from contact predictions.** *Bioinformatics* 2007, **23**(18):2376-2384.
- 3 21. Nair R, Rost B: **Better prediction of sub-cellular localization by combining**  
4 **evolutionary and structural information.** *Proteins: Structure, Function, and*  
5 *Bioinformatics* 2003, **53**(4):917-930.
- 6 22. Nair R, Rost B: **Mimicking cellular sorting improves prediction of subcellular**  
7 **localization.** *Journal of Molecular Biology* 2005, **348**(1):85-100.
- 8 23. Ofra Y, Rost B: **Protein-protein interaction hot spots carved into sequences.**  
9 *PLoS Computational Biology* 2007, **3**(7):e119.
- 10 24. Ofra Y, Rost B: **ISIS: Interaction Sites Identified from Sequence.** *Bioinformatics*  
11 2007, **23**(2):e13-16.
- 12 25. Marino Buslje C, Teppa E, Di Domenico T, Delfino JM, Nielsen M: **Networks of high**  
13 **mutual information define the structural proximity of catalytic sites: implications**  
14 **for catalytic residue identification.** *PLoS Comput Biol* 2010, **6**(11):e1000978.
- 15 26. Bromberg Y, Rost B: **SNAP: predict effect of non-synonymous polymorphisms on**  
16 **function.** *Nucleic Acids Research* 2007, **35**(11):3823-3835.
- 17 27. Adzhubei IA, Schmidt S, Peshkin L, Ramensky VE, Gerasimova A, Bork P, Kondrashov  
18 AS, Sunyaev SR: **A method and server for predicting damaging missense**  
19 **mutations.** *Nat Methods* 2010, **7**(4):248-249.
- 20 28. Hayat S, Sander C, Marks DS, Elofsson A: **All-atom 3D structure prediction of**  
21 **transmembrane beta-barrel proteins from sequences.** *Proceedings of the National*  
22 *Academy of Sciences of the United States of America* 2015, **112**(17):5413-5418.
- 23 29. Marks DS, Hopf TA, Sander C: **Protein structure prediction from sequence**  
24 **variation.** *Nat Biotechnol* 2012, **30**(11):1072-1080.
- 25 30. Marks DS, Colwell LJ, Sheridan R, Hopf TA, Pagnani A, Zecchina R, Sander C: **Protein**  
26 **3D structure computed from evolutionary sequence variation.** *PLoS One* 2011,  
27 **6**(12):e28766.
- 28 31. Morcos F, Pagnani A, Lunt B, Bertolino A, Marks DS, Sander C, Zecchina R, Onuchic  
29 JN, Hwa T, Weigt M: **Direct-coupling analysis of residue coevolution captures**  
30 **native contacts across many protein families.** *Proc Natl Acad Sci U S A* 2011,  
31 **108**(49):E1293-1301.
- 32 32. Suzek BE, Wang Y, Huang H, McGarvey PB, Wu CH, UniProt C: **UniRef clusters: a**  
33 **comprehensive and scalable alternative for improving sequence similarity**  
34 **searches.** *Bioinformatics* 2015, **31**(6):926-932.
- 35 33. Altschul SF, Madden TL, Schaeffer AA, Zhang J, Zhang Z, Miller W, Lipman DJ:  
36 **Gapped Blast and PSI-Blast: a new generation of protein database search**  
37 **programs.** *Nucleic Acids Research* 1997, **25**:3389-3402.
- 38 34. Remmert M, Biegert A, Hauser A, Soding J: **HHblits: lightning-fast iterative protein**  
39 **sequence searching by HMM-HMM alignment.** *Nat Methods* 2012, **9**(2):173-175.
- 40 35. Steinegger M, Meier M, Mirdita M, Voehringer H, Haunsberger SJ, Soeding J: **HH-**  
41 **suite3 for fast remote homology detection and deep protein annotation.** *bioRxiv*  
42 2019:560029.
- 43 36. Steinegger M, Soding J: **MMseqs2 enables sensitive protein sequence searching**  
44 **for the analysis of massive data sets.** *Nature biotechnology* 2017, **35**(11):1026.





- 1 37. Uversky VN, Radivojac P, Iakoucheva LM, Obradovic Z, Dunker AK: **Prediction of**  
2 **intrinsic disorder and its use in functional proteomics.** *Methods Mol Biol* 2007,  
3 **408**:69-92.
- 4 38. Dunker AK, Babu MM, Barbar E, Blackledge M, Bondos SE, Dosztanyi Z, Dyson HJ,  
5 Forman-Kay J, Fuxreiter M, Gsponer J *et al.*: **What's in a name? Why these proteins**  
6 **are intrinsically disordered.** *Intrinsically Disord Proteins* 2013, 1(1):e24157.
- 7 39. Perdigo N, Heinrich J, Stolte C, Sabir KS, Buckley MJ, Tabor B, Signal B, Gloss BS,  
8 Hammang CJ, Rost B *et al.*: **Unexpected features of the dark proteome.** *Proceedings*  
9 *of the National Academy of Sciences of the United States of America* 2015.
- 10 40. Schafferhans A, O'Donoghue SI, Heinzinger M, Rost B: **Dark Proteins Important for**  
11 **Cellular Function.** *Proteomics* 2018, **18**(21-22):e1800227.
- 12 41. Peters ME, Neumann M, Iyyer M, Gardner M, Clark C, Lee K, Zettlemoyer L: **Deep**  
13 **contextualized word representations.** *arXiv* 2018, arXiv:1802.05365.
- 14 42. Schils E, Haan Pd: **Characteristics of sentence length in running text.** *Literary and*  
15 *linguistic computing* 1993, **8**(1):20-26.
- 16 43. Hochreiter S, Schmidhuber J: **Long short-term memory.** *Neural Computation* 1997,  
17 **9**(8):1735-1780.
- 18 44. Kabsch W, Sander C: **Dictionary of protein secondary structure: pattern**  
19 **recognition of hydrogen bonded and geometrical features.** *Biopolymers* 1983,  
20 **22**:2577-2637.
- 21 45. Klausen MS, Jespersen MC, Nielsen H, Jensen KK, Jurtz VI, Sonderby CK, Sommer  
22 MOA, Winther O, Nielsen M, Petersen B *et al.*: **NetSurfP-2.0: Improved prediction of**  
23 **protein structural features by integrated deep learning.** *Proteins* 2019.
- 24 46. Almagro Armenteros JJ, Sonderby CK, Sonderby SK, Nielsen H, Winther O: **DeepLoc:**  
25 **prediction of protein subcellular localization using deep learning.** *Bioinformatics*  
26 2017, **33**(24):4049.
- 27 47. Asgari E, Mofrad MR: **Continuous Distributed Representation of Biological**  
28 **Sequences for Deep Proteomics and Genomics.** *PLoS One* 2015,  
29 **10**(11):e0141287.
- 30 48. Boutet E, Lieberherr D, Tognolli M, Schneider M, Bansal P, Bridge AJ, Poux S,  
31 Bougueleret L, Xenarios I: **UniProtKB/Swiss-Prot, the Manually Annotated Section**  
32 **of the UniProt KnowledgeBase: How to Use the Entry View.** *Methods Mol Biol*  
33 2016, **1374**:23-54.
- 34 49. Mikolov T, Chen K, Corrado G, Dean J: **Efficient estimation of word representations**  
35 **in vector space.** *ArXiv* 2013:arXiv:1301.3781.
- 36 50. Rives A, Goyal S, Meier J, Guo D, Ott M, Zitnick CL, Ma J, Fergus R: **Biological**  
37 **structure and function emerge from scaling unsupervised learning to 250 million**  
38 **protein sequences.** *bioRxiv* 2019:622803.
- 39 51. Chou KC, Wu ZC, Xiao X: **iLoc-Euk: a multi-label classifier for predicting the**  
40 **subcellular localization of singleplex and multiplex eukaryotic proteins.** *PLoS*  
41 *One* 2011, **6**(3):e18258.
- 42 52. Maaten Lvd, Hinton G: **Visualizing data using t-SNE.** *Journal of machine learning*  
43 *research* 2008, **9**(Nov):2579-2605.

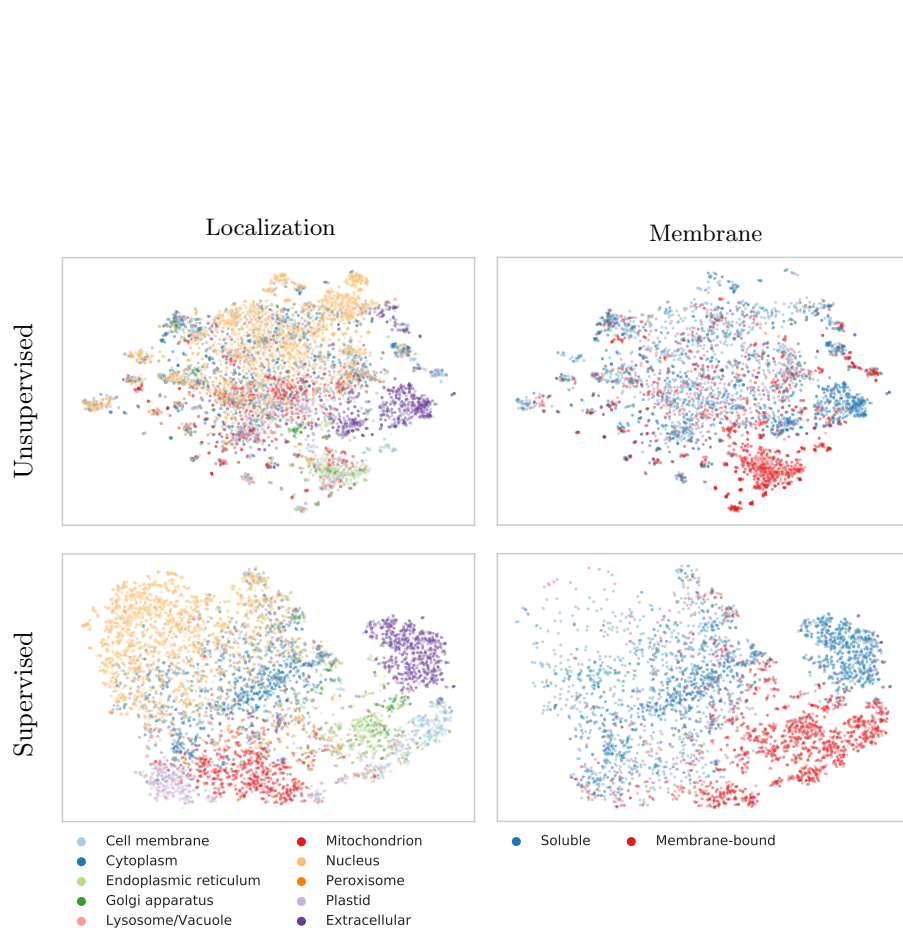
- 1 53. Fox NK, Brenner SE, Chandonia J-M: **SCOPE: Structural Classification of**  
2 **Proteins—extended, integrating SCOP and ASTRAL data and classification of**  
3 **new structures.** *Nucleic acids research* 2013, **42**(D1):D304-D309.
- 4 54. Kosloff M, Kolodny R: **Sequence-similar, structure-dissimilar protein pairs in the**  
5 **PDB.** *Proteins: Structure, Function, and Bioinformatics* 2008, **71**(2):891-902.
- 6 55. Bairoch A: **The ENZYME database in 2000.** *Nucleic acids research* 2000, **28**(1):304-  
7 305.
- 8 56. Velankar S, Dana JM, Jacobsen J, Van Ginkel G, Gane PJ, Luo J, Oldfield TJ,  
9 O'donovan C, Martin M-J, Kleywegt GJ: **SIFTS: structure integration with function,**  
10 **taxonomy and sequences resource.** *Nucleic acids research* 2012, **41**(D1):D483-  
11 D489.
- 12 57. Heffernan R, Yang Y, Paliwal K, Zhou Y: **Capturing non-local interactions by long**  
13 **short-term memory bidirectional recurrent neural networks for improving**  
14 **prediction of protein secondary structure, backbone angles, contact numbers**  
15 **and solvent accessibility.** *Bioinformatics* 2017, **33**(18):2842-2849.
- 16 58. Wang S, Peng J, Ma J, Xu J: **Protein secondary structure prediction using deep**  
17 **convolutional neural fields.** *Scientific reports* 2016, **6**:18962.
- 18 59. Wang S, Li W, Liu S, Xu J: **RaptorX-Property: a web server for protein structure**  
19 **property prediction.** *Nucleic acids research* 2016, **44**(W1):W430-W435.
- 20 60. Drozdetskiy A, Cole C, Procter J, Barton GJ: **JPred4: a protein secondary structure**  
21 **prediction server.** *Nucleic acids research* 2015, **43**(W1):W389-W394.
- 22 61. Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, Shindyalov IN,  
23 Bourne PE: **The protein data bank.** *Nucleic acids research* 2000, **28**(1):235-242.
- 24 62. Wang G, Dunbrack Jr RL: **PISCES: a protein sequence culling server.**  
25 *Bioinformatics* 2003, **19**(12):1589-1591.
- 26 63. Yang Y, Gao J, Wang J, Heffernan R, Hanson J, Paliwal K, Zhou Y: **Sixty-five years**  
27 **of the long march in protein secondary structure prediction: the final stretch?**  
28 *Briefings in bioinformatics* 2016, **19**(3):482-494.
- 29 64. Cuff JA, Barton GJ: **Evaluation and improvement of multiple sequence methods**  
30 **for protein secondary structure prediction.** *Proteins: Structure, Function, and*  
31 *Genetics* 1999, **34**(4):508-519.
- 32 65. Abriata LA, Tamò GE, Monastyrskyy B, Kryshchak A, Dal Peraro M: **Assessment**  
33 **of hard target modeling in CASP12 reveals an emerging role of alignment-based**  
34 **contact prediction methods.** *Proteins: Structure, Function, and Bioinformatics* 2018,  
35 **86**:97-112.
- 36 66. Goldberg T, Hamp T, Rost B: **LocTree2 predicts localization for all domains of life.**  
37 *Bioinformatics* 2012, **28**(18):i458-i465.
- 38 67. Blum T, Briesemeister S, Kohlbacher O: **MultiLoc2: integrating phylogeny and Gene**  
39 **Ontology terms improves subcellular protein localization prediction.** *BMC*  
40 *Bioinformatics* 2009, **10**:274.
- 41 68. Briesemeister S, Blum T, Brady S, Lam Y, Kohlbacher O, Shatkay H: **SherLoc2: a**  
42 **high-accuracy hybrid method for predicting subcellular localization of proteins.**  
43 *J Proteome Res* 2009, **8**(11):5363-5366.
- 44 69. Yu CS, Chen YC, Lu CH, Hwang JK: **Prediction of protein subcellular localization.**  
45 *Proteins* 2006, **64**(3):643-651.




- 1 70. Horton P, Park KJ, Obayashi T, Fujita N, Harada H, Adams-Collier CJ, Nakai K: **WoLF**  
2 **PSORT: protein localization predictor**. *Nucleic Acids Res* 2007, **35**(Web Server  
3 issue):W585-587.
- 4 71. Briesemeister S, Rahnenfuhrer J, Kohlbacher O: **YLoc - an interpretable web server**  
5 **for predicting subcellular localization**. *Nucleic Acids Res* 2010, **38** Suppl:W497-  
6 502.
- 7 72. Li W, Godzik A: **Cd-hit: a fast program for clustering and comparing large sets of**  
8 **protein or nucleotide sequences**. *Bioinformatics* 2006, **22**(13):1658-1659.
- 9 73. Fu L, Niu B, Zhu Z, Wu S, Li W: **CD-HIT: accelerated for clustering the next-**  
10 **generation sequencing data**. *Bioinformatics* 2012, **28**(23):3150-3152.
- 11 74. Moussa M, Mandoiu, Il: **Single cell RNA-seq data clustering using TF-IDF based**  
12 **methods**. *BMC Genomics* 2018, **19**(Suppl 6):569.
- 13 75. Bailey TL, Boden M, Buske FA, Frith M, Grant CE, Clementi L, Ren J, Li WW, Noble  
14 WS: **MEME SUITE: tools for motif discovery and searching**. *Nucleic Acids Res*  
15 2009, **37**(Web Server issue):W202-208.
- 16 76. Bernard G, Chan CX, Ragan MA: **Alignment-free microbial phylogenomics under**  
17 **scenarios of sequence divergence, genome rearrangement and lateral genetic**  
18 **transfer**. *Sci Rep* 2016, **6**:28970.
- 19 77. Nakai K, Horton P: **PSORT: a program for detecting sorting signals in proteins**  
20 **and predicting their subcellular localization**. *Trends Biochem Sci* 1999, **24**(1):34-  
21 36.
- 22 78. Kuang R, le E, Wang K, Wang K, Siddiqi M, Freund Y, Leslie C: **Profile-based string**  
23 **kernels for remote homology detection and motif extraction**. *J Bioinform Comput*  
24 *Biol* 2005, **3**(3):527-550.
- 25 79. Leslie C, Eskin E, Weston J, Noble WS: **Mismatch string kernels for SVM protein**  
26 **classification**. *Bioinformatics* 2003:in press.
- 27 80. Noble WS, Kuang R, Leslie C, Weston J: **Identifying remote protein homologs by**  
28 **network propagation**. *FEBS J* 2005, **272**(20):5119-5128.
- 29 81. Hamp T, Rost B: **Evolutionary profiles improve protein-protein interaction**  
30 **prediction from sequence**. *Bioinformatics* 2015, **31**(12):1945-1950.
- 31 82. Asgari E, McHardy AC, Mofrad MRK: **Probabilistic variable-length segmentation of**  
32 **protein sequences for discriminative motif discovery (DiMotif) and sequence**  
33 **embedding (ProtVecX)**. *Sci Rep* 2019, **9**(1):3577.
- 34 83. Kim S, Lee H, Kim K, Kang J: **Mut2Vec: distributed representation of cancerous**  
35 **mutations**. *BMC medical genomics* 2018, **11**(2):33.
- 36 84. Xu Y, Song J, Wilson C, Whisstock JC: **PhosContext2vec: a distributed**  
37 **representation of residue-level sequence contexts and its application to general**  
38 **and kinase-specific phosphorylation site prediction**. *Scientific reports* 2018, **8**.
- 39 85. Bojanowski P, Grave E, Joulin A, Mikolov T: **Enriching word vectors with subword**  
40 **information**. *Transactions of the Association for Computational Linguistics* 2017,  
41 **5**:135-146.
- 42 86. Pennington J, Socher R, Manning C: **Glove: Global vectors for word representation**.  
43 In: *Proceedings of the 2014 conference on empirical methods in natural language*  
44 *processing (EMNLP): 2014*. 1532-1543.

- 1 87. Devlin J, Chang M-W, Lee K, Toutanova K: **Bert: Pre-training of deep bidirectional**  
2 **transformers for language understanding**. *arXiv preprint arXiv:181004805* 2018.
- 3 88. Kim Y, Jernite Y, Sontag D, Rush AM: **Character-aware neural language models**. In:  
4 *Thirtieth AAAI Conference on Artificial Intelligence: 2016*.
- 5 89. Reddi SJ, Kale S, Kumar S: **On the convergence of adam and beyond**. *arXiv preprint*  
6 *arXiv:190409237* 2019.
- 7 90. Kingma DP, Ba J: **Adam: A method for stochastic optimization**. *arXiv preprint*  
8 *arXiv:1412.6980* 2014.
- 9 91. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R: **Dropout: a**  
10 **simple way to prevent neural networks from overfitting**. *The Journal of Machine*  
11 *Learning Research* 2014, **15**(1):1929-1958.
- 12 92. Henikoff S, Henikoff JG: **Amino acid substitution matrices from protein blocks**.  
13 *PNAS* 1992, **89**:10915-10919.
- 14 93. Ioffe S, Szegedy C: **Batch normalization: Accelerating deep network training by**  
15 **reducing internal covariate shift**. *arXiv preprint arXiv:1502.03167* 2015.
- 16 94. Matthews BW: **Comparison of the predicted and observed secondary structure of**  
17 **T4 phage lysozyme**. *Biochimica et Biophysica Acta* 1975, **405**:442-451.
- 18 95. Gorodkin J: **Comparing two K-category assignments by a K-category correlation**  
19 **coefficient**. *Computational biology and chemistry* 2004, **28**(5-6):367-374.
- 20






(a) Amino acids




- ▲ Hydrophobic (aromatic)
- ▼ Hydrophobic (aliphatic)
- ⊕ Positive
- ⊖ Negative
- Polar neutral
- ◆ Special cases
- Small (<130 Dalton)
- Medium
- Big (>150 Dalton)

(b) SCOPe




- All alpha
- All beta
- Alpha & beta (a|b)
- Alpha & beta (a+b)
- Multi-domain
- Membrane, cell surface
- Small proteins

(c) Enzyme Commission number (E.C.)



- Oxidoreductases
- Transferases
- Hydrolases
- Lyases
- Isomerases
- Ligases
- Translocases

(d) Kingdoms



- Eukaryota
- Bacteria
- Archaea
- Viruses



