

# Domain-PFP: Protein Function Prediction Using Function-Aware Domain Embedding Representations

Nabil Ibtehaz<sup>1</sup>, Yuki Kagaya<sup>2</sup>, & Daisuke Kihara<sup>1,2,\*</sup>

1 Department of Computer Science, Purdue University, West Lafayette, IN, United States

2 Department of Biological Sciences, Purdue University, West Lafayette, IN, United States

Corresponding Author, [dkihara@purdue.edu](mailto:dkihara@purdue.edu)

**Keywords:** self-supervised learning, protein sequence embedding, protein function prediction, gene ontology, deep learning

## Abstract

Domains are functional and structural units of proteins that govern various biological functions performed by the proteins. Therefore, the characterization of domains in a protein can serve as a proper functional representation of proteins. Here, we employ a self-supervised protocol to derive functionally consistent representations for domains by learning domain-Gene Ontology (GO) co-occurrences and associations. The domain embeddings we constructed turned out to be effective in performing actual function prediction tasks. Extensive evaluations showed that protein representations using the domain embeddings are superior to those of large-scale protein language models in GO prediction tasks. Moreover, the new function prediction method built on the domain embeddings, named Domain-PFP, significantly outperformed the state-of-the-art function predictors. Additionally, Domain-PFP demonstrated competitive performance in the CAFA3 evaluation, achieving overall the best performance among the top teams that participated in the assessment.

## Introduction

Protein function prediction is one of the long-standing, fundamental topics of bioinformatics, which involves profiling the activities and interactions of proteins<sup>1</sup>. Although protein functions are eventually determined by experiments, the experimental effort and expense slow down the process of function discovery, which is in contrast to the ever-increasing volume of sequenced proteins<sup>2</sup>. At present, not even 1% of sequenced proteins have functional annotation<sup>3</sup>. Unlike relatively cheaper sequencing technologies, there is a deficit of scalable, high-throughput experimental assays to functionally annotate proteins<sup>4</sup>. This has led to the demand for in-silico methods of automated protein function prediction<sup>5</sup>. Protein functions have been determined naturally from sequence similarity to known proteins<sup>6</sup> and other characteristics of proteins that can trace functional relevance. Such information includes structural configuration<sup>7-9</sup>, phylogenetic information<sup>10,11</sup>, domain distribution<sup>12-14</sup>, protein networks<sup>3,15</sup>, and combinations of multiple sources<sup>16,17</sup>. Recently, various deep learning-based methods were proposed to learn a functional representation of proteins<sup>8,16,18-22</sup>. Such methods demonstrated substantial improvement over traditional database search-based methods<sup>23,24</sup>.

Proteins consist of domains, which are functional and structural units responsible for specific functions and interactions<sup>25</sup>. Therefore, it is compelling to infer the functions of a protein based on the presence and distribution of the various domains in it. InterPro2GO is an ongoing project that assigns GO annotations to specific domains in the InterPro database, and this annotation is done manually by experts<sup>26,27</sup>. Although the domain-GO mapping by InterPro2GO provides curated information on protein function, the coverage is severely limited. For example, there are approximately 38k InterPro entries and 48k GO terms, but the current version of InterPro2GO (version-date: 2022/03/16) mapping only includes 16,443 unique InterPro entries and 6,482 GO terms. Despite the lack of annotations, several methods have tried to leverage protein domain information for function prediction. Messih et al. analyzed the recurrence and order of protein domains and their influence on protein functions<sup>13</sup>. Rojano et al. attempted to associate domains and functions through tripartite graphs<sup>14</sup>. Besides such domain-focused studies, protein domains have been consistently used as a source of complementary functional information in a number of ensemble methods<sup>3,16,17</sup>, and some analyses even revealed that domain information is the most crucial one<sup>16</sup>.

As in many other areas in bioinformatics, deep learning has been applied for function prediction from domain information. However, the effective use of domains is critically constrained by low coverage of functional assignments, high dimensionality, and acute data imbalance. For instance, in a recent competitive deep-learning-based model, DeepGOZero<sup>22</sup>, a 26,406-dimensional input of InterPro feature vectors was reduced to 1024 dimensions using a single multi-layer perceptron (MLP) layer, which results in significant information loss. A similar situation is observed in DeepGraphGO<sup>21</sup> as well.

Here, we introduce Domain-PFP, a protein function prediction method that uses functional representation of proteins through domain-GO association learned by a self-supervised method from protein databases. Self-supervised learning is based on the idea of leveraging the inherent co-occurrence relationship of complementary information in the data to learn new labels in a semi-automatic process<sup>28</sup>. We used self-supervised learning because it can directly learn domain and GO co-occurrence from abundant protein sequences and is able to alleviate the problem of current domain databases, where many domains do not have function annotation. Following the underlying concepts of self-supervised learning, we first learned pseudo-labels of GO prediction probability from individual domain terms. Then, we derived the dense representation of domains

consistent with functional information to characterize protein sequences and used the representation to predict protein functions. The embeddings learned both at the domain and protein level have turned out to be functionally meaningful as the embedding distance showed substantial negative correlations with functional similarity<sup>29</sup> of GO terms that are present in the domains and protein sequences. Moreover, a systematic comparison with large-scale Protein Language Model (PLM) representations<sup>30,31</sup>, which use variants of Transformers<sup>32</sup> and BERT<sup>33</sup> architectures, and have demonstrated success in function prediction<sup>38,60</sup>, revealed that our embeddings are more applicable for function prediction, despite being a fraction of the aforementioned PLM complexity. This improvement is further vividly observed in challenging cases of predicting rare and more specific functions. In addition, using a straightforward K-Nearest Neighbors (KNN) model with the learned embeddings along with sequence similarity and interaction information, Domain-PFP significantly outperforms more complex state-of-the-art methods. Most notably, Domain-PFP achieved an increase in the area under precision-recall curve (AUPR) by 2.43%, 14.58%, and 9.57% over the state-of-the-art method for molecular function (MF), biological process (BP), and cellular components (CC), respectively. Domain-PFP has also demonstrated competitive performance when compared with top-scoring methods in the CAFA3 evaluation<sup>34</sup>.

## Results

### Construction of Domain-GO Embeddings

#### Dataset of domains and GO annotations

We collected 568,002 protein sequences from Swiss-Prot (release 2022\_3)<sup>35</sup> and assigned InterPro domains using InterProScan 5 REST API<sup>36</sup>. Despite InterPro maximizing domain coverage by combining entries from 13 databases, 36,403 proteins had no InterPro annotations, so we discarded them. Concurrently, we collected GO terms for protein sequences from UniProt. We considered both experimentally and computationally assigned functions since IEA (Inferred from Electronic Annotation) terms demonstrated increased accuracy in our previous works<sup>6</sup>. We also propagated the parent GO terms using the core ontology release 2021-01-01. In summary, our dataset

contained 531,599 proteins with 32,471 unique domains and 33,199 unique GO terms (8,297, 21,805, and 3,097 MF, BP, and CC terms, respectively).

### Self-supervised learning for domain representation

Using the domain and GO term assignments to protein sequences, we computed the conditional probability of a protein that contains  $domain_i$  having the  $GO_j$  function:

$$p(GO_j | domain_i) = \frac{p(domain_i \cap GO_j)}{p(domain_i)} \quad (\text{Eq. 1})$$

Here,  $p(domain_i)$  represents the probability of a protein containing  $domain_i$ , while  $p(domain_i \cap GO_j)$  represents the joint probability of a protein with  $domain_i$  performing  $GO_j$ . We can calculate both probabilities from the co-occurrence relationships of domains and GO terms in the dataset by counting the occurrences. These probabilities serve as the pseudo labels or target function for our self-supervised learning method.

Our ultimate goal is to predict protein functions. To achieve this, we aim to develop a representation of domains that, in conjunction with a learned representation of GO terms, is consistent with the domain-GO co-occurrence conditional probability. In other words, we seek to design two representations or embeddings,  $\phi$  and  $\psi$ , which separately represent domains and GO terms, respectively, and a bivariate function  $f$  that can map the conditional probability of the co-occurrence of any  $domain_i$  and  $GO_j$ :

$$f(\phi(domain_i), \psi(GO_j)) \rightarrow p(GO_j | domain_i) \quad (\text{Eq. 2})$$


In our case, we utilized two 256-dimensional embedding matrices  $\phi$  and  $\psi$ , as representations for domains and GO terms, respectively. The bivariate function  $f$  was modeled as a two-layer densely connected network that takes the Hadamard product of  $\phi(domain_i)$  and  $\psi(GO_j)$  as input, decomposes the values in a 128-dimensional space, and finally predicts the conditional probability  $p(GO_j | domain_i)$ . The network architecture is presented in Figure 1(a),

where the function  $f$  is represented as an array of circles in light blue. Concretely,  $f$  takes the following form:

$$f = W_2(\text{RELU}(W_1(\phi(\text{domain}_i) \odot \psi(\text{GO}_j)) + b_1)) + b_2 \quad (\text{Eq. 3})$$

where  $(W_1, b_1)$  and  $(W_2, b_2)$  are weights and biases from the first and the second layer of the network, respectively. The Hadamard product of the two embeddings is represented by the symbol  $\odot$ . The network is regularized by dropout, and the domain embedding matrix  $\phi$  is further regularized by L1-norm to impose sparsity. The  $\phi$  embedding matrix for each domain, as well as the  $\psi$  embedding for each GO term, were learned through backpropagation with the mean squared error (MSE) loss using the Adam optimizer with default settings<sup>37</sup>. We intended to keep the function  $f$  simple so that the domain embeddings could effectively learn functional relevance, rather than letting the function  $f$  learn the correlation between domain and GO term co-occurrence. This is inspired by a recent work, which demonstrated that a strong encoder in conjunction with weak decoder results in a strong representation learner<sup>61</sup>. The function  $f$  provides the association probability between a domain and a GO term (Eq. 1), which we name DomainGO-prob. We trained three different versions of DomainGO-prob for the three sub-ontologies, MF, BP and CC, respectively.

The overall pipeline for learning the domain embeddings is summarized in Fig. 1b. We started by collecting annotated protein sequences from Swiss-Prot, along with domain and GO term assignments. Domains were obtained from InterProScan, while GO terms were collected from Swiss-Prot. Next, we calculated the conditional probabilities of all the domain-GO associations by counting their co-occurrences in the dataset. Finally, the domain embeddings ( $\phi$ ) and GO term embeddings ( $\psi$ ) were computed using the network shown in Fig. 1a. The network was trained and validated on the aforementioned dataset of 32,471 unique domains and 33,199 unique GO terms. We randomly selected 80% of the domain-GO pairs for training and used the remaining 20% for validation. Three different models, i.e., three different sets of embeddings were developed for the three sub-ontologies. The details of the network training are described in the Methods.



**Figure 1.** Overview of Domain-PFP. **a**, the network architecture used for self-supervised learning of domain embeddings. **b**, the overall pipeline of learning the functionally aware domain embeddings. **c**, the steps of computing the embeddings of a protein and inferring the functions.

### Predicting GO terms for a query protein (Domain-PFP)

Using the computed domain embeddings, we represented a protein, which may be composed of several domains, as the average of the embeddings of all the domains in it. This is similar to how PLM encodes proteins by averaging the individual residue level representations<sup>31</sup>. For a protein  $P_k$  with domains  $d_{P_k}$  the protein embedding is computed as

$$D(P_k) = \frac{\sum_{d \in d_{P_k}} \phi(d)}{|d_{P_k}|} \quad (\text{Eq. 4})$$

With the protein embedding, we can use supervised classifiers to infer protein functions. Here, we used a KNN classifier, following the convention of BLAST or PPI network scoring<sup>16,21</sup>. KNN models using protein language models have also been shown to be on par with top methods of Critical Assessment of Functional Annotation 3 (CAFA3)<sup>38</sup>. The confidence score of annotating a protein  $p_i$  with the GO term  $GO_j$ ,  $S_D(p_i, GO_j)$  is computed as follows:

$$S_D(p_i, GO_j) = \frac{\sum_{p_k \in K_{neigh}} I(p_k, GO_j) \times \|D(p_i) - D(p_k)\|^2}{\sum_{p_k \in K_{neigh}} \|D(p_i) - D(p_k)\|^2} \quad (\text{Eq. 5})$$

where  $K_{neigh}$  is a neighborhood of  $K$  proteins, and  $I(p_k, GO_j)$  is 1 if the protein  $p_k$  is annotated with  $GO_j$ , and 0 otherwise.

The steps of computing protein embeddings and predicting functions are outlined in Fig. 1c. For a given query protein sequence, domains are assigned using InterProScan, and their individual domain embeddings are obtained. The embedding of the query protein is then computed by taking the average of the assigned domain embeddings (Eq. 4). Finally, the protein embedding is used to find known proteins that are close in the embedding space (Eq. 5) using a supervised classifier (KNN for our approach) to infer its functions.


## Correlation of embedding distance and functional similarity of domains and proteins

To start with, we analyzed how the distance of the domain embeddings correlates with the functional similarity of domains and proteins. Having functionally similar domains close in the embedding space is essential for the embeddings to be useful for function prediction. As a measure of the embedding distance, we adopted the Manhattan distance as it is discussed to be more meaningful in high-dimensional spaces than, for example, the Euclidean distance<sup>39</sup>. As for functional similarity, we computed the Jaccard Index following a previous work<sup>3</sup>. For a domain, we considered GO terms assigned to the domain if they have a conditional probability no less than 0.5, i.e.,  $GO\ Terms = \{GO_i : p(GO_i | domain) \geq 0.5\}$ . This set of assigned GO terms for

domain A and B are denoted as  $GO\ Terms_{domainA}$  and  $GO\ Terms_{domainB}$  in the following equation. The Jaccard Index for two domains, A and B is defined as

$$Domain\ Functional\ Similarity\ (domain_A, domain_B) = \frac{|GO\ Terms_{domainA} \cap GO\ Terms_{domainB}|}{|GO\ Terms_{domainA} \cup GO\ Terms_{domainB}|} \quad (\text{Eq. 6})$$

We randomly selected 100,000 pairs of domains and computed their functional similarity relative to the embedding distance in Fig. 2a. Domain functional similarity was computed separately for each of the three GO categories. Overall, a negative correlation was observed between the embedding distance and functional similarity for all three GO categories. Substantial Jaccard Index values, such as those over 0.5, were observed mainly for domain embedding pairs that were close in distance, for example, less than 10. Almost all domain pairs with a large distance, for example, a distance of 20 or higher for MF and CC and over 10 for BP, had a small functional similarity value of less than 0.2. A perfect Jaccard Index of 1.0 was only observed for domain pairs with a relatively small embedding distance. Thus, it is evident that our model generates similar embeddings for functionally similar domains.



**Figure 2.** Domain and GO associations using DomainGO-prob. **a**, Functional consistency of domain embeddings. The domain functional similarity was quantified by the Jaccard Index of GO terms relative to the Manhattan distance of domain embeddings. Three GO categories, MF, BP, CC, are separately shown. **b**, Functional coherence in the protein level. Protein pairs were split into bins based on their embedding distance and the mean *funSim* score for each bin was plotted. Bins with less than 100 proteins were discarded. The last bin includes protein pairs with a distance larger than 25. The color of a data point indicates the number of protein pairs in the bin and the size of circles indicates the standard deviation of the *funSim* score. **c**, Predicted scores of GO terms for domains in InterPro2GO. The score distribution GO terms for domains were taken from DomainGO-prob. We included the scores from both the standard model (trained on the entire dataset) and the model trained in the adversarial manner (trained after removing the InterPro2GO pair information), which are represented with orange and blue bars, respectively.

We have also examined protein-level functional similarity relative to the embedding similarity (Fig. 2b). As the measure of the functional similarity of proteins, which are annotated by multiple GO terms in the three categories, we used the *funSim* score<sup>40</sup>. *funSim* essentially computes the average of semantic similarity of best matching GO terms from two proteins for each GO category, and then averages the score over the three GO categories (for the concrete definition, see Methods). *funSim* score ranges from 0 to 1 with 1 as the maximum score.

We took 1,000,000 random pairs of proteins and computed their embeddings for MF, BP, and CC separately and concatenated them to obtain the overall embedding. In Fig. 2b, mean *funSim* score of protein pairs were plotted relative to the Manhattan distance of protein embeddings. We can see the overall trend that *funSim* score drops as protein embeddings become more distant from each other. Large *funSim* scores were observed only for close protein embeddings, e.g. a Manhattan distance of less than 5.

Overall, in this section, we confirmed that functionally similar domains and proteins are placed close to each other in the embedding space.

## Learning InterPro2GO annotations

Next, we examined how well our domain embeddings align with expert-curated GO mappings of InterPro2GO. For this analysis, we used the InterPro2GO mapping of version-date 2022/03/16<sup>41</sup>, which comprises 35,046 mappings between 16,443 unique InterPro domain entries and 6,482

unique GO terms. We considered 34,832 InterPro2GO annotations, excluding 214 mappings with domains or GO terms that are not included in our dataset.

For all the domain-GO pairs in the InterPro2GO mappings, we predicted the conditional probability using DomainGO-prob (Fig. 1a) that the GO exists in the domain. The results are shown in Fig. 2c (using orange bars). As shown, for over 80% of cases, existing GO term-domain associations have a high score of over 0.9 (the rightmost bar) for all three GO categories. Thus, DomainGO-prob was able to associate GO terms to protein domains using the self-supervised learning protocol that associated GO terms and domains from the co-occurrences in full protein sequences.

We further conducted experiments with an adversarial version of this analysis. Namely, to test the generalization ability to learn from the context of related, co-occurring domains and GO terms alone, we removed all the probability values of domain-GO pairs that exist in the InterPro2GO mapping and then re-trained the DomainGO-prob models. Formally, from the original dataset  $\mathcal{D} = \{(domain_i, GO_j)\}$  we constructed a new dataset  $\mathcal{D}' = \{(domain_i, GO_j) : (domain_i, GO_j) \notin \text{InterPro2GO}\}$ . With this dataset  $\mathcal{D}'$ , we re-trained the DomainGO-prob models and examined the conditional probability of GO terms that exist in InterPro2GO. The results are represented by the blue bars in Fig. 2c. Under this setting, DomainGO-prob predicted a score higher than 0.5 for 66.5%, 81.9%, and 86.5% for MF, BP, and CC, respectively. Thus, even without explicit knowledge, DomainGO-prob was able to extract the meanings of domain-GO relationships only from the contextual information of co-occurrences and hierarchies. Among the three GO categories, the counts of domain-GO associations with the highest probability bin (0.9 to 1.0) of MF terms showed the largest decrease when compared with the results with full training data,  $\mathcal{D}$  (orange bar). This is probably because MF terms (e.g., enzymatic function) are associated with a domain at a residue level unlike BP and CC, which are more contextual<sup>42</sup>.

## Examples of domain-GO associations learned by the network

In this section, we discuss several examples that illustrate how DomainGO-prob learns domain-GO associations. We used the aforementioned adversarial version, i.e., the model trained with  $\mathcal{D}'$ , and we examined how the model likely learned the GO terms solely from the co-occurrence of different domains. The examples show that DomainGO-prob recovered the correct

domain and GO relationships in InterPro2GO from other domain and GO associations in a way that is consistent with the hierarchical and associative relationships of domains and GO terms. This is analogous to the way grammatical structures and word relations aid in masked language modeling<sup>33</sup> in NLP.

The first example (Fig. 3a) is IPR000010, a domain in cysteine protease inhibitors<sup>43</sup>, which was annotated with GO:0004869 (cysteine-type endopeptidase inhibitor activity) and GO:0004866 (endopeptidase inhibitor activity) both with a high probability of 0.9 by DomainGO-prob. The GO term GO:0004869 in the MF category represents binding to and preventing the activity of cysteine-type endopeptidase. Looking at the domain structure, IPR000010 has three subdomains, among which only one subdomain, IPR025764, a Fetuin-B-type cystatin domain, has an annotated GO term with experimental evidence, GO:0004866 (endopeptidase inhibitor activity)<sup>44</sup>. From this domain hierarchy, in addition to GO:0004866, a child term, GO:0004869 with an ‘is a’ relationship with GO: 0004866, was correctly transferred to IPR000010 by DomainGO-prob.


The second example is a recovered annotation of a CC term, GO:0005634, which represents nuclear localization, with a probability of 0.99 assigned to IPR000690. In InterPro2GO, GO:0005634 is the only CC term associated with this domain. IPR000690 is Matrin/U1-C, C2H2-type zinc finger, which co-occurs with the homologous superfamily IPR036236 in 86.6% of protein sequences (Fig. 3b). IPR036236 is zinc finger C2H2-type superfamily, and 98.6% of its proteins are annotated with the CC term of nuclear localization. For instance, the protein A5PJN8 has both the two InterPro entries and is also annotated with the GO:0005634 term. Therefore, DomainGO-prob extracted the CC term from the co-occurrence of these domains in proteins and correctly annotated IPR000690.

The next example in Fig. 3c illustrates the transfer of a GO term from multiple co-occurring domains in proteins. DomainGO-prob annotated IPR000081 with the function GO:0016032 (viral process) with a probability of 1.0, which refers to a multi-organism process by a virus. All proteins with this domain (for example, P03303) also have domains IPR007094 (encoded in RNA-containing viruses), IPR001205 (found in RNA viruses), IPR000605 (found in DNA viruses), IPR029053 (forms icosahedral virus shell), IPR002527 (alters membrane permeability), IPR014838 (poliovirus replication) or 8 other domains related to various viral activities. Although not all such co-occurring domains have exactly GO:0016032, they all have related terms, such as

GO:0039694 (viral RNA genome replication). Therefore, DomainGO-prob was able to learn the viral process function GO:0016032 by combining such supplementary information.

Some domains are responsible for multiple different functions. For example, the domain IPR000081, which has just been analyzed for viral activity in the previous example, was also correctly assigned with proteolysis (GO:0006508) by DomainGO-prob with a predicted probability of 1.0 (Fig. 3d). However, this information was not learned from the aforementioned co-occurring domains, but rather from the homologous superfamily IPR009003 (Peptidase S1, PA clan), which all proteins with IPR000081 is a part of. For example, the protein P06209 not only contains the domain IPR000081 but also is a member of IPR009003 homologous superfamily. Cysteine peptidase from IPR009003 hydrolyses a peptide bond using the thiol group<sup>45</sup> and thus has the GO:0006508 function, which was derived to IPR000081. It should be noted that despite IPR009003 completely overlapping with IPR000081, DomainGO-prob did not associate IPR009003 with viral activity (GO: 0016032). For the domain IPR009003 DomainGO-prob predicted a small probability of 0.33 for GO:0016032 (viral process), which was likely induced from the several co-occurring domains involved in viral activities, for example, IPR007094, IPR002527, IPR014838. On the contrary, the actual function for IPR000081, i.e., GO:0006508 was predicted with a probability of 1.0. Therefore, for this example DomainGO-prob was capable of contrasting between complementary information.

There are cases where DomainGO-prob failed to associate GO terms to domains. For instance, in Fig. 3e, IPR000174 represents two different Chemokine receptors from the CXC family, namely CXCR1 and CXCR2<sup>46</sup>. Therefore, proteins from this family are annotated with the function GO:0016494 (CXC chemokine receptor activity). Since Chemokine receptors are part of the G protein-coupled receptor (GPCR) family, proteins from the IPR000174 family (for example, P21109) are also members of IPR000276 (G protein-coupled receptor, rhodopsin-like) and have the IPR017452 (GPCR, rhodopsin-like, 7TM) domain. Although this context provides information about the GPCR family, it is difficult to narrow it down to the CXC family, without individual precise information. This precise information is absent in our adversarial mode of training. As a result, DomainGO-prob predicted a low score of 0.34 for GO:0016494 but managed to assign GO:0004930 (G protein-coupled receptor activity) to IPR000174 with 0.95 probability from the co-occurrence of IPR017452.



**Figure 3.** Examples to explain how domain-GO associations learned by DomainGO-prob can be interpreted. The target domain in discussion in each example are shown in bold. Yellow arrows indicate the GO term (blue) that was transferred to the target domain. The associated number with the arrow is the predicted probability that the domain has the GO term. Other GO terms in discussion are colored purple. **a**, an example where a GO term was obtained from subdomain. GO:0004866 is a direct parental term of GO:0004869 with “is a” relationship. **b**, an example of learning GO terms from co-occurred domains. **c**, examples that a GO term was obtained from multiple co-occurring domains that have a common context. **d**, an example where DomainGO-prob was able to distinguish correct and incorrect GO terms that exist in co-occurring domains. **e**, an example where the contextual information is not sufficient to retrieve a more specific GO term.

## Comparison with large protein language models in GO function prediction

We evaluated the performance of DomainGO-prob embedding in comparison with 12 large Protein Language Models (PLMs) following the benchmark study performed by Unsal et al.<sup>47</sup>. The 12 PLMs we compare against are ProtT5-XL<sup>30</sup>, ProtALBERT<sup>30</sup>, SeqVec<sup>48</sup>, ProtBERT-BFD<sup>30</sup>, ESM-1b<sup>31</sup>, ProtXLNet<sup>30</sup>, TAPE-BERT-PFAM<sup>49</sup>, CPCPProt<sup>50</sup>, MSA-Transformer<sup>51</sup>, UniRep<sup>52</sup>, Learned-Vec<sup>53</sup>, and ProtVec<sup>54</sup>. These PLMs were trained on unsupervised tasks such as predicting a segment of masked residues given the rest of the protein<sup>30</sup> or predicting the next residue from all the residues before it<sup>48</sup>, on a large protein sequence dataset, e.g. the entire UniProt. Supplementary Table 1 summarizes how these PLMs were trained.

To use a PLM for GO prediction, Unsal et al. converted the residue-level embedding to protein-level by computing the mean of the embeddings along the residues and used a linear Support Vector Machine model. The benchmark by Unsal et al. was performed on the PROBE benchmark dataset they constructed, which provides GO terms of different difficulties to predict. In the Probe dataset, GO terms are divided into three categories based on the frequency in the PROBE benchmark dataset (low, middle, high having 2-30, 100-500, >1000 annotated proteins, respectively) and specificity (shallow, normal, specific for the ontologies being within the depth of 1/3<sup>rd</sup>, 2/3<sup>rd</sup> and bottom rest, respectively). Therefore, based on frequency and specificity,  $3 \times 3 = 9$  groups of GO terms can be constructed for the three GO categories, i.e.  $3 \times 9 = 27$  groups. Among them, as there were no GO terms that fall under the high-specific group, the benchmark ended up with 25 groups. For each group, at most 5 GO terms were selected based on dissimilarity according to the Lin's similarity measure<sup>40</sup>, which resulted in a total of 117 GO terms to predict. The PROBE dataset contains 19,995 human proteins clustered at a 50% identity cutoff and only experimental GO annotation. The human proteins falling under these criteria were used for benchmarking GO function prediction by undergoing a 5-fold cross-validation test. Unsal et al. provided a convenient CodeOcean distribution (<https://PROBE.kansil.org>, version November 3, 2022), where given the embeddings of the test proteins, GO predictions are made, and the performance is evaluated on the PROBE dataset. We used it to test our DomainGO-pair-based protein embedding (Eq. 4).

For this benchmark, we trained the domain and GO embeddings (Eq. 2) for the three GO categories separately on Swiss-Prot, after removing all the human proteins. We removed these

proteins to avoid overlap between the test proteins and the proteins used for training. However, note that the PLMs we compared against almost certainly have these human proteins in their training set, as they used an entire public protein sequence dataset for training. Since our embedding dimension is only 256, which is quite small compared to the PLMs we compared against, we concatenated the embeddings from the three GO categories and performed mean normalization to balance them. This resulted in a 768-dimensional protein embedding vector, as follows:

$$\text{Embedding}(p) = \text{MeanNorm}(\text{Concat}(D_{MF}(p), D_{BP}(p), D_{CC}(p))) \quad (\text{Eq. 7})$$

Here,  $D_{MF}(p)$ ,  $D_{BP}(p)$ ,  $D_{CC}(p)$  are computed embeddings for a protein  $p$  for MF, BP, CC sub-ontologies, respectively (using Eq. 4.)


The results are presented in Fig. 4(a), where we compared the GO prediction performance of our model (Eq. 7) with 12 models on the PROBE benchmark. The numerical values are provided in Supplementary Table 2. Our model based on DomainGO-prob outperformed all the PLMs in all three categories. For MF, BP, and CC, DomainGO-prob resulted in 0.02, 0.06, and 0.06 higher weighted F1 scores than ProtT5-XL, the previous top method, respectively, and 0.04 when the average across the three GO categories was considered. Notably, this improvement was obtained from a much simpler model with 768-dimensional embeddings with merely a fraction of the parameters of the PLM models by adopting a functionally informed learning protocol. As shown in Supplementary Table 1, ProtT5-XL has 1024-dimensional embeddings and was trained with a network with 3 billion parameters, while the three networks we used (Eq. 7) have only 31 million parameters in total.

An important consideration when training machine learning models for protein analysis is to remove redundancy, i.e. similar sequences from the training set relative to the test set. Therefore, although we have already omitted human protein sequences from our training dataset, we retrained our models after removing proteins with 75%, 50%, and 25% sequence identity with the test set using MMseqs2<sup>55</sup>. The results are shown in Fig. 4b, in comparison with ProtT5-XL. As expected, the F1 score decreased slightly as more sequences were removed. However, this is most likely due to the fact that we were losing some domain and GO association information when we removed


the proteins. Nevertheless, in all cases, our embedding performed better than that of ProtT5-XL, even at the identity cutoff of 25%.

In Fig. 4c, we examined how performance changed when considering GO terms of varying levels of difficulty to predict. Weighted F1 scores for the different GO groups in the PROBE benchmark, classified by GO depth and frequency, were separately shown. As we moved from high to low frequency or shallow to specific GO terms, the classification task became more difficult. We compared our model's performance with the best-performing PLM, ProtT5-XL. Even in this evaluation, it was evident that our model substantially outperformed ProtT5-XL. Interestingly, the margin of the advantage of our model increased as we considered more difficult GO groups. In most of the easier cases, DomainGO-prob was at least similar to or slightly better than ProtT5-XL. In difficult cases, a substantial improvement was observed. For example, for low-frequency and specific CC terms, DomainGO-prob was 20% better. It is apparent that although the PLM was able to comprehend frequent GO terms from unsupervised learning on a large volume of protein sequence data, such models failed to account for rare GO terms and suffered from limited specificity. On the contrary, our self-supervised learning approach seemed to decipher the functional identity of proteins better, regardless of the rarity and specificity of the GO terms to a degree.

**a**



**b**



**c**

MF			BP			CC				
	Shallow	Normal	Specific	Shallow	Normal	Specific	Shallow	Normal	Specific	
High	DG:0.87 T5:0.89	DG:0.94 T5:0.94	N/A	High	DG:0.61 T5:0.63	DG:0.73 T5:0.73	<u>DG:0.68</u> T5:0.66	High	DG:0.74 T5:0.75	DG:0.71 T5:0.72
	DG:0.90 T5:0.91	<u>DG:0.94</u> T5:0.91	<b>DG:0.95</b> T5:0.90		<u>DG:0.74</u> T5:0.72	<b>DG:0.89</b> T5:0.83	<u>DG:0.81</u> T5:0.78		DG:0.81 T5:0.83	<u>DG:0.80</u> T5:0.77
	<b>DG:0.89</b> T5:0.79	DG:0.89 T5:0.94	<b>DG:0.98</b> T5:0.93		<u>DG:0.64</u> T5:0.60	<b>DG:0.47</b> T5:0.24	<b>DG:0.90</b> T5:0.71		<b>DG:0.86</b> T5:0.67	<b>DG:0.65</b> T5:0.57
Middle				Middle				Middle		
Low				Low				Low		

**Figure 4.** Comparison with 12 large PLMs in the PROBE benchmark. **a**, the weighted F1 Score obtained by our model that is based on DomainGO-prob with 12 PLMs on the PROBE benchmark. To compute the weighted F1 score, the values of F1 score for the individual GO terms were averaged, weighted by the number of samples having that particular GO term, for all three GO categories and the average across the categories. **b**, analysis of the performance of DomainGO-prob when trained on non-redundant datasets. The model with DomainGO-prob was retrained on training datasets after removing proteins with 75%, 50%, and 25% sequence identity to the test proteins. None indicates the result using the original training set without further removing training proteins. The dashed line shows the weighted F1 score of ProtT5-XL. **c**, comparative performance on the nine groups of GO terms with different difficulty levels against the best-performing ProtT5-XL. DG, prediction with the model (Eq. 7) using DomainGO-prob. N/A indicates that

there is no GO term in that particular group. DG is shown in bold and underlined when the improvement over ProtT5-XL is greater than and less than 5%, respectively.

## GO function prediction by Domain-PFP in comparison with existing methods

Subsequently, we benchmarked the GO prediction performance of Domain-PFP (Eq. 5) on the NetGO dataset<sup>16</sup> to compare it with state-of-the-art protein function prediction methods from recent literature. We used the data split of the NetGO dataset into training, validation, and test sets provided in the work of DeepGOZero<sup>22</sup>, who followed the same data split protocol of NetGO2.0<sup>16</sup>. The NetGO dataset consists of 64,279, 91,443, and 83,004 proteins for MF, BP, and CC categories, respectively, with a specified training, validation, and test set split (Supplementary Table 3). We trained DomainGO-prob on the NetGO2.0 training dataset and created a weighted K Nearest Neighbor (KNN) model based on the learned embedding. The number of K-neighbors used for MF, BP, and CC was 1000, 800, and 1200, respectively, which were tuned based on the performance on the validation data split of the NetGO benchmark (Supplementary Figure 1). The performance of the various models on the test dataset is presented in Table 1. The evaluation results of the existing methods, from BLAST to NetGO2.0 (Server) in Table 1, were taken from the paper of DeepGOZero<sup>22</sup>.

DeepGOPlus<sup>18</sup> infers protein functions through a combination of DiamondBLAST<sup>56</sup> and DeepGOCNN, which employs a 1D convolutional neural network to predict GO from the amino acid sequence. TALE+<sup>19</sup> similarly fuses DiamondBLAST with sequence representation learned from a Transformer. Other top-performing methods are either based on domain information or used as a component. For instance, DeepGOZero<sup>22</sup> leverages a model-theoretic approach to predict ontologies from InterPro domains, which can be further improved by incorporating DiamondBLAST. DeepGraphGO<sup>21</sup> associates InterPro features with protein-protein interaction (PPI) networks employing a graph convolutional neural network. NetGO2.0<sup>16</sup> is an all-encompassing ensemble method that incorporates BLAST, domain, PPI, GO term frequency, PubMed publications, and sequence information both in form of k-mers and embedding. Among the existing methods, NetGO2.0 has shown the highest evaluation values for MF and the best  $S_{min}$  value in BP<sup>34</sup> (note that the NetGO2.0 results are from the current server, ran by the authors of NetGOZero in their paper).

In the latter half of Table 1 we show results by Domain-PFP and Domain-PFP that incorporate BLAST and PPI information to compare with the other state-of-the-art methods that combine diverse information sources. The scores of GO term  $j$  for protein  $i$  from BLAST and PPI information are defined as

$$S_B(p_i, GO_j) = \frac{\sum_{p_k \in D} I(p_k, GO_j) \times B(p_i, p_k)}{\sum_{p_k \in D} B(p_i, p_k)} \quad (\text{Eq. 8})$$

$$S_N(p_i, GO_j) = \frac{\sum_{p_k \in D} I(p_k, GO_j) \times \omega(p_i, p_k)}{\sum_{p_k \in D} \omega(p_i, p_k)} \quad (\text{Eq. 9})$$

Here,  $B(p_i, p_k)$  and  $\omega(p_i, p_k)$  are the bit-score from DiamondBLAST with ‘more-sensitive’ setting<sup>56</sup>, and edge weight from STRING PPI network (ver. 11.0)<sup>57</sup>, respectively. We used the same STRING version as DeepGraphGO<sup>21</sup>.

The final score is a simple average of the terms from the three sources:

$$S(p_i, GO_j) = \frac{S_D(p_i, GO_j) + I_B(p_i) S_B(p_i, GO_j) + I_N(p_i) S_N(p_i, GO_j)}{1 + I_B(p_i) + I_N(p_i)} \quad (\text{Eq. 10})$$

$I_B$  and  $I_N$  are identity functions, which results in 1 if BLAST and String Network matches are found for the protein  $i$ , respectively.

We compared the performance of the methods using the three CAFA evaluation metrics, namely  $F_{\max}$ , AUPR, and  $S_{\min}$ <sup>34</sup> (see Methods).  $F_{\max}$  computes the maximum possible protein-centric F1 score, overall prediction thresholds. AUPR, the area under PR curve, on the other hand, is a suitable metric for imbalanced data and penalizes the false positive predictions, which is highly applicable to function prediction<sup>16</sup>. Finally,  $S_{\min}$  is a measure of semantic distance between predicted and actual annotation, based on the information content of the individual GO terms<sup>18</sup>, i.e., this metric signifies the capability of predicting rare GO terms.

**Table 1.** Comparison with state-of-the-art function prediction methods.

Method	Features Used	F <sub>max</sub>			AUPR			S <sub>min</sub>		
		MF	BP	CC	MF	BP	CC	MF	BP	CC
BLAST	BLAST	0.627	0.407	0.625	0.427	0.272	0.412	5.503	25.918	9.351
DeepGOCNN	Sequence	0.589	0.337	0.624	0.565	0.271	0.623	6.417	27.235	10.617
DeepGOPlus	Sequence, BLAST	0.661	0.419	0.655	0.667	0.342	0.663	5.407	25.603	9.374
DeepGOZero	Domain	0.662	0.396	0.662	0.668	0.337	0.645	5.322	25.838	9.834
DeepGOZero + BLAST	Domain, BLAST	0.655	0.432	0.675	0.665	0.356	0.654	5.337	25.439	9.391
DeepGraphGO	Domain, PPI	0.671	0.418	<u>0.679</u>	0.647	0.364	0.669	5.374	25.866	<u>9.165</u>
TALE+	Sequence, BLAST	0.466	0.382	0.661	0.441	0.31	0.681	8.136	26.308	9.599
NetGO2.0 (Server)	Domain, Kmer, RNN, PPI, BLAST, Pubmed, Frequency	<b>0.698</b>	0.431	0.662	<u>0.701</u>	0.343	0.627	<u>5.187</u>	<u>25.076</u>	9.473
Domain-PFP	Domain	<u>0.675</u>	0.41	0.675	0.676	0.344	<u>0.697</u>	5.259	25.838	9.709
DPFP + BLAST	Domain, BLAST	0.674	<u>0.434</u>	<u>0.681</u>	<u>0.693</u>	<u>0.367</u>	<b>0.717</b>	<u>5.188</u>	25.11	<u>9.239</u>
DPFP + PPI	Domain, PPI	0.666	<u>0.435</u>	0.673	0.689	<u>0.379</u>	0.675	5.404	<u>25.002</u>	9.35
DPFP + BLAST + PPI	Domain, BLAST, PPI	<u>0.685</u>	<b>0.452</b>	<b>0.686</b>	<b>0.718</b>	<b>0.393</b>	<u>0.687</u>	<b>5.146</b>	<b>24.292</b>	<b>9.084</b>

The best, 2<sup>nd</sup> best, and 3<sup>rd</sup> best results are indicated by bold, double-underline, and single-underline, respectively. We also include the features used by the predictors. DPFP, Domain-PFP; Kmer, the K-mer distribution in the protein sequences; RNN, protein sequence embedding computed by a recurrent neural network; Frequency, the frequency of the GO terms in the database.

Firstly, we compared Domain-PFP with sequence-only or domain-based methods, e.g., DeepGOCNN and DeepGOZero. This is a fair comparison as the base Domain-PFP uses only domain information which is inferred from sequence information. It can be observed from the table that Domain-PFP outperforms these methods in terms of F<sub>max</sub>, AUPR, and S<sub>min</sub> in all the three sub-ontologies. Notably, Domain-PFP achieved an AUPR of 0.697 for CC, whereas DeepGOZero, a recent method based on domain information scored 0.645, i.e., a large improvement of 0.052. In

terms of Fmax, Domain-PFP outperformed DeepGOZero and DeepGOCNN by achieving 0.013-0.014 and 0.051-0.086 higher scores, respectively.

Adding different features generally improves the performance of function prediction. When BLAST information was combined, Domain-PFP improved the overall performance, except for a slight drop of 0.001 in Fmax for MF. Fmax for BP increased from 0.41 to 0.434, and AUPR for CC increased from 0.697 to 0.717. Furthermore, Domain-PFP with BLAST consistently outperformed DeepGOZero+BLAST, which also uses the same information, in all 9 metrics. For example, DeepGOZero+BLAST achieved AUPR scores of 0.665, 0.356, and 0.654 for MF, BP, and CC, respectively, whereas Domain-PFP+BLAST achieved 0.693, 0.367, and 0.717, representing improvements of 0.028, 0.011, and 0.063, respectively. When compared with DeepGOPlus or TALE+, both of which use BLAST, the improvements made by Domain-PFP+BLAST appeared consistent as well.

Next, we experimented with including PPI information with Domain-PFP. However, this only improved the performance in BP, as expected, since BP involves multiple related and interacting functions that can be captured by PPIs. On the other hand, the performance of MF and CC was negatively affected. This situation is similar to the findings of NetGO2.0<sup>16</sup>, where the authors reported that PPI information performed better than domain information for predicting BP terms but not for MF and CC terms. For example, the Fmax of MF and CC dropped by 0.009 and 0.002, respectively. Despite this, Domain-PFP+PPI still outperformed DeepGraphGO, a method using domain and PPI information in a much more sophisticated graph neural network, in 5 out of 9 metrics.

Finally, we experimented with integrating both BLAST and PPI simultaneously. This brought improvements in all the metrics except for AUPR of CC. Notably, Fmax and AUPR of BP improved by 0.042 and 0.049, respectively. This integration of BLAST and PPI features enabled Domain-PFP to consistently perform superior to all the existing methods. For example, the current state-of-the-art method NetGO2.0 was surpassed by Domain-PFP in 8 out of 9 metrics (except for Fmax for MF). In terms of Fmax for BP and CC, Domain-PFP achieved 0.021 and 0.024 higher scores, respectively. For AUPR, the improvements were 0.017, 0.050, and 0.060 for MF, BP, and CC, respectively. Similarly, in terms of Smin, Domain-PFP+BLAST+PPI achieved 0.041, 0.784, and 0.389 smaller scores for MF, BP, and CC, respectively, implying that non-trivial GO terms were captured better.

This comparative evaluation with state-of-the-art function prediction methods further supports our self-supervised approach of learning functionally informed representations for protein domains. We observed that a simple KNN model with DomainGO-prob embedding not only outperformed more sophisticated deep learning models (e.g., DeepGraphGO) but also methods with access to more information sources (e.g., NetGO2.0). The only case where we fell behind the previous state-of-the-art, NetGO2.0, is in Fmax for MF, which we hypothesize is due to the inclusion of Pubmed publication information that is likely to contain precise information vital for MF prediction.

In order to assess the performance of Domain-PFP relative to structure-based protein function predictors we considered two recent methods, DeepFRI<sup>8</sup> and GAT-GO<sup>62</sup>. These methods use 3D protein structure information in a graph neural network and protein sequence information with a language model. Both methods were evaluated on a common benchmark dataset, composed of 29,902, 3,323, and 3,416 proteins for training, validation, and testing, respectively. The train and test proteins possess a total of 2,752 GO terms.


We retrained Domain-PFP on this training dataset and observed the performance across the three sub-ontologies in the test set. The results are presented in Supplementary Table 5. It shows that Domain-PFP outperforms the two much complex graph neural network-based function predictors with access to structural information on all the metrics except for Fmax in CC and AUPR in MF. The performance of Domain-PFP was further improved by including BLAST predictions, which results in the best score for all the metrics.

## Evaluation on CAFA3 benchmark

We further evaluated Domain-PFP on the CAFA3 benchmark<sup>34</sup>. We trained the network model of Domain-PFP using the CAFA3 training dataset and evaluated the results using the official evaluation code. The training dataset comprised 66,841 protein sequences annotated before September 2016, with 677, 3992, and 551 MF, BP, and CC GO terms, respectively (Supplementary Table 4). The test set contained 3328 proteins annotated between September 2016 to February 2017. To include sequence similarity information using BLAST in our pipeline, we constructed a new BLAST database with the CAFA3 training sequences. However, we could not use PPI information from the STRING database for this benchmark because STRING v10.a (the version during the competition timeline) lacked sufficient interaction data of the CAFA3 test proteins. We

did not perform any additional hyperparameter tuning and kept the same hyperparameters computed from the NetGO benchmark validation data.

The results of Domain-PFP on the CAFA3 benchmark are presented in Fig. 5 in comparison with the top 10 performing methods as published by the organizers of CAFA3<sup>34</sup>. DomainPFP+BLAST consistently showed a higher Fmax than Domain-PFP alone. For both BP and CC, Domain-PFP+BLAST outperformed the existing methods. Domain-PFP+BLAST achieved a Fmax score of 0.63 for CC, which is 0.02 higher than the CAFA3 top model, Zhu Lab. For BP, Domain-PFP+BLAST showed a slightly higher Fmax of 0.398 than the CAFA3 top model (Fmax: 0.397). For MF, our Fmax score, 0.59, was second to Zhu Lab (Fmax: 0.62), with a substantial margin to the next method, orengo-funfams (Fmax: 0.54). The top method by Zhu Lab combined more diverse information using an ensemble approach, including sequence, domain, homology, biophysical information, which likely gave that a competitive edge, similar to NetoGO2.0. We also mention that both DomainPFP and DomainPFP+BLAST showed higher Fmax scores than DeepGOPlus, which reported Fmax scores of 0.557, 0.390, and 0.614 for MF, BP, and CC, respectively, in their paper<sup>18</sup>.



**Figure 5.** Comparison with CAFA3 methods. The F<sub>max</sub> for the top 10 methods and the 2 baseline methods (Naïve, BLAST) are presented. All the scores were collected from the official CAFA3 result (Zhou et al., 2019). The F<sub>max</sub> score of Domain-PFP was also computed using the official CAFA3 evaluation codes.

## Discussion

Despite protein domains carrying the functional signatures of proteins, they have not been used to their full potential to date. Look-up-table-based domain to GO assignments tend to lack coverage. On the other hand, deep learning-based approaches using domains as high-dimensional input suffer from limitations in training data and information loss in network bottlenecks. Therefore, our motivation has been twofold: improving coverage and reducing information loss. Based on recent advancements in self-supervised learning, it has become motivating to apply such concepts in protein domain learning to alleviate these issues. Our method follows one of the core concepts of self-supervised learning where pseudo-labels are first learned to initialize model parameters, which are then used to perform the actual task using a supervised or unsupervised method<sup>63</sup>. Our approach is consistent with this definition as we first use the domain-go association probabilities as pseudo labels, which initializes our domain embedding parameters; then, we use this embedding later in a supervised learning protocol and we predict the functions of the proteins. This strategy also holds in the benchmarks on the NetGO and CAFA3 dataset we performed. To the best of our knowledge, this work is the first to apply self-supervised learning in the domain of protein function prediction. Based on co-occurrence contextual information between domain and GO terms, we devise embeddings for domains so that functionally related domains have similar embeddings. Since co-occurrences were learned from entire protein sequences, the domain embedding, DomainGO-prob, encodes GO associations that are not explicitly described in the domain database. Remarkably, our rather simple model, Domain-PFP, along with BLAST and PPI information, demonstrated superior performance over all state-of-the-art function predictors.

One likely limitation of this work could be the case of unknown domains. All existing methods based on protein domains fail to predict anything if the domain seen during inference was absent in the training data, in which case they predict a default value. This limitation can possibly be resolved by generating the functionally aware domain representation and localization end-to-end from the protein sequence directly using a larger deep learning model. Another limitation is that the current protein embedding considers domains in a protein equally (Eq. 4), although each domain may have different levels of contribution to protein function. Also, the order of appearance of domains in a protein is not considered, which is known to be relevant to function<sup>13</sup>. To address these, attention mechanism maybe applicable. These are improvements we wish to explore in our

subsequent works. In the current work, we practically alleviated these issues by augmenting with BLAST and PPI information-based predictions.

Another possible future direction could be to combine with general protein language models<sup>31</sup>, which were shown to perform well in protein tertiary structure prediction and other tasks. Additionally, we wish to analyze the suitability of our model in a zero-shot learning scenario. Specifically, our goal is to train DomainGO-prob on pretrained GO embeddings based on GO tree hierarchy and observe if GO terms absent in the training data can be retrieved this way.

## Methods

### Neural network architecture

We have designed a neural network to learn the domain-GO co-occurrence conditional probability distribution (Fig. 1a). The domain and GO terms are received as one-hot-encoded inputs, which are passed through two separate embedding layers to generate the 256-dimensional domain and GO embeddings, respectively. Then, from the computed domain and GO embeddings, we calculate the Hadamard product as a measure of correlation between the two types of embeddings and pass them through a densely connected layer of 128 neurons. The neurons are regularized through dropout ( $p=0.05$ ) and activated by RELU. Finally, we use a linear layer to predict the  $p(\text{domain}|\text{GO})$  score. The domain embedding matrix is extracted to generate the representations of domains. In order to increase the sparsity of the domain embeddings, we apply L1-regularization on that embedding layer ( $\lambda = 0.1$ ).

### Network training

Similar to word2vec embedding training<sup>58</sup>, we have a comparatively much larger number of domain-GO terms out of context, i.e.,  $p(\text{GO}|\text{domain}) = 0$ . Thus, we employed negative sampling by randomly selecting 1000-2000 non co-occurring GO terms for each domain. The network was trained by minimizing the MSE (mean squared error) loss with Adam optimizer<sup>37</sup> with a learning rate of 0.001 (the other parameters were kept as default) and a batch size of 163,840 for 200 epochs. 20% of domain-GO pairs, which were randomly selected, were used as the validation set. The

experiments were performed 10 times and the best model based on validation performance was selected.

### funSim score

*funSim* score is popularly used for quantifying similarity of GO term annotation of two proteins<sup>29,59</sup>. *funSim* score uses the relevance semantic similarity score  $sim_{REL}$  for the similarity of GO terms of the same category<sup>40</sup>:

$$sim_{Rel}(GO_1, GO_2) = \max_{go \in Ancestors(GO_1, GO_2)} \left( \frac{2 \log p(go)}{\log p(GO_1) + \log p(GO_2)} \times (1 - p(go)) \right) \text{(Eq. 11)}$$

where common ancestral GO terms of  $GO_1$  and  $GO_2$  are explored to maximize the score and  $p(GO)$  is the probability of GO term in the entire Swiss-Prot database. Then, a set of GO annotations in a GO category for two proteins,  $a$  and  $b$ , are defined as

$$GO_{score}(protein_a, protein_b) = \max \left( \frac{1}{N} \sum_{j=1}^N \max_{1 \leq j \leq M} s_{ij}, \frac{1}{M} \sum_{i=1}^M \max_{1 \leq i \leq N} s_{ij} \right) \text{(Eq. 12)}$$

where  $s_{ij}$  is  $sim_{REL}$  score of  $GO_i$  and  $GO_j$  of  $protein_a$  and  $protein_b$ , respectively, computed in an all-vs-all fashion.

Finally, *funSim* score is the average of the  $GO_{score}$  from the three GO categories.

$$funSim = \frac{1}{3} \left[ \left( \frac{MF_{score}}{\max(MF_{score})} \right)^2 + \left( \frac{BP_{score}}{\max(BP_{score})} \right)^2 + \left( \frac{CC_{score}}{\max(CC_{score})} \right)^2 \right] \text{(Eq. 13)}$$

*funSim* score ranges from 0 to 1 with 1 as the maximum score.

### Evaluation Metrics

For the PROBE benchmark, similar to the original benchmark by Unsal et al.<sup>47</sup>, we used Weighted F1 Score as the evaluation metric. The values were computed using their official CodeOcean distribution.

To compare with state-of-the-art methods, we used the CAFA protein-centric evaluation metrics Fmax, Smin, and AUPR<sup>2</sup>. We used the same evaluation codes as used by<sup>22</sup> to ensure consistency.

$F_{\max}$  is the maximum possible protein-centric F1 score, computed over all prediction thresholds.

$$F_{\max} = \max_{0 \leq \tau \leq 1} \frac{2 \ pr(\tau) \ re(\tau)}{pr(\tau) + re(\tau)} \quad (\text{Eq. 14})$$

Here,  $pr(\tau)$  and  $re(\tau)$  are precision and recall scores, respectively, computed at the cut-off value of  $\tau$ . The precision and recall values are computed as

$$pr(\tau) = \frac{1}{h(\tau)} \sum_{j=1}^{h(\tau)} \frac{\sum_i I(S(G_i, P_j) \geq \tau) \cdot I(G_i, P_j)}{\sum_i I(S(G_i, P_j) \geq \tau)} \quad (\text{Eq. 15})$$

$$rc(\tau) = \frac{1}{N_T} \sum_{j=1}^{N_T} \frac{\sum_i I(S(G_i, P_j) \geq \tau) \cdot I(G_i, P_j)}{\sum_i I(G_i, P_j)} \quad (\text{Eq. 16})$$

Here,  $N_T$  is the total number of proteins and  $h(\tau)$  is the number of proteins with a prediction score no smaller than  $\tau$  for at least one GO term.  $I$  is the identity function which returns 1 if the condition is true, 0 otherwise.  $I(G_i, P_j)$  therefore, implies if the protein  $P_j$  whether has the GO term  $G_i$  or not.  $S(G_i, P_j)$  denotes the prediction score of  $P_j$  having the  $G_i$  term.

The area under precision-recall curve, i.e., AUPR score is computed from the computed precision and recall scores using the trapezoidal rule.

$$AUPR = \frac{\Delta x}{2} (f(x_0) + 2f(x_1) + 2f(x_2) + \dots + 2f(x_{N-1}) + f(x_N)) \quad (\text{Eq. 17})$$

Here,  $x_0, x_1, \dots, x_N$  are various recall values, whereas  $f(x_0), f(x_1), \dots, f(x_N)$  are values of precision at those recalls, and  $\Delta x$  is the step size.

$S_{min}$  is a measure of semantic distance between the ground truth and prediction annotations based on information content of the GO classes. The information content  $IC(c)$  for a class  $c$  is computed based on the annotation probability of class  $c$  relative to its parent class  $P(c)$

$$IC(c) = -\log (probability(c|P(c))) \quad (\text{Eq. 18})$$

The two terms remaining uncertainty ( $ru$ ) and average misinformation ( $mi$ ) are defined as

$$ru(t) = \frac{1}{n} \sum_{i=1}^n \sum_{c \in \{T_i - P_i(t)\}} IC(c) \quad (\text{Eq. 19})$$

$$mi(t) = \frac{1}{n} \sum_{i=1}^n \sum_{c \in \{P_i(t) - T_i\}} IC(c) \quad (\text{Eq. 20})$$

The value of  $S_{min}$  is computed as

$$S_{min} = \min_t (\sqrt{ru(t)^2 + mi(t)^2}) \quad (\text{Eq. 21})$$

## Code availability

The Domain-PFP program is freely available for academic use from GitHub at (<https://github.com/kiharalab/Domain-PFP>). Furthermore, the program is available to run on Google Colab Notebook ([bit.ly/domain-pfp-colab](https://bit.ly/domain-pfp-colab)).

## Data availability

The embeddings of the proteins from the PROBE benchmark dataset, computed by DomainGO-prob, GO term prediction by DomainPFP on the CAFA3 dataset, trained DomainGO-prob model weights and Domain-PFP KNN models are accessible at <https://github.com/kiharalab/Domain-PFP> and made also available at <https://kiharalab.org/domainpfp/>.

## Reference

1. Cruz, L. M., Trefflich, S., Weiss, V. A., & Castro, M. A. A. (2017). Protein Function Prediction. In M. Kaufmann, C. Klinger, & A. Savelsbergh (Eds.), *Functional Genomics. Methods in Molecular Biology* (Vol. 1654, pp. 55–75). [https://doi.org/10.1007/978-1-4939-7231-9\\_5](https://doi.org/10.1007/978-1-4939-7231-9_5).
2. Radivojac, P., Clark, W. T., Oron, T. R., Schnoes, A. M., Wittkop, T., Sokolov, A., Graim, K., Funk, C., Verspoor, K., Ben-Hur, A., Pandey, G., Yunes, J. M., Talwalkar, A. S., Repo, S., Souza, M. L., Piovesan, D., Casadio, R., Wang, Z., Cheng, J., ... Friedberg, I. (2013). A large-scale evaluation of computational protein function prediction. *Nature Methods*, 10(3), 221–227. <https://doi.org/10.1038/nmeth.2340>.
3. Torres, M., Yang, H., Romero, A. E., & Paccanaro, A. (2021). Protein function prediction for newly sequenced organisms. *Nature Machine Intelligence*, 3(12), 1050–1060. <https://doi.org/10.1038/s42256-021-00419-7>.
4. Clark, W. T., & Radivojac, P. (2011). Analysis of protein function and its prediction from amino acid sequence. *Proteins: Structure, Function, and Bioinformatics*, 79(7), 2086–2096. <https://doi.org/10.1002/prot.23029>.
5. Hawkins, T., & Kihara, D. (2007). Function Prediction of Uncharacterized Proteins. *Journal of Bioinformatics and Computational Biology*, 05(01), 1–30. <https://doi.org/10.1142/S0219720007002503>.
6. Hawkins, T., Chitale, M., Luban, S., & Kihara, D. (2009). PFP: Automated prediction of gene ontology functional annotations with confidence scores using protein sequence data. *Proteins: Structure, Function, and Bioinformatics*, 74(3), 566–582. <https://doi.org/10.1002/prot.22172>.
7. Dawson, N. L., Orengo, C., & Gáspári, Z. (2020). *Assessing Protein Function Through Structural Similarities with CATH* (pp. 43–57). [https://doi.org/10.1007/978-1-0716-0270-6\\_4](https://doi.org/10.1007/978-1-0716-0270-6_4).
8. Gligorijević, V., Renfrew, P. D., Kosciolék, T., Leman, J. K., Berenberg, D., Vatanen, T., Chandler, C., Taylor, B. C., Fisk, I. M., Vlamakis, H., Xavier, R. J., Knight, R., Cho, K., & Bonneau, R. (2021). Structure-based protein function prediction using graph convolutional networks. *Nature Communications*, 12(1), 3168. <https://doi.org/10.1038/s41467-021-23303-9>.
9. Kagaya, Y., Flannery, S. T., Jain, A., & Kihara, D. (2022). ContactPFP: Protein Function Prediction Using Predicted Contact Information. *Frontiers in Bioinformatics*, 2. <https://doi.org/10.3389/fbinf.2022.896295>.
10. Jain, A., & Kihara, D. (2019). Phylo-PFP: improved automated protein function prediction using phylogenetic distance of distantly related sequences. *Bioinformatics*, 35(5), 753–759. <https://doi.org/10.1093/bioinformatics/bty704>.
11. Sahraeian, S. M., Luo, K. R., & Brenner, S. E. (2015). SIFTER search: a web server for accurate phylogeny-based protein function prediction. *Nucleic Acids Research*, 43(W1), W141–W147. <https://doi.org/10.1093/nar/gkv461>.

12. Forslund, K., & Sonnhammer, E. L. L. (2008). Predicting protein function from domain content. *Bioinformatics*, 24(15), 1681–1687. <https://doi.org/10.1093/bioinformatics/btn312>.
13. Messih, M. A., Chitale, M., Bajic, V. B., Kihara, D., & Gao, X. (2012). Protein domain recurrence and order can enhance prediction of protein functions. *Bioinformatics*, 28(18), i444–i450. <https://doi.org/10.1093/bioinformatics/bts398>.
14. Rojano, E., Perkins, J. R., Sillitoe, I., Orengo, C., García Ranea, J. A., & Seoane, P. (2019). Associating Protein Domains with Biological Functions: A Tripartite Network Approach. *IWBBIO 2019: Bioinformatics and Biomedical Engineering*, 155–164. [https://doi.org/10.1007/978-3-030-17935-9\\_15](https://doi.org/10.1007/978-3-030-17935-9_15).
15. Zhao, B., Zhang, Z., Jiang, M., Hu, S., Luo, Y., & Wang, L. (2020). NPF:network propagation for protein function prediction. *BMC Bioinformatics*, 21(1), 355. <https://doi.org/10.1186/s12859-020-03663-7>.
16. Yao, S., You, R., Wang, S., Xiong, Y., Huang, X., & Zhu, S. (2021). NetGO 2.0: improving large-scale protein function prediction with massive sequence, text, domain, family and network information. *Nucleic Acids Research*, 49(W1), W469–W475. <https://doi.org/10.1093/nar/gkab398>.
17. You, R., Zhang, Z., Xiong, Y., Sun, F., Mamitsuka, H., & Zhu, S. (2018). GOLabeler: improving sequence-based large-scale protein function prediction by learning to rank. *Bioinformatics*, 34(14), 2465–2473. <https://doi.org/10.1093/bioinformatics/bty130>.
18. Kulmanov, M., & Hoehndorf, R. (2019). DeepGOPlus: improved protein function prediction from sequence. *Bioinformatics*. <https://doi.org/10.1093/bioinformatics/btz595>.
19. Cao, Y., & Shen, Y. (2021). TALE: Transformer-based protein function Annotation with joint sequence–Label Embedding. *Bioinformatics*, 37(18), 2825–2833. <https://doi.org/10.1093/bioinformatics/btab198>.
20. Wan, C., & Jones, D. T. (2020). Protein function prediction is improved by creating synthetic feature samples with generative adversarial networks. *Nature Machine Intelligence*, 2(9), 540–550. <https://doi.org/10.1038/s42256-020-0222-1>.
21. You, R., Yao, S., Mamitsuka, H., & Zhu, S. (2021). DeepGraphGO: graph neural network for large-scale, multispecies protein function prediction. *Bioinformatics*, 37(Supplement\_1), i262–i271. <https://doi.org/10.1093/bioinformatics/btab270>.
22. Kulmanov, M., & Hoehndorf, R. (2022). DeepGOZero: improving protein function prediction from sequence and zero-shot learning based on ontology axioms. *Bioinformatics*, 38(Supplement\_1), i238–i245. <https://doi.org/10.1093/bioinformatics/btac256>.
23. Bonetta, R., & Valentino, G. (2020). Machine learning techniques for protein function prediction. *Proteins: Structure, Function, and Bioinformatics*, 88(3), 397–413. <https://doi.org/10.1002/prot.25832>.

24. Ibtehaz, N., & Kihara, D. (2023). Application of Sequence Embedding in Protein Sequence-Based Predictions. In *Machine Learning in Bioinformatics of Protein Sequences* (pp. 31–55). WORLD SCIENTIFIC. [https://doi.org/10.1142/9789811258589\\_0002](https://doi.org/10.1142/9789811258589_0002).
25. Doerks, T., Copley, R. R., Schultz, J., Ponting, C. P., & Bork, P. (2002). Systematic Identification of Novel Protein Domain Families Associated with Nuclear Functions. *Genome Research*, 12(1), 47–56. <https://doi.org/10.1101/gr.203201>.
26. Burge, S., Kelly, E., Lonsdale, D., Mutowo-Muellenet, P., McAnulla, C., Mitchell, A., Sangrador-Vegas, A., Yong, S.-Y., Mulder, N., & Hunter, S. (2012). Manual GO annotation of predictive protein signatures: the InterPro approach to GO curation. *Database*, 2012(0), bar068–bar068. <https://doi.org/10.1093/database/bar068>.
27. Camon, E. B., Barrell, D. G., Dimmer, E. C., Lee, V., Magrane, M., Maslen, J., Binns, D., & Apweiler, R. (2005). An evaluation of GO annotation retrieval for BioCreAtIVe and GOA. *BMC Bioinformatics*, 6(S1), S17. <https://doi.org/10.1186/1471-2105-6-S1-S17>.
28. Liu, X., Zhang, F., Hou, Z., Mian, L., Wang, Z., Zhang, J., & Tang, J. (2021). Self-supervised Learning: Generative or Contrastive. *IEEE Transactions on Knowledge and Data Engineering*, 1–1. <https://doi.org/10.1109/TKDE.2021.3090866>.
29. Wei, Q., Khan, I. K., Ding, Z., Yerneni, S., & Kihara, D. (2017). NaviGO: interactive tool for visualization and functional similarity and coherence analysis with gene ontology. *BMC Bioinformatics*, 18(1), 177. <https://doi.org/10.1186/s12859-017-1600-5>.
30. Elnaggar, A., Heinzinger, M., Dallago, C., Rehawi, G., Wang, Y., Jones, L., Gibbs, T., Feher, T., Angerer, C., Steinegger, M., Bhowmik, D., & Rost, B. (2021). ProtTrans: Towards Cracking the Language of Lifes Code Through Self-Supervised Deep Learning and High Performance Computing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–1. <https://doi.org/10.1109/TPAMI.2021.3095381>.
31. Rives, A., Meier, J., Sercu, T., Goyal, S., Lin, Z., Liu, J., Guo, D., Ott, M., Zitnick, C. L., Ma, J., & Fergus, R. (2021). Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15). <https://doi.org/10.1073/pnas.2016239118>.
32. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is All you Need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems*. Curran Associates, Inc.
33. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1*, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>.
34. Zhou, N., Jiang, Y., Bergquist, T. R., Lee, A. J., Kacsoh, B. Z., Crocker, A. W., Lewis, K. A., Georghiou, G., Nguyen, H. N., Hamid, M. N., Davis, L., Dogan, T., Atalay, V., Rifaioglu, A. S., Dalkiran, A., Cetin Atalay, R., Zhang, C., Hurto, R. L., Freddolino, P. L.,

... Friedberg, I. (2019). The CAFA challenge reports improved protein function prediction and new functional annotations for hundreds of genes through experimental screens. *Genome Biology*, 20(1), 244. <https://doi.org/10.1186/s13059-019-1835-8>.

35. Bateman, A., Martin, M.-J., Orchard, S., Magrane, M., Agiveto, R., Ahmad, S., Alpi, E., Bowler-Barnett, E. H., Britto, R., Bursteinas, B., Bye-A-Jee, H., Coetzee, R., Cukura, A., Da Silva, A., Denny, P., Dogan, T., Ebenezer, T., Fan, J., Castro, L. G., ... Teodoro, D. (2021). UniProt: the universal protein knowledgebase in 2021. *Nucleic Acids Research*, 49(D1), D480–D489. <https://doi.org/10.1093/nar/gkaa1100>.

36. Blum, M., Chang, H.-Y., Chuguransky, S., Grego, T., Kandasamy, S., Mitchell, A., Nuka, G., Paysan-Lafosse, T., Qureshi, M., Raj, S., Richardson, L., Salazar, G. A., Williams, L., Bork, P., Bridge, A., Gough, J., Haft, D. H., Letunic, I., Marchler-Bauer, A., ... Finn, R. D. (2021). The InterPro protein families and domains database: 20 years on. *Nucleic Acids Research*, 49(D1), D344–D354. <https://doi.org/10.1093/nar/gkaa977>.

37. Kingma, D. P., & Ba, J. (2015, December 22). Adam: A Method for Stochastic Optimization. *3rd International Conference for Learning Representations*.

38. Littmann, M., Heinzinger, M., Dallago, C., Olenyi, T., & Rost, B. (2021). Embeddings from deep learning transfer GO annotations beyond homology. *Scientific Reports*, 11(1), 1160. <https://doi.org/10.1038/s41598-020-80786-0>.

39. Aggarwal, C. C., Hinneburg, A., & Keim, D. A. (2001). On the Surprising Behavior of Distance Metrics in High Dimensional Space. In Jan Van den Bussche & Victor Vianu (Eds.), *Database Theory - ICDT 2001* (pp. 420–434). Springer Berlin Heidelberg. [https://doi.org/10.1007/3-540-44503-X\\_27](https://doi.org/10.1007/3-540-44503-X_27).

40. Schlicker, A., Domingues, F. S., Rahnenführer, J., & Lengauer, T. (2006). A new measure for functional similarity of gene products based on Gene Ontology. *BMC Bioinformatics*, 7(1), 302. <https://doi.org/10.1186/1471-2105-7-302>.

41. Mitchell, A., Chang, H.-Y., Daugherty, L., Fraser, M., Hunter, S., Lopez, R., McAnulla, C., McMenamin, C., Nuka, G., Pesseat, S., Sangrador-Vegas, A., Scheremetjew, M., Rato, C., Yong, S.-Y., Bateman, A., Punta, M., Attwood, T. K., Sigrist, C. J. A., Redaschi, N., ... Finn, R. D. (2015). The InterPro protein families database: the classification resource after 15 years. *Nucleic Acids Research*, 43(D1), D213–D221. <https://doi.org/10.1093/nar/gku1243>.

42. Hill, D. P., Smith, B., McAndrews-Hill, M. S., & Blake, J. A. (2008). Gene Ontology annotations: what they mean and where they come from. *BMC Bioinformatics*, 9(S5), S2. <https://doi.org/10.1186/1471-2105-9-S5-S2>.

43. Abrahamson, M., Alvarez-Fernandez, M., & Nathanson, C.-M. (2003). Cystatins. *Biochemical Society Symposium*, 70, 179–199. <https://doi.org/10.1042/bss0700179>.

44. Gaudet, P., Livstone, M. S., Lewis, S. E., & Thomas, P. D. (2011). Phylogenetic-based propagation of functional annotations within the Gene Ontology consortium. *Briefings in Bioinformatics*, 12(5), 449–462. <https://doi.org/10.1093/bib/bbr042>.

45. Barrett, A. J., & Rawlings, N. D. (2001). Evolutionary lines of cysteine peptidases. *Biological Chemistry*, 382(5), 727–733. <https://doi.org/10.1515/BC.2001.088>.
46. Bizzarri, C., Beccari, A. R., Bertini, R., Cavicchia, M. R., Giorgini, S., & Allegretti, M. (2006). ELR+ CXC chemokines and their receptors (CXC chemokine receptor 1 and CXC chemokine receptor 2) as new therapeutic targets. *Pharmacology & Therapeutics*, 112(1), 139–149. <https://doi.org/10.1016/j.pharmthera.2006.04.002>.
47. Unsal, S., Atas, H., Albayrak, M., Turhan, K., Acar, A. C., & Doğan, T. (2022). Learning functional properties of proteins with language models. *Nature Machine Intelligence*, 4(3), 227–245. <https://doi.org/10.1038/s42256-022-00457-9>.
48. Heinzinger, M., Elnaggar, A., Wang, Y., Dallago, C., Nechaev, D., Matthes, F., & Rost, B. (2019). Modeling aspects of the language of life through transfer-learning protein sequences. *BMC Bioinformatics*, 20(1), 723. <https://doi.org/10.1186/s12859-019-3220-8>.
49. Rao, R., Bhattacharya, N., Thomas, N., Duan, Y., Chen, X., Canny, J., Abbeel, P., & Song, Y. S. (2019). Evaluating Protein Transfer Learning with TAPE. *Advances in Neural Information Processing Systems*, 32, 9689–9701. <https://pubmed.ncbi.nlm.nih.gov/33390682>.
50. Choy, C. T., Wong, C. H., & Chan, S. L. (2018). Infer related genes from large scale gene expression dataset with embedding. *BioRxiv*, 362848. <https://doi.org/10.1101/362848>.
51. Rao, R. M., Liu, J., Verkuil, R., Meier, J., Canny, J., Abbeel, P., Sercu, T., & Rives, A. (2021). MSA Transformer. In M. Meila & T. Zhang (Eds.), *Proceedings of the 38th International Conference on Machine Learning* (Vol. 139, pp. 8844–8856). PMLR. <https://proceedings.mlr.press/v139/rao21a.html>.
52. Alley, E. C., Khimulya, G., Biswas, S., AlQuraishi, M., & Church, G. M. (2019). Unified rational protein engineering with sequence-based deep representation learning. *Nature Methods*, 16(12), 1315–1322. <https://doi.org/10.1038/s41592-019-0598-1>.
53. Yang, K. K., Wu, Z., Bedbrook, C. N., & Arnold, F. H. (2018). Learned protein embeddings for machine learning. *Bioinformatics*, 34(15), 2642–2648. <https://doi.org/10.1093/bioinformatics/bty178>.
54. Asgari, E., & Mofrad, M. R. K. (2015). Continuous Distributed Representation of Biological Sequences for Deep Proteomics and Genomics. *PLOS ONE*, 10(11), e0141287. <https://doi.org/10.1371/journal.pone.0141287>.
55. Steinegger, M., & Söding, J. (2017). MMseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nature Biotechnology*, 35(11), 1026–1028. <https://doi.org/10.1038/nbt.3988>.
56. Buchfink, B., Xie, C., & Huson, D. H. (2015). Fast and sensitive protein alignment using DIAMOND. *Nature Methods*, 12(1), 59–60. <https://doi.org/10.1038/nmeth.3176>.
57. Szklarczyk, D., Gable, A. L., Nastou, K. C., Lyon, D., Kirsch, R., Pyysalo, S., Doncheva, N. T., Legeay, M., Fang, T., Bork, P., Jensen, L. J., & von Mering, C. (2021). The STRING database in 2021: customizable protein–protein networks, and functional characterization

of user-uploaded gene/measurement sets. *Nucleic Acids Research*, 49(D1), D605–D612. <https://doi.org/10.1093/nar/gkaa1074>.

58. Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed Representations of Words and Phrases and their Compositionality. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, & K.Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems*. Curran Associates, Inc.

59. Chitale, M., Hawkins, T., Park, C., & Kihara, D. (2009). ESG: extended similarity group method for automated protein function prediction. *Bioinformatics*, 25(14), 1739–1745. <https://doi.org/10.1093/bioinformatics/btp309>.

60. Yuan, Q., Xie, J., Xie, J., Zhao, H., & Yang, Y. (2023). Fast and accurate protein function prediction from sequence through pretrained language model and homology-based label diffusion. *Briefings in Bioinformatics*, 24(3). <https://doi.org/10.1093/bib/bbad117>.

61. He, K., Chen, X., Xie, S., Li, Y., Dollar, P., & Girshick, R. (2022). Masked Autoencoders Are Scalable Vision Learners. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2022-June. <https://doi.org/10.1109/CVPR52688.2022.01553>.

62. Lai, B., & Xu, J. (2022). Accurate protein function prediction via graph attention networks with predicted structure information. *Briefings in Bioinformatics*, 23(1). <https://doi.org/10.1093/bib/bbab502>.

63. Doersch, Carl; Zisserman, Andrew (2017). [Multi-task Self-Supervised Visual Learning](#). *2017 IEEE International Conference on Computer Vision (ICCV)*. IEEE. pp. 2070–2079.

## Acknowledgments

This work was partly supported by the National Science Foundation (DBI2003635, DBI2146026, IIS2211598, DMS2151678, CMMI1825941, and MCB1925643) and by the National Institutes of Health (R01GM133840, 3R01 GM133840-02S1).

## Author contributions

D.K. conceived the study. N.I. designed and implemented the domain-based protein embedding and Domain-PFP. N.I. performed the computation. N.I. and D.K. analyzed the data. Y.K. participated in function evaluation on the CAFA3 dataset. N.I. drafted the manuscript and D.K. edited it. All authors read and approved the manuscript.

## Competing interests

The authors declare no competing interests.