RESEARCH ARTICLE

PROTEINS WILEY

Learning a functional grammar of protein domains using natural language word embedding techniques

Daniel W. A. Buchan [ORCID] | **David T. Jones**

Department of Computer Science, University College London, London, UK

Correspondence
David T. Jones, Department of Computer Science, University College London, Gower Street, London WC1E 6BT, UK.
Email: d.t.jones@ucl.ac.uk

Abstract

In this paper, using Word2vec, a widely-used natural language processing method, we demonstrate that protein domains may have a learnable implicit semantic "meaning" in the context of their functional contributions to the multi-domain proteins in which they are found. Word2vec is a group of models which can be used to produce semantically meaningful embeddings of words or tokens in a fixed-dimension vector space. In this work, we treat multi-domain proteins as "sentences" where domain identifiers are tokens which may be considered as "words." Using all InterPro (Finn et al. 2017) pfam domain assignments we observe that the embedding could be used to suggest putative GO assignments for Pfam (Finn et al. 2016) domains of unknown function.

KEYWORDS
function prediction, machine learning, protein domains, semantic embedding, word2vec

1 | INTRODUCTION

Word2vec[1] is a group of models which can be used to learn the embeddings of words in a continuous fixed-dimension vector space, given a corpus of sentences as training data. Often Natural Language Processing (NLP) tasks consider words as sets of unrelated tokens, subjecting them to no-more rigorous analysis than simple frequency counting. While this is mathematically and computationally convenient, it ignores the fact that most words have degrees of similarity, such as verbs with differing tenses, adverbs with differing endings or words which share the same suffixes. Word2vec aims to produce embeddings of words in a vector space where distance in the vector space correctly encodes the degree to which words or terms are similar or can be used in similar semantic context. Although a great degree has been written about these methods it remains unclear exactly why these models are so performant.[2] Nevertheless, they do show good performance in the task of clustering words with related semantic meaning, and interested readers should consult the original paper for further details of the model.[1] Since lexical word embeddings have become popular, they have been adapted and applied directly to protein and gene sequences. prot2vec, gene2vec, and seq2vec are examples of such methods.[3,4] Another prior application of Word2vec is the

work of Viehweger,[5] applying protein domain embeddings as a method to encode whole genomes.

Proteins are often composed of discrete domains, and these can either be conceptualized as sub-sequences of independent protein sequences which share homology (and by extension evolutionary origin),[6] or alternatively, domains may be considered structurally, where they are subsections of the proteins which are compact, independently folding and observed to be shared between a variety of proteins.[7-9] An extension of this observation, that proteins can be decomposed into sets of domains, is the hypothesis that domains act as sub-functional units and when composed together, a protein's given combination of domains is what gives rise to the protein's overall specific function[10,11] In the following study, we show that protein domains can be embedded in a "semantically" meaningful vector space and that this embedding space reflects meaningful information about the functional roles (in terms of GO term assignments) of the individual protein domains.

Protein function prediction has received a great deal of attention in the preceding 20 years[12] and a great number of function prediction methods have been developed. Many of these make use of sequence comparison and some manner of nearest neighbor functional assignment.[13,14] As the field has progressed work has been carried out to integrate more sophisticated statistical methods and models with

many contemporary methods leveraging machine learning with ensemble or meta-prediction methodologies. Current state of the art in protein function is measured by the Critical Assessment in Function Annotation (CAFA) community experiment.[15] In this experiment groups, attempt to predict experimentally validated Gene Ontology (GO) terms[16] over a blind set of unannotated protein sequences. The most successful methods in CAFA employ a wide variety of predictive methodologies. Among the most common are methods which integrate data and annotations from a wide variety of sources including blast searches, protein–protein interaction networks, multiple sequence alignment analysis, sequence analysis, expression data, and many more.[17-20] A number of other successful methodologies eschew integrating heterogenous data sources and make use of more focused analyses, such as phylogenetic analysis,[21] literature analysis,[22] MSA analysis,[23] domain function analysis [24, 25]. Information about protein domains is typically only included indirectly, such as in the methods INGA and PFPDB which make use of Pfam to derive phylogenetic or domain architecture patterns. Less common are methods which directly attempt to annotate domains with function and then leverage this information for function prediction. Both the SIFTER, CATH-Funfam,[24] and Superfamily-dcGO[25] methods in CAFA were successful methods which directly leverage such domain function annotations. It is clear that understanding the relationship between protein domains and their function can make a significant contribution to accurate function prediction. Nevertheless, even with the wide range of prediction methodologies, performance and progress in the CAFA experiment indicates that protein function prediction remains a challenging problem in the field of bioinformatics.

In the following work, we discuss the use of Word2vec in protein domain embedding. We prepare such a domain embedding and attempt to explore the its properties to discern whether such embeddings encode biological information that may be useful in either a predictive or analytic context. Such embeddings may be a useful adjuncts or input features in protein function prediction as it may give a homology-free way to characterize and functionally cluster protein domains. At the end of the paper we note that such an embedding could be used for the purposes of homology-free GO term inheritance and we show a naïve application of this for Pfam Domains of Unknown Function.

2 | METHOD

2.1 | Datasets

InterPro 62[26] was downloaded along with the associated GO and protein domain assignments. The files were parsed to extract only the eukaryotic proteins and their GO and Pfam domain assignments. This work looks only at eukaryotic proteins as there are few examples of proteins with multiple domains with independent evolutionary histories in the bacterial and archaeal kingdoms, as such little domain context information would be available for proteins from those kingdoms. Only GO assignments with the following evidence codes were retained: EXP, IBA, IDA, IEP, IGC, IGI, IMP, and IPI. These are

(respectively); inferred from EXPeriment, Inferred from Biological Aspect of ancestor, inferred from Direct Assay, Inferred from Expression Pattern, Inferred from Genomic Context, Inferred from Genetic Interaction, Inferred from Mutant Phenotype and Inferred from Physical Interaction. This eliminates all the high throughput and more tenuous computational annotation assignments. The resulting dataset contains 9 030 650 eukaryotic proteins, which have domain assignments over 11 355 of the available Pfam domain families and these proteins are associated with annotations from 2358 GO Terms.

Not all regions within each protein have been assigned to domains (see Table 1). In large part because not all domains are known and assigned but also because many eukaryotic proteins possess regions of intrinsic disorder,[27] regions of low complexity or coiled-coiled sequences. All such unassigned regions were compiled (see below). As Word2vec analyses words based on the semantic context of neighboring words representing unassigned regions in our corpus could contain important domain context information, and so we wished to preserve this.

These data were then used to derive which Pfam domains are seen to be associated to which GO terms. For every Pfam domain, we associated all GO terms assigned to all the proteins the Pfam domain was observed in. This assigns a varied bag of GO terms to each Pfam domain and this bag of terms can be viewed as representing the spectrum of observed functional diversity for that Pfam domain.

2.2 | Unassigned sequence region assignments

The sequence database for InterPro 62 was masked for both coiled coil and low complexity regions using pfilt.[28] Disordered regions were derived directly from the existing InterPro disorder annotations. Gap regions which did not contain disorder annotations, coiled-coil or low complexity sequence were assigned given the length of the unassigned regions. These remaining gap regions were binned into size bins based on their lengths (see Figure 1). The majority of gap regions are around 100 residues in length, as the typical structural domain size is around 100 residues five gap types were created to represent unassigned regions of various sizes which are approximate multiples of the typical domain size, see Table 2. All non-domain regions: gaps, disordered, low complexity, and coiled-coil regions were

TABLE 1 Table of the total residue counts across the eukaryotic Interpro protein set and the number of residues assigned to each class of domain or region

Class	Residue count	Percentage
Total	5 001 517 961	—
Domains	1 256 832 058	25.1
Gaps	3 405 089 896	68.1
Disordered	167 103 753	3.3
Coiled coil	3 309 167	0.06
Low complexity	2 079 334	0.04

TABLE 2 Names and sizes of gap pseudo-domains and the number of interpro proteins where we observe at least one of these regions

Gap region ID	Size (residues)	Protein count
G100	20–100	4 234 931
G200	101–200	2 635 225
G300	201–300	1 168 553
G400	301–400	575 517
G500	401–>500	926 673

then compiled as a set of adjunct domain-like sequence regions to complement the Pfam domain assignments.

2.3 | Building the word embedding

To build Word2vec embeddings, we treat protein sequences and their domain assignments as "sentences." The Pfam IDs and other sequence region assignments are used as tokens/pseudo-words in such a pseudo-sentence. For instance, a typical protein may be converted to a sentence such as "PF00170 PF003534 G200 LowComplexity PF00678." Which would indicate two leading Pfam domains followed by a gap region up to 200 residues, a region of low complexity sequence finally terminating in a Pfam domain (see Figure 2). We compile such sentences for every eukaryotic protein in InterPro62 and this set of sentences becomes the corpus we use to create the word embedding.

Python library gensim (https://radimrehurek.com/gensim/) was used to create the word2vec model from the corpus. The size parameter was set to 100, representing the dimensionality of the vector space to project the words in to. The minimum word count was set to 0, indicating that all words would be positioned in the vector space.
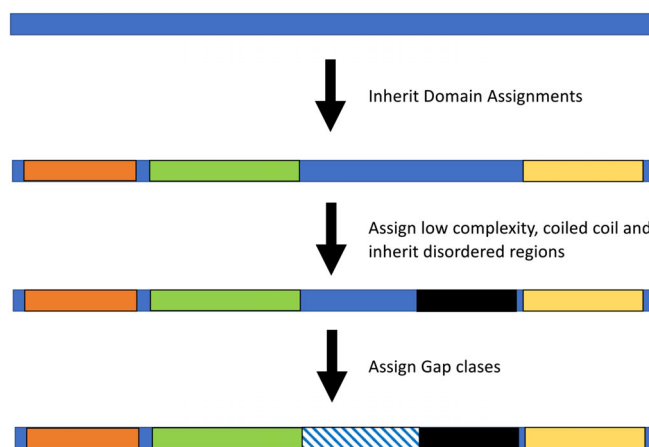
FIGURE 2 The example of the domain and sequence region assignment. Pfam domains and disorder regions are derived from InterPro annotations. Low complexity and coiled coil regions are calculated by pfilt and gaps are assigned given their size [Color figure can be viewed at wileyonlinelibrary.com]

This ensures that all domains, including important infrequent ones are considered, also the embedding uses the skip-gram algorithm and model to build the embedding. The goal of Word2vec is to learn the weights in the hidden layer of a simple neural network, this hidden layer is an n by m matrix, where n is the number of input words in the corpus and m is the size parameter (eg, 100). To train these weights the network is given a training task, the skip-gram task, which asks the network to predict, for each word in turn to output the probability that other words from the corpus are near to the input word (ie, within a given window size, in this instance a window of 5). Once the training is complete the output probabilities are discarded and only the weights of the hidden layer are retained as this matrix is regarded as the word embedding. It is possible to develop alternative training

tasks to learn the embedding matrix. A target behavior of Word2vec is that words which fulfill similar semantic roles should be near one another in the embedding and it is believed that the skip-gram task, by having the network learn about which words are local to one another, in turn is encoding this information in the weights of the hidden layer.

The embedding process is illustrated full in Figure 3. For the benchmark below an all-against-all distance matrix of domains was derived.

2.4 | Benchmark

We are interested in whether Word2vec embeds Pfam domains in a manner which is biologically meaningful. This would in turn would indicate that there is some manner of semantic meaning in the positioning or sequence context for protein domains. To investigate the embedding, initially we attempted to project the domain vectors into three dimensions (data not shown) using multi-dimensional scaling. However, the resulting projection did not yield any trivially interpretable result.

An alternative means of investigating whether the embedding is biologically meaningful would be to establish if functionally related domains are placed near one another in the embedding. To investigate this, we assigned GO terms to the Pfam domains. This was carried out by allowing Pfam domains to inherit all GO terms assigned to the proteins each Pfam domain is observed in. Pfam domains inherit an average of 19.6 GO terms, although some domains may have upwards of 100 terms associated, see Figure 4. Although this is somewhat imprecise, as GO annotations reflect protein functions

rather than domain function, each domain's "bag" of GO terms will reflect the functional diversity for the contexts a domain is observed in. A total of 2358 GO terms were assigned over the 11 355 Pfam domains observed in the eukaryotic proteins. These assignments could then be used for a nearest-neighbor benchmark test.

3 | RESULTS

3.1 | Nearest-neighbor performance

Performance in nearest neighbor functional annotation was calculated to assess whether the vector embedding of domains displayed any meaningful structure. That is, domains with similar functionality were placed near one another in the embedding. Each domain was in turn considered by inheriting the GO terms from its k-nearest neighbors and comparing these predicted terms to the known terms assigned via InterPro annotations. Table 3 gives the precision and Matthew's Correlation Coefficients (MCC) scores for the nearest-neighbor benchmark. The MCC value indicates the predicted terms are non-random (greater than 0) which in turn suggests that there is some meaningful structure in the embedding of domains in a vector space. Mean accuracy is high and this is a consequence of there being a very large number of GO terms where typically only a few (relatively) are used to annotate any given protein or domain. This in turn means any given domain has very large numbers of true negatives most of which are called correctly. As K is increased recall also increases as the bag of assigned terms gets very large but this comes at the cost of a sharply declining precision.
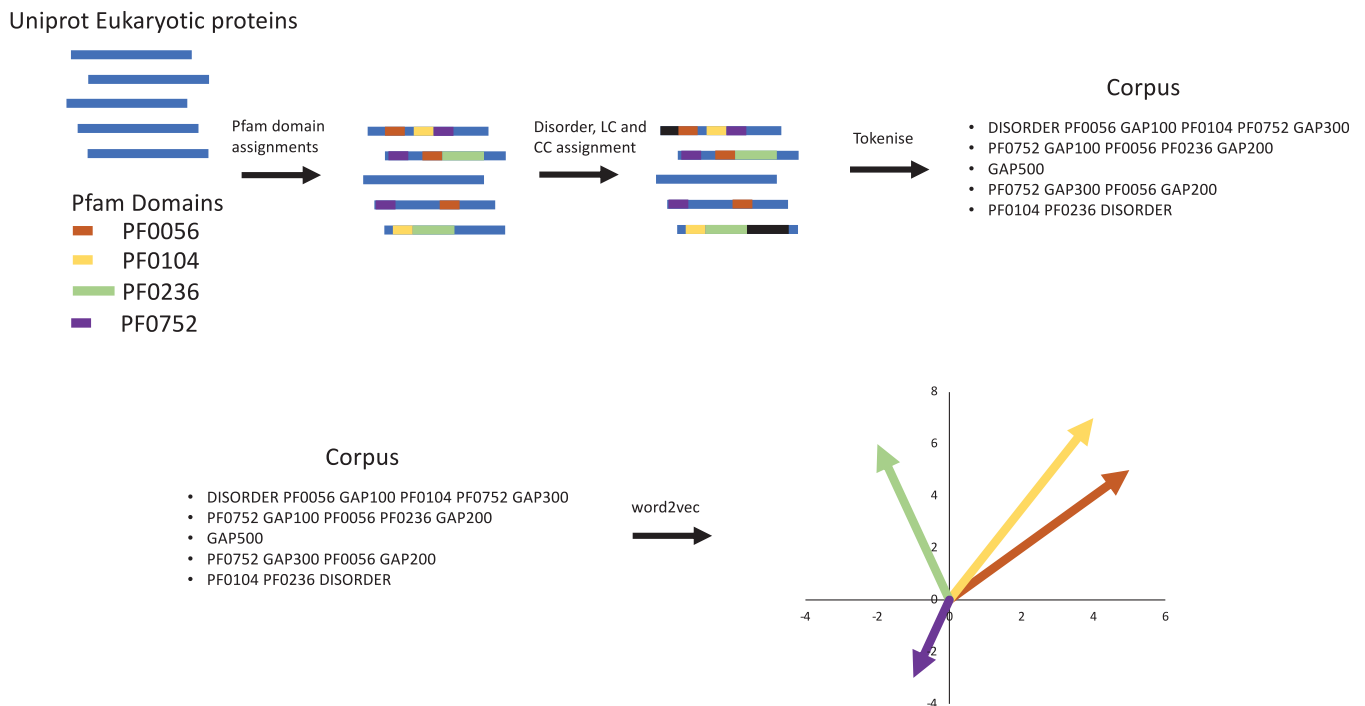
FIGURE 3 Compiling protein "sentences." InterPro compiles assignments of domains on Uniprot protein sequences. We take only the Pfam domain assignments the InterPro database stores and complement those with the assignments of disorder and our own low complexity (LC) and coiled-coil (CC) region assignments. These are then tokenized to create a corpus of "sentences." The corpus can then be used as input to Word2vec. The output is a vector space which places each token at a point within that space, here stylized in two-dimensional. Tokens which appear in similar syntactic contexts in the corpus should be placed near one another in the vector space

Word2vec is designed to embed human language words in a vector space such that words which occur in similar semantic contexts are close to one another in the vector space. That our domain embedding is non-random implies that multidomain proteins exhibit some form of semantic structure. That is, certain domains appear in contexts near or adjacent to other domains and it may be possible to learn grammar-like rules which govern this.

It is worth noting that increasing the number of neighbors (increasing K) from which functional roles can be inherited degrades performance in this function-annotation task. Domains are typically involved in a large number of possible different protein functions. By increasing the number of neighbors, GO terms can be inherited from the number of false positives is greatly increased and so performance degrades.

3.2 | Per ontology results

MCC values were also calculated for each of the three GO Ontologies (see Table 4). Of the 2358 GO terms used to annotate eukaryotic sequences in InterPro: 1,018 are from the molecular function ontology, 1,026 are from the biological process ontology and 314 from the cellular component ontology. The MCC values indicate different functional inheritance performance for each ontology with. In the context of the vector embedding this may imply that the simple syntax contained in the domain orderings contains some additional information about where a protein is located within the cell. Given the results of the previous CAFA experiment[15] it may, more simply, be that cellular component prediction is an easier task.

In general, we believe the MCC calculated may underestimate the quality of the domain embedding. Given the figures in Table 1 we see that nearly 70% of the proteins are gap regions. This indicates many domain assignments and domain types may be missing. We would expect with better domain coverage we would also have a more robust and biologically meaningful embedding.

Alongside this, using GO assignments to genes to annotate domains is inherently noisy. GO annotations may not be good descriptors of the specific role a domain plays in a given protein. For instance, GO:0051987 (Chaperone Binding), assigned to 92 Pfam domains, might be regarded as property or function of a whole protein rather than just a specific domain. An alternative issue is illustrated by Pfam domain PF00176 which is assigned both GO:0009916 (alternative oxidase activity) and GO:0001733 (galactosylceramide sulfotransferase activity). These assignments come through differing InterPro proteins but represent different catalytic reaction chemistries this domain is unlikely to possess both of these. Within the context of a multidomain proteins, domains provide specific sub-functionality such as providing catalytic sites, presenting one or more small molecule binding sites, providing membrane anchoring and so forth. It seems plausible if domains were annotated at a level, that better reflected these more specific sub-functional roles (rather than the protein's role), then the nearest-neighbor assignment would return better results. The lack of a computer readable "domain ontology" remains a barrier for large scale studies of domain functionality and evolution.

3.3 | Comparison to first-order Markov representation

As sets of domains are sequences of symbols or states, it is possible to represent the information contained in the corpus of domain strings as a Markov process. We also investigated whether the Word2vec domain embedding was a more robust representation of the information contained in the domain corpus than a first-order Markov process. Parsing the corpus of proteins, a table of the transition probabilities of all domains against all domains was prepared. A given domain's immediate context can be read from the table as the rows give the probabilities of

FIGURE 4 Distribution of Gene Ontology term assignments

TABLE 4 Matthew's correlation coefficients (MCC) values for nearest neighbor inheritance of Gene Ontology (GO) terms, calculated for each separate GO ontology

	k			
Ontology	**1**	**3**	**5**	**10**
Biological process	0.27	0.20	0.19	0.17
Molecular function	0.30	0.23	0.22	0.19
Cellular component	0.33	0.22	0.22	0.20

TABLE 3 Mean precision and accuracy and Matthew's Correlation Coefficients (MCC) given nearest neighbor inheritance of Gene Ontology terms

k-Nearest neighbors	Mean precision	Mean recall	Mean accuracy	Mean MCC
1	0.33	0.30	0.99	0.28
3	0.23	0.42	0.98	0.28
5	0.18	0.49	0.98	0.26
10	0.12	0.57	0.96	0.23

the following domain and columns indicate the probabilities of preceding domains. It follows that pairs of domains which share both similar row and column vectors are used in the same context in multidomain proteins. A distance matrix of Euclidean distances between all domains' vectors was prepared and the nearest-neighbor assignment analysis was described above was performed, the results can be seen in Table 5. These results indicate that the Word2vec domain embedding is substantially better at encoding the biological information contained in the corpus of multidomain proteins. The comparison may not be completely equivalent, Markov probabilities take in to account only the preceding symbol (or symbols in higher order chains) whereas the Word2vec method considers a window of tokens around each domain, and this feature is likely a better match for modeling protein domain placement. Considering the incoming and outgoing probabilities for each domain could be considered equivalent to considering a window of three domains. The default window size for Word2vec is 5. This comparison may under report the performance of a Markov process to model this data. However, the corpus of multidomain proteins only contains a tiny fraction of the possible 3- and 5-mers of domains and with many unassigned regions it getting accurate probabilities may not be possible.

3.4 | Vector arithmetic on the domain embeddings

One observation of semantic embeddings of natural languages is that arithmetic operations on the vectors frequently have semantic or lexical meanings, one classic example being:

$$King - Man + Woman = Queen.$$

We wished to investigate if simple vector arithmetic or translations for the protein domain embedding might have similar lexical meaning.

In the King to Queen example (see Figure 5), subtracting Man from King takes you to a space in the embedding with the meaning of man "removed" such that adding the Woman vector will take you to Queen. We can perform similar vector subtractions for the domain embedding. In this context, we would treat a domain's set of GO terms as equivalent to its "meaning," although, as discussed, this is a lossy way to conceptualize the meaning of a domain. Nevertheless, if we subtract two domain vectors we would hope the third vector is in a space where the remaining set of GO terms is the set difference of the two domains.

TABLE 5 Comparison of Matthew's correlation coefficients (MCC) performance between first-order Markov encoding and the Word2vec embedding of the domain corpus

	Mean MCC	
k-Nearest Neighbors	Word2vec	Markov
1	0.28	0.13
5	0.28	0.14
5	0.26	0.14
10	0.23	0.11

We took the most common 20 Pfam domains, removing the one that is not present in eukaryotes and in turn subtracted all possible domain vectors. For the resulting third vector, we found the nearest domain and tested the GO term overlaps with the initial two domains. In nearly all cases the resulting domain has minimal GO term overlaps with its parents. It is clear that this operation moves us to a region in the vector space where the domains' "meaning" is profoundly altered, much as removing Man from King might be thought of as moving to a gender-neutral space. What is not clear is what is the functional meaning of this in protein domain terms.

To investigate whether we could find more meaningful movements in the vector space we looked instead for translations in the vector space between mutually exclusive binary annotations. King and Queen are typically used as mutually exclusive labels that straddle some conceptual binary assignment (ie, gender) and much the same is true of many GO terms. For instance, in the cellular component ontology annotation, terms such as intracellular and extracellular might be viewed as a similar mutually exclusive binary.

We chose three binary cellular component term pairs; intracellular (GO:0005622) vs extracellular (GO:0005615), nucleus (GO: 0005634) vs cytoplasm (GO: 0005737), and cytoplasm (GO: 0005737) vs transmembrane (GO: 0009279). For each pairing, we identified proteins with domains annotated exclusively with one term and not the other term. Then for the first term we calculated the vector which moves from the location of the domain with the first term to the closest domain annotated with the second term. As with the prior analysis not having a detailed domain ontology prevents us from knowing if this closest domain is the most appropriate domain to move to. This led to a population of translation vectors which we could test to measure if the translation from a domain with one term to a domain with the other term was always vector oriented in a similar direction. We compared all Intracellular to extracellular vectors in an all against all fashion and did the same for the other two pairs of terms (see Figure 6). If the translation is preserved in the vector space, we would expect that all the vectors to have a small angle of deflection between them. In the transmembrane case, there was no such alignment and no trend in the angles between the vectors. In both, the intracellular to extracellular and the nucleus to cytoplasmic cases, there is a clear distribution which peaks around 1.5 rad, indicating that in general the translation is commonly orthogonal and is not preserved in the vector space. This stands somewhat at odds with the prior observation that vector arithmetic which encodes semantic translations is a general property of these embeddings. The caveat to make here is that our embedding may not of high enough quality to perform this analysis productively. As noted above there may not be enough domain coverage to robustly place the domains in the embedding space. Alternatively when choosing the domain pairs, the closest paired domain may not be the correct domain to calculate the angle between either we have selected the wrong extant domain or the correct domain is yet to be added to Pfam.

However, the intracellular to extracellular histogram shows a small leading tail below 1 rad (see Figure 7) indicative of a small population of vectors which do approach alignment. And indeed, we are able to find small numbers of genes in InterPro which share Pfam domains and where the difference is a substitution of one or more intracellular

FIGURE 5 Example demonstrating semantically meaningful vector algebra. In A, four terms are placed in the vector space. If we subtract the **Man** vector from **King** (graph B), we move to an undefined point in the vector space. Adding the **Woman** vector C, moves to the **Queen** vector [Color figure can be viewed at wileyonlinelibrary.com]


FIGURE 6 Comparing translation vector from one binary Gene Ontology property to another. A, Putative vector embedding of intracellular (blue dots) and extracellular (orange crosses) labeled domains. B, Vectors which translate each intracellular domain to its closest extracellular labeled domain. C, Vectors are extracted and pooled D, angle between each vector is compared to find vectors that point in the same direction [Color figure can be viewed at wileyonlinelibrary.com]

annotated domains for extracellular domains. Two examples, such as G3I6X9 (solute carrier family 25 member 46) and A0A0L6WZ71 (glycogen debranching enzyme) or I3L0A0 (Human Transcript TMEM189-UBE2BV1) and G7Y5H3 (Ubiquitin-conjugating enzyme E2 L3), see Figure 8. The first pair, G3I6X9 and A0A0L6WZ71, have respectively extracellular and intracellular functions. The second pair; G7Y5H3 has a cytoplasmic function but it is less clear what the role of I3L0A0 might be. The fact that this appears to work in some limited cases may suggest that an embedding based on a dataset with much greater domain coverage might be more accurate.

3.5 | Domains of unknown function

As the Word2vec embedding has some meaningful structure with regards GO term inheritance we can also use a nearest neighbor approach to suggest putative sets of GO terms that each eukaryotic Pfam domain of unknown function (Pfam DUFs) may take part in. This allows a homology-free way to estimate GO assignments. Our corpus of eukaryotic genes contained annotations from 3918 DUFs. Using a single nearest neighbor inheritance method, 1292 of these domains could be assigned new GO terms (ie, their nearest neighbor in the embedding was annotated and was not a gap or other sequence region). On average each DUF gets 11 novel GO terms assigned.
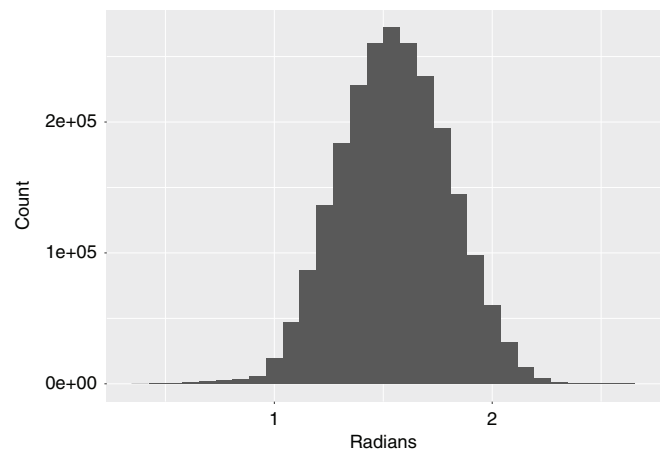

FIGURE 7 Histogram of transformation vector angles. For intracellular to extracellular

Surveying the GO assignments, we note that the mean ontology depth for each assigned term (ie, the shortest number of steps from an assigned term to the root of the ontology) is a depth on the graph 4.9 steps from the root of the ontology. The distribution of assigned term depths is also somewhat positively skewed (data not shown).

FIGURE 8 Diagram of intra/extra-cellular domain swaps. Both proteins share Pfam domain PF00179. In protein I3L0A0 domain PF10520 has been assigned the Gene Ontology (GO) extracellular GO term (GO:0005615). In protein G7Y5H3 the substituted domains, PF014699 and PF14701, are both labeled with the intracellular GO term (GO:0005622) [Color figure can be viewed at wileyonlinelibrary.com]

FIGURE 9 Frequency of the number of Gene Ontology terms assigned to domains of unknown functions
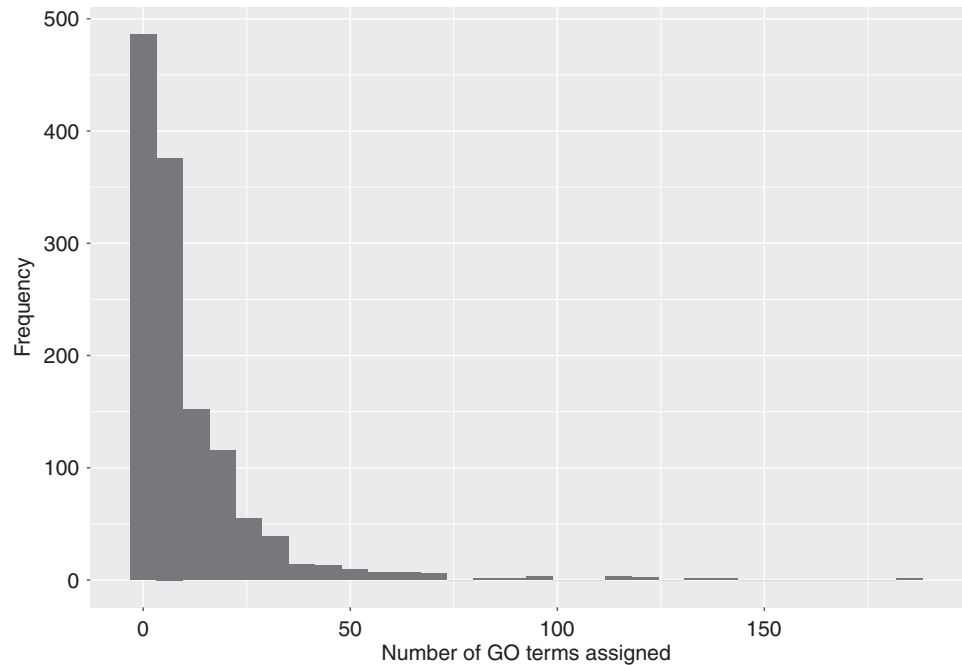
The BP, MF, and CC ontologies have maximum depths of 16, 16, and 11, respectively. This indicates that the typical term assignments are somewhat general, closer to generic terms such as "protein binding" rather than terms which indicate explicit functional roles, such as catalytic mechanisms. In Figure 9, the distribution of terms indicates that the majority of DUFs receive only a handful of putative GO assignments. We suggest that such assignments could be used as general starting points for Pfam domain annotations and with relatively fewer terms to confirm in most these should not make such annotation tasks more onerous or obfuscated. We make these annotations available (see Appendix S1) and note they could make a starting point for future annotation of these domains in Pfam.

4 | DISCUSSION

Applying Word2vec to protein domains, making the assumption that multi-domain proteins are sentence-like, reveals that domains display some manner of semantic or lexical structure. Given this, it should be possible in future to elucidate statistical or semantic rules for domain placement in multi-domain proteins using grammatical inference methods. This would have applications in protein design and modeling.

The Word2vec algorithm was designed to work over very large corpuses of human language, and while the 9 million eukaryotic Inter-Pro sequences used in this study is a relatively large corpus, the corpus of "sentences" currently has too sparse a level of GO annotation to allow us to develop a high-quality embedding of word-tokens which maps well to GO term defined function. A further limitation lies in the amount of domain coverage. Nearly, 70% of the proteins remain unassigned to domains and without greater domain coverage a truly robust domain embedding may not be possible. Additionally, multi-domain proteins typically have fewer than six domains, and often just two or three, whereas human sentences comprise longer sequences. This may mean sequential sets of domains are unlikely to provide sufficient contextual information to produce an informative vector embedding. All these issues might be addressed by retuning the Word2vec model to make it more appropriate for domain data. Word2vec offers several trainable parameters which may allow the method to be adapted for better performance with protein domains, however, it may be the case that an entirely different architecture will be needed.

Using GO annotations to annotate domains is necessarily noisy. It is not clear that they are the best way to encode the lexical "meaning" of an isolated domain in its multi-domain context. In future, a finer grained annotation of domains' sub-functional roles will be necessary

to correctly interpret the lexical meaning of arithmetic transformations of vectors in the embedding space. Nevertheless, this work does open up the tantalizing possibility that protein domains have contextual lexical meaning that could be learned and in turn could be used to derive rules for multidomain protein evolution. However, even in light of these limitations the vector embedding allows us to suggest preliminary function roles for many, as yet, unannotated Pfam domains, and combined with other sources of functional information, this could help improve our overall ability to assign functions to proteins and the genes which encode them.

4.1 | Code and data

All code is available on GitHub and the domain assignments, genism model, token distance matrix and DUF assignments are available via our webserver:

https://github.com/psipred/domain_word2vec_scripts

https://bioinfadmin.cs.ucl.ac.uk/downloads/word2vec/.

ORCID

Daniel W. A. Buchan https://orcid.org/0000-0001-7391-4696

REFERENCES

1. Mikolov, T., Chen, K., Corrado, G. Dean, J. *Efficient Estimation of Word Representations in Vector Space*. arXiv: 1301.3781v1, 2013.
2. Goldberg, Y. & Levy, O., *word2vec Explained: Deriving Mikolov Et al's Negative-Sampling Word-Embedding Method*. arXiv: abs/1402.3722 2014.
3. Asgari E, Mofrad MR. Continuous distributed representation of biological sequences for deep proteomics and genomics. *PLoS One*. 2015;10(11):e0141287.
4. Yang KK, Wu Z, Bedbrook CN, Arnold FH. Learned protein embeddings for machine learning. *Bioinformatics*. 2018;34(15):2642-2648.
5. Viehweger A, Krautwurst S, Parks DH, König B, Marz M. An encoding of genome content for machine learning. *Biorxiv*. 2019. https://doi.org/10.1101/524280.
6. Finn RD, Coggill P, Eberhardt RY, et al. The Pfam protein families database: towards a more sustainable future. *Nucleic Acids Res*. 2016;44(D1):D279-D285.
7. Andreeva A, Howorth D, Chothia C, Kulesha E, Murzin AG. SCOP2 prototype: a new approach to protein structure mining. *Nucleic Acids Res*. 2014;42(Database issue):D310-D314.
8. Cheng H, Schaeffer RD, Liao Y, et al. ECOD: an evolutionary classification of protein domains. *PLoS Comput Biol*. 2014;10(12):e1003926.
9. Dawson NL, Lewis TE, Das S, et al. CATH: an expanded resource to predict protein function through structure and sequence. *Nucleic Acids Res*. 2017;45(D1):D289-d295.
10. Das S, Oregngo CA. Protein function annotation using protein domain family resources. *Methods*. 2016;93:24-34.
11. Nepomnyachiy S, Ben-Tal N, Kolodny R. Complex evolutionary footprints revealed in an analysis of reused protein segments of diverse lengths. *Proc Natl Acad Sci U S A*. 2017;114(44):11703-11708.
12. Friedberg I. Automated protein function prediction—the genomic challenge. *Brief Bioinform*. 2006;7(3):225-242.
13. Watson JD, Laskowski RA, Thornton JM. Predicting protein function from sequence and structural data. *Curr Opin Struct Biol*. 2005;15(3):275-284.
14. Loewenstein Y, Raimondo D, Redfern OC, et al. Protein function annotation by homology-based inference. *Genome Biol*. 2009;10(2):207.
15. Radivojac P, Clark WT, Oron TR, et al. A large-scale evaluation of computational protein function prediction. *Nat Methods*. 2013;10(3):221-227.
16. Consortium GO. Expansion of the gene ontology knowledgebase and resources. *Nucleic Acids Res*. 2017;45(D1):D331-d338.
17. Cozzetto D, Buchan DWA, Bryson K, Jones DT. Protein function prediction by massive integration of evolutionary analyses and multiple data sources. *BMC Bioinformatics*. 2013;14(Suppl 3):S1.
18. Lan L et al. MS-kNN: protein function prediction by integrating multiple data sources. *BMC Bioinformatics*. 2013;14(Suppl 3):S8.
19. Goldberg T, Hecht M, Hamp T, et al. LocTree3 prediction of localization. *Nucleic Acids Res*. 2014;42(Web Server issue):W350-W355.
20. Khan IK, Wei Q, Chapman S, KC DB, Kihara D. The PFP and ESG protein function prediction methods in 2014: effect of database updates and ensemble approaches. *Gigascience*. 2015;4:43.
21. Almeida-e-Silva DC, Vencio RZ. SIFTER-T: a scalable and optimized framework for the SIFTER phylogenomic method of probabilistic protein domain annotation. *Biotechniques*. 2015;58(3):140-142.
22. Van Landeghem S et al. Exploring biomolecular literature with EVEX: connecting genes through events, homology, and indirect associations. *Adv Bioinformatics*. 2012;2012:582765.
23. Falda M et al. Argot2: a large scale function prediction tool relying on semantic similarity of weighted Gene Ontology terms. *BMC Bioinformatics*. 2012;13(Suppl 4):S14.
24. Das S, Lee D, Sillitoe I, Dawson NL, Lees JG, Orengo CA. Functional classification of CATH superfamilies: a domain-based approach for protein function annotation. *Bioinformatics*. 2016;32(18):2889.
25. Fang H, Gough J. A domain-centric solution to functional genomics via dcGO predictor. *BMC Bioinformatics*. 2013;14(Suppl 3):S9.
26. Finn RD, Attwood TK, Babbitt PC, et al. InterPro in 2017-beyond protein family and domain annotations. *Nucleic Acids Res*. 2017;45(D1):D190-d199.
27. Walsh I, Giollo M, di Domenico T, Ferrari C, Zimmermann O, Tosatto SCE. Comprehensive large-scale assessment of intrinsic protein disorder. *Bioinformatics*. 2015;31(2):201-208.
28. Jones DT. Protein secondary structure prediction based on position-specific scoring matrices. *J Mol Biol*. 1999;292(2):195-202.