

# Learning Representations for Gene Ontology Terms by Contextualized Text Encoders

Dat Duong, Ankith Uppunda, Chelsea Ju, James Zhang, Muhao Chen, Eleazar Eskin, Jingyi Jessica Li, Kai-Wei Chang.

University of California Los Angeles, California, USA.

[datdb@cs.ucla.edu](mailto:datdb@cs.ucla.edu), [eeskin@cs.ucla.edu](mailto:eeskin@cs.ucla.edu), [jli@stat.ucla.edu](mailto:jli@stat.ucla.edu), [kwchang@cs.ucla.edu](mailto:kwchang@cs.ucla.edu)

## Abstract

Functions of proteins are annotated by Gene Ontology (GO) terms. As the amount of new sequences being collected is rising at a faster pace than the number of sequences being annotated with GO terms, there have been efforts to develop better annotation techniques. When annotating protein sequences with GO terms, one key auxiliary resource is the GO data itself. GO terms have definitions consisted of a few sentences describing some biological events, and are also arranged in a tree structure with specific terms being child nodes of generic terms. The definitions and positions of the GO terms on the GO tree can be used to construct vector representations for the GO terms. These GO vectors can then be integrated into existing prediction models to improve the classification accuracy. In this paper, we adapt the Bidirectional Encoder Representations from Transformers (BERT) to encode GO definitions into vectors. We evaluate BERT against the previous GO encoders in three tasks: (1) measuring similarity between GO terms (2) asserting relationship for orthologs and interacting proteins based on their GO annotations and (3) predicting GO terms for protein sequences. For task 3, we show that using GO vectors as additional prediction features increases the accuracy, primarily for GO terms with low occurrences in the manually annotated dataset. In all three tasks, BERT often outperforms the previous GO encoders.

## 1 Introduction

The Gene Ontology (GO) provides descriptions for functions of genes and proteins [8]. This database <sup>1</sup> contains terms referred to as GO terms, each term has a definition describing some biological events. To clearly annotate the locations and functions of proteins, this database is further divided into three smaller ontologies: cellular components (CC), molecular functions (MF) and biological processes (BP). In each smaller ontology, the GO terms are arranged into a directed tree with one single root (GO tree), where terms describing more specific biology functions are child nodes of more generic terms.

In late 2017, [29] reported that only about 1% of the proteins in the GO database have manually verified GO annotations. With the advancement of sequencing technology, this fraction is expected to drop in the incoming years; hence, there have been great efforts in developing methods to automatically assign GO terms for unknown sequences [10, 17, 24, 29, 30]. The manually annotated data, which is often used as training and evaluation sets, have many GO terms annotating only a few proteins; for example, [24] estimate that about half of the GO terms annotate about 10 proteins in Human and Yeast

---

<sup>1</sup><https://www.uniprot.org/>

database. To increase the prediction accuracy, automatic annotation methods have been using two additional data resources.

The first data resource is the sequence-to-sequence relationship. For example, protein-protein interaction network and structural homology are often used to constrain the fact that closely related proteins should have similar GO labels [4, 29]. The second data resource is in fact the Gene Ontology itself. For example, distance metric for GO terms can be inferred from the GO tree or their definitions, and then be used as the intuitive constraint that forces similar terms to have equivalent prediction probabilities for a given protein sequence [24]. More importantly, for rare label prediction problems, works in other research domains have shown that using vector representations of labels as one of the features can boost the classification accuracy [2, 19, 27].

This paper will focus on the second data resource, the Gene Ontology itself. On this end, there have been efforts in developing distance metric for GO terms [7, 12, 14, 16, 18, 21, 28]. Most traditional methods for computing semantic similarity of GO terms rely on the Information Content (IC) and the GO tree. For two GO terms, the key idea is to first retrieve the shared common ancestors and then weigh these nodes with their IC values. For example, the most basic method [18] takes the maximum IC of the shared ancestors as the similarity score for two given GO terms. Methods based on shared ancestors and IC values have two drawbacks. First, they do not consider the definitions of the GO terms which have been shown to yield better semantic similarity scores in many cases [7, 14]. Second, they are unable to create vector representations of GO terms which then can be integrated into other annotation models to predict functions for protein sequences.

In recent years, with the advancement of computing power, neural network (NN) encoders have been introduced to map GO terms into vectors based on the principle that the vectors of related GO terms should have similar values [7, 21]. Once the GO vectors are created, then their distance metric naturally follows; for example, cosine similarity or Euclidean distance can be applied. Thus, NN encoders solve the same problem as IC models, and also provide GO vectors which later can be integrated into known annotation methods. Typically, NN encoders are divided into two classes; they either transform GO definitions or the GO entities (e.g. GO names) into vectors. For example, consider two recent methods [7] and [21]. [7] apply Long-short Term Memory on the GO definitions; whereas, *Onto2vec* in [21] apply *Word2vec* on axioms, for example "GO:0060611 is\_subclass GO:0060612", to capture relatedness of GO entities by using the vectors representing their GO names.

In principle, both types of encoders solve the same problem; however, one key question is: in practice, which type of encoder tends to be better? Unfortunately, there have not been any extensive works comparing these encoders. Moreover, despite these methods providing the GO vectors, there have not been works assessing how do these vectors affect the prediction of GO labels for protein sequences; for example, do these GO vectors indeed increase the function annotation accuracy? Also, the recent GO definition encoder in [7] uses LSTM as its key component. In the past year, a new NN method Bidirectional Encoder Representations from Transformers (BERT) has achieved the state-of-the-art results when compared to LSTM-based models in several language tasks such as textual entailment, name entity recognition, sentimental analysis, and language translation [6]. It is therefore important to determine how well can BERT encode GO terms. This paper addresses the

key points mentioned here.

The original BERT implementation produces an embedding matrix for the words in the input (usually a few sentences), but does not provide a vector representation for each of the sentences in the input. In this paper, we introduce 4 adaptations based on BERT to transform GO terms into vectors. Given a sentence (e.g. GO definition) BERT provides a matrix embedding where the column  $j$  corresponds to the  $j^{th}$  word. Instead of using LSTM or a Convolution Layer with fixed window size, BERT uses a 12-layer attention mechanism [6, 23]. Loosely speaking, for one input sentence having  $L$  words, at the layer  $i$  and the word  $j$ , the word vector  $w_{ij}$  is a weighted average of the vectors  $w_{i-1,k}$  for all  $k \in [1, L]$ . We test 4 adaptations of BERT to transform GO terms into vectors. In BERT1, conditioned on two given GO terms, we train BERT to (1) predict missing words in the two definitions and (2) test if the second definition follows the first one (e.g. if the two GOs are child-parent terms). To extract the vector for one GO definition, we average the token embedding in layer 11 of BERT (to be explained later). BERT2 and BERT3 build upon BERT1; here we continue training BERT1 but with a new objective function. In BERT2, we average the word embedding in layer 12 of BERT to represent one GO definition. In BERT3, we use the header token of the GO definition to represent the entire definition. Given two GO terms, we train BERT2 and BERT3 to minimize the cosine loss of the two vectors representing these GO term. In BERT4, we reuse BERT1 framework, and simply replace the GO definitions with the GO names so that we convert GO names into vectors; this idea is similar to Onto2vec [21].

Second, we evaluate the 4 BERT encoders, the bidirectional Long Short Term Memory (biLSTM) in [7], Graph Convolution Network in [9], and Onto2vec in [21] in three downstream tasks: (1) measuring semantic similarity between GO terms, (2) asserting relationship for orthologs and interacting proteins based on their GO annotations and (3) predicting GO terms for protein sequences. For tasks 1 and 2, we include two IC methods: Resnik and Aggregated Information Content (AIC) [18, 22]. IC methods do not return vector representations for GO terms, and so are not used in task 3. In task 1 and 2, GO encoders can outperform IC methods only when the data are well annotated with GO terms having high IC values. In task 3, using GO vectors as prediction features increases the annotation accuracy, primarily for rare GO terms. For all three tasks, GO definition encoders are often better than entity encoders, where BERT2 usually outranks the other encoders. Our code and data are at <https://github.com/datduong/EncodeGeneOntology>.

## 2 Encoders

In this paper, we use the word *encoders* to refer to NN methods that transform GO terms into vectors. Typically, there are two types of encoders. *Sentence encoders* convert the definitions of GO terms into vectors; by default, terms describing related biology events will have similar vectors. *Entity encoders* treat a GO term as a single entity and encode it into a vector based on its position in the GO tree without using its definition. Here, terms within the same neighborhood in the GO tree will have similar vector values. We will first describe the sentence encoders, and then the entities encoders.

## 2.1 Sentence encoders

### 2.1.1 BiLSTM

We first describe the Bidirectional Long Short Term Memory (BiLSTM) model to encode sentences into vectors. BiLSTM provides contextualized vector for each word in a sentence, so that the same word will have different vectors depending on its position in the sentence. We begin with the input of BiLSTM which is very often the Word2vec encoder. Word2vec assigns similar vectors to words with related meanings or are likely to occur in the same sentence [13, 20]. We train our own Word2vec using open accessed papers on Pubmed following the setting in [7] where the word dimension is 300. For this process, we keep stop-words (e.g. but, and not) and symbols like + and – because they may have important biological meanings.

Given one sentence, when using Word2vec, we would convert the sentence into a matrix  $M$  where each column  $M_i$  is vector for the word at position  $i$  in the sentence. Regardless of the sentence, the same word is always assigned to the same vector. To capture the fact that the same word can have different meanings depending on its position in the sentence, we apply  $\tilde{M} = \text{biLSTM}(M)$  where the same word in  $\tilde{M}$  will have different vectors. For example, consider the word vector  $M_i$  at position  $i$  in a sentence of length  $L$ . BiLSTM computes the forward and backward LSTM model to produce the output vectors  $\vec{h}_i = \text{LSTM}(\vec{h}_{i-1}, M_i)$  and  $\overleftarrow{h}_i = \text{LSTM}(\overleftarrow{h}_{i+1}, M_i)$  and then returns  $\tilde{M}_i = [\vec{h}_i; \overleftarrow{h}_i]$  where  $[\vec{h}_i; \overleftarrow{h}_i]$  indicates the concatenation of the two vectors into one single vector.

To encode a matrix of words into a vector of a sentence, we take max-pooling across the columns of  $\tilde{M}$ ,  $\text{maxpool}(\tilde{M})$  [5]. Next, we apply one linear transformation to this *aggregated* vector to produce a final representation of the GO definition. We set the biLSTM hidden layer at 1024, and the final linear layer at 768. During training, we freeze the input  $M$  and update only BiLSTM parameters.

### 2.1.2 BERT

We provide a high-level explanation for BERT. Like BiLSTM, BERT converts words in an input sequence (which can be more than one sentence) into a contextualized embedding where the same word has different vector representations depending on its position in the sequence. Unlike BiLSTM, BERT's key internal structure is the Transformer framework which relies on attention mechanism and will be described below.

We will use this BERT architecture to capture the relationship of the GO definitions. Consider the example in Figure 1, where the input sequence is the child-parent description *perforation plate* (GO:0005618) and *cellular anatomical entity* (GO:0110165). We are using the short descriptions in this example, but in the experiment we will use the complete descriptions. To capture the relationship that *perforation plate* is a *cellular anatomical entity*, we input both sentences into BERT (Fig. 1).

In the first step, BERT splits each word into smaller segments called tokens; for example the word *perforation* is segmented into 3 tokens *per ##fo ##ration*. We use the same segmentation rule as in the original paper [6]. The symbol ## is only a naming convention and bares no significant meaning. For our example, BERT processes the GO terms into

the format *[CLS] per ##fo ##ration plate [SEP] cellular an ##ato ##mic ##al entity [SEP]*, where the special token [CLS] denotes the start of the whole input and [SEP] denotes end of each sentence [6]. BERT internal structure is the Transformer encoder which is described in detail in [23]. Here, we briefly describe the key idea in Transformer. Transformer has several independent heads, each using its own attention mechanism. Loosely speaking, for each head  $h$  in the layer  $i$ , the output vector  $o_{i,j}^h$  for the word at position  $j$  is computed as a weighted average

$$o_{i,j}^h = \sum_{k \in \{1:L\}} a_{ik}^h V^h(w_{ik}) \quad (1)$$

$$a_{ik}^h = \text{softmax}(e_{ik}^h) \quad (2)$$

$$e_{ik}^h = Q^h(w_{ij})^\top K^h(w_{ik}) \quad (3)$$

where  $L$  is input length, and  $V^h, Q^h, K^h$  are transformation functions. To merge all the heads at the layer  $i$ . Transformer concatenates the output  $o_{i,j}^h$  at the position  $j$  of all the heads, and then applies a linear transformation on this concatenated vector. The output of this linear transformation at position  $j$  is then passed onto the next layer  $i + 1$ .

The input of the first layer denoted as  $w_{0j}$  is a summation of the token, position and type embedding. Token embedding is analogous to the Word2vec embedding where every token is assigned to exactly one vector. Position embedding assigns a vector to each location in the sentence; in our example, we would add the position vector  $p_1$  to the token at position 1 which is *per*. Type embedding assigns the same vector  $v_1$  to tokens in the first sentence, and the same vector  $v_2$  to tokens in the second sentence. In our example, we would add the same vector  $v_1$  to each token vector in the first sentence which are *[CLS] per ##fo ##ration plate [SEP]*.

We use the same hyper-parameters as the original BERT in [6], where the Transformer encoder has 12 heads, 12 layers, and the linear transformation matrix produces a vector size 768. The final results are 12 layers of embedding, each with size  $768 \times L$ . Based on the framework of Transformer, the final output vector of the token [CLS] is a function of all the other words in both GO definitions, and therefore can be viewed as an aggregated representation of both GO definitions. For this reason, the embedding of [CLS] in layer 12 is passed through a full connected layer to predict if *perforation plate* is a *cellular anatomical entity*. To ensure that BERT returns high probability for this example case, we will need to fine-tune the original BERT with respect to the context of the Gene Ontology.

We use the Pytorch BERT code <sup>2</sup> and fine-tune a BERT pretrained on Pubmed [11]. BERT is tuned with two tasks: masked language model and next-sentence prediction. Masked language task randomly removes words in a sentence, and then uses the remaining words to predict the missing words. Next-sentence task estimates if two sentences are sequential or chosen randomly from the corpus. Next-sentence prediction uses the [CLS] embedding in layer 12 to make the final decision. In our example, the next-sentence task should confirm that the two sentences are sequential. To fine-tune, we create our own data with respect to the context of the Gene Ontology. To create one *document*, we concatenate the definitions of all GO terms in one single branch of the GO tree, starting from the leaf

<sup>2</sup><https://github.com/huggingface/pytorch-pretrained-BERT>



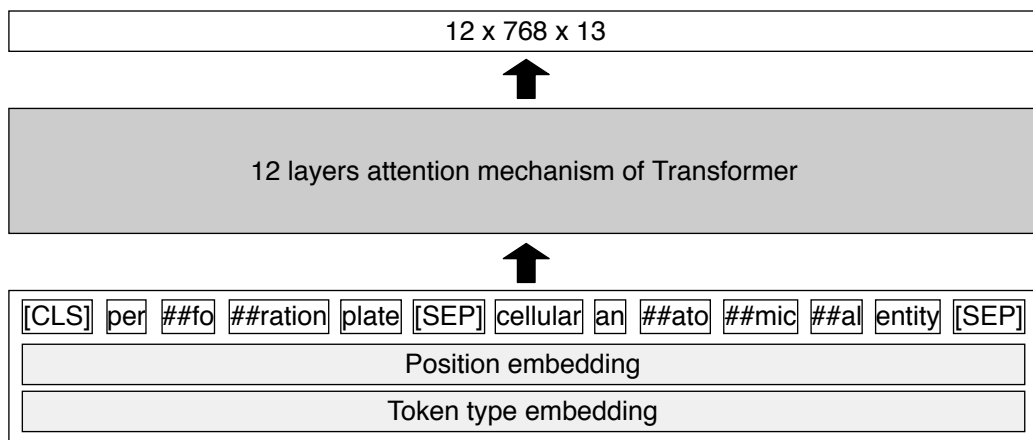


Figure 1: Consider child-parent terms GO:0005618 (perforation plate) and GO:0110165 (cellular anatomical entity). We input into BERT the tokenized words of *perforation plate* and *cellular anatomical entity*. This illustration uses the short descriptions, but in the experiment we will use the complete descriptions. The [CLS] specifies the start of the whole input, and [SEP] specifies the end of each sentence. At the lowest level, BERT has three embedding layers: word, position, and token type. For example, [SEP] appears twice but will have different position vector representations. Each token in the first segment [CLS] *per ##fo ##ration plate* [SEP] will have token type 1 embedding, whereas the rest of the input will have token type 2 embedding. In general, the type 1 embedding is assigned to words in the first sentence, and type 2 embedding is meant for words in the second sentence. For each token, BERT adds the embedding of the three layers and then sends this summation into the Transformer encoder. BERT outputs 12 layers of embedding size  $12 \times 768 \times 13$  for this example, where 13 is the total length of both GO terms including the [CLS] and [SEP] tokens.

node to the root. We consider only is-a relation, and randomly select only one parent if the given node has many parents. Our fine-tuned BERT will capture the relationships of words within a sentence, and also the relationships of GO definitions that are on the same path to the root nodes.

We emphasize that by default, BERT does not provide a vector representation for a given GO definition. BERT only provides the matrix embedding for the words in a GO definition. After tuning BERT, we test two methods to retrieve the vector representations for the GO definitions from the word embedding matrix.

For our first method (BERT1), we follow Bert-as-service [26]. To transform the GO description *perforation plate* into a vector, we input it into BERT as [CLS] *per ##fo ##ration plate* [SEP] without any second sentence which would be *cellular anatomical entity* in the example in Fig. 1. Then we average the vectors of all the tokens, [CLS] and [SEP] in layer 11. [26] recommends this layer because layer 12 may be too affected by the masked language model and next-sentence prediction task instead of our key objective which is to extract the sentence representation. We do not further train the fine-tuned BERT. Because we use the same hyper-parameters as the original BERT, by default, BERT1 returns a GO vector of length 768.

For the second method, we continue training our fine-tuned BERT but apply a different

approach. We remove masked language model, and use only the next-sentence prediction. For two GO terms, we first use BERT to produce the vectors for their definitions, and then apply cosine loss on these two vectors. We test two different ways to represent the GO definitions (BERT2 and BERT3). BERT2 follows similar idea in BERT1. For a GO term, we send only its definition through BERT (not appending definitions of parent terms), and then average the token embedding in the layer 12. We add one extra linear layer to transform this *mean* vector to retrieve the final representation of the GO definition.

In BERT3, for each GO term, we again send only its definition through BERT. Next, we use the *pooled output* of layer 12 which is simply the [CLS] token in layer 12 transformed by a linear layer with Tanh activation. We pass this *pooled output* through one more linear layer to produce the final vector representation of the GO definition. In BERT2 and BERT3, the final linear layer returns an output of size 768 to match BERT1 output which is 768

For BERT2 and BERT3, given two GO terms, we independently transform each of their definitions into a vector (by individually sending each definition through BERT, and not by concatenating them first as one long sentence). Then for training, our objective is to maximize or minimize the cosine distance of these two vectors depending on whether the terms are child-parent or randomly chosen.

## 2.2 Entity encoders

Because GO terms are arranged in a directed tree, we can treat a GO term as a single entity and encode it into a vector without using its definition. In this paper, we test Graph Convolution Network (GCN) and Onto2vec. There are other node embedding methods, but GCN has shown to work well in practice for prediction tasks when labels have low occurrence frequencies [9, 19, 27].

### 2.2.1 GCN

Graph Convolution Network encodes each GO term in the tree into a vector [9]. Let  $A$  be the adjacency matrix, where  $A_{ij} = 1$  if  $GO_i$  is the parent of  $GO_j$ . Compute  $\hat{A} = A + I$ , where  $I$  is identity matrix. Compute the degree matrix  $\hat{D}$ , where  $\hat{D}_{ii} = \sum_j \hat{A}_{ij}$ . Next scale  $A$  into  $\hat{A} = \hat{D}^{-\frac{1}{2}} \hat{A} \hat{D}^{-\frac{1}{2}}$ . Let  $W_1$  and  $W_2$  be two transformation matrices. Define  $X$  to be the initial vector embedding for the GO terms, where a column in  $X$  corresponds to a GO vector. Before training,  $X$  is initialized with random numbers. During training  $X$  is transformed into a new matrix  $E = W_2 \hat{A} \text{relu}(W_1 \hat{A} X)$ . Loosely speaking, one column  $i$  in  $\hat{A} X$  is some summation of all its neighbor nodes and itself.  $W_1 \hat{A} X$  then transforms this summation into a new vector. We repeat this transformation twice as in [9, 19]. At the end, column  $i$  in  $E$  is the vector for  $GO_i$  and is a function of its child nodes. We train GCN to minimize the cosine distance loss of the column vectors in  $E$  using the data in section 2.3. We set GCN to produce the final vector representation of size 768, same as Bert-as-service.

## 2.2.2 Onto2vec

Onto2vec encodes GO terms into vectors by transforming their relationships on the GO tree into sentences, which are referred to as axioms in the original paper [21]. For example, the child-parent GO terms GO:0060611 and GO:0060612 are rewritten into the following sentence "GO:0060611 is\_subclass GO:0060612". Onto2vec then applies word2vec on these sentences, so that GO names occurring in the same sentence are encoded into similar vectors. Because the training *sentences* are constructed from the GO trees without GO definitions, Onto2vec can conceptually be viewed as method that encodes nodes on graph into vectors like GCN. Because word2vec objective function is based on cosine similarity, for Onto2vec, GO terms in close proximity will have high cosine similarity score. We set Onto2vec to produce the final vector representation of size 768, same as Bert-as-service.

## 2.2.3 BERT as entity encoder

Following Onto2vec, we apply BERT as an entity encoder (our BERT4) where the key objective is to encode the GO names into vectors. We create training data as follows. For each GO term, we select one path from that term to the root node via only *is\_a* relation. For each path, we split the set of GO terms into half so that they represent the *first* and *second* sentence. BERT1 and BERT4 have similar idea. In BERT1, the training step requires GO definitions, whereas this phase in BERT4 uses only the GO names. For example, consider the path GO:0000023, GO:0005984 GO:0044262, GO:0044237, and GO:0008152. In BERT4, we format it into the input `[CLS] GO:0000023 GO:0005984 [SEP] GO:0044262 GO:0044237 GO:0008152 [SEP]`. Next, we set the words in the BERT vocabulary as the GO names. Then, we train masked language model and next-sentence prediction on this data, so that we can capture relatedness among the GO names like Onto2vec. We use the same hyper-parameters as bioBERT, where the final token embedding is size 768. The final vector representation for GO terms is the BERT's initial token embedding. We do not take the last layer output because we do not want the contextualized vectors of the GO names.

## 2.3 Training data

We train BERT1 and BERT4 using the data and fine-tune procedure described in section 2.1.2 and 2.2.3. We train the other encoders using the data described here. Our objective is to first use the encoders to transform GO terms into vectors, and then to maximize and minimize cosine distance for child-parent and unrelated GO pairs sampled from the GO database. For our training data, we treat the BP, MF and CC ontology one connected network; this approach has shown to increase accuracy for downstream tasks [7]. We randomly pair a GO with one of its parents, treating the follow one-directional relationship "is a", "part of", "regulates", "negatively regulates", and "positively regulates" as the same edge. To ensure that these child-parent terms are very similar, we compute their Aggregated Information Content (AIC) scores and retain pairs with scores above the median [22].

To create the negative dataset where each sample is a pair of two randomly chosen GO terms, we sample two types of unrelated pairs. For the first type, we randomly pick about



half the GOs seen in the positive dataset. We pair each term  $c$  in this set with a randomly chosen term  $d$ . For the second type, we pair the same term  $d$  with another randomly chosen term  $e$ . This strategy helps the encoders by letting the same GO terms to be seen under different scenarios. Next, to ensure that these random pairs are very dissimilar, we retain pairs with AIC scores below the median. This training data is available at <https://github.com/datduong/EncodeGeneOntology>.

### 3 Evaluation

We evaluate the GO encoders in three tasks. First, we measure the semantic similarity for two types of GO pairs: child-parent and unrelated terms. The objective is to determine which encoders can best distinguish the two kinds of GO pairs. Here, we also observe how the number of neighbors (degrees) and ICs of the GO terms affect their similarity scores. Second, we assert the relationship for orthologs and interacting proteins based on their GO annotations. Here, we consider only manual annotations data. This task is similar the first task; that is, if an encoder does well in task 1 then it is likely to do well in task 2. However, task 2 provides a more holistic picture because in practice, genes and proteins are not often manually annotated by uninformative GO terms which have high degrees and low IC values. Hence, a method can possibly perform well in task 2 even if it does not do well in task 1. Third, we edit the DeepGO model so that it takes the GO vectors as an extra input. We test if the GO vectors indeed boost the accuracy for predicting GO labels of protein sequences. Moreover, suppose the GO vectors improve the accuracy, then we want to know if the increment occurs for rare GO labels. In some essence, being able to well predict rare GO labels is important because these terms are often located lower in the GO tree, describe more detail biology events, and are closer to the true properties of the proteins. For example, predicting GO:0005618 *perforation plate* is more precise to a protein's location than predicting its parent term GO:0110165 *cellular anatomical entity* or its ancestor GO:0005575 *cellular component*.

#### 3.1 Semantic similarity task

Theoretically speaking, a good GO encoder will clearly separate child-parent GO terms from unrelated pairs regardless of the degrees and ICs for these GO terms. We shall see in practice, such proposition does not hold true; however, the GO encoders which align most closely with the theoretical expectation will be considered best. In general, a GO term's Information Content is negatively correlated with its number of neighbors (or degree) in the GO tree. We estimate this correlation to be  $-0.445$  for 20,283 Human terms. For this experiment, we observe how the degrees and ICs of the terms affect the similarity scores for child-parent and unrelated GO pairs, by seeing how well the inter-quantile ranges (IQRs) of the scores for these two groups stay separated at different degree and IC values. We randomly select Human GO pairs  $A, B$  with  $\max(\deg_A, \deg_B) \leq 100$  and  $\min(\text{IC}_A, \text{IC}_B) \geq 1$ ; the final set contains 3069 child-parent pairs and 3069 unrelated pairs. For each GO pair  $A, B$ , we compare its  $\max(\deg_A, \deg_B)$  and  $\min(\text{IC}_A, \text{IC}_B)$  against its similarity score. We include AIC method which does not encode GO terms into vectors;

it is however informative to compare AIC against the GO encoders.

In Figure 2, the performances of all the methods are inversely correlated to the degrees of GO terms. When  $\max(\text{deg}_A, \text{deg}_B)$  is small, that is when terms are near the leaves or have few neighbors, then all methods except for BERT4 perform well, where the IQRs for child-parent and unrelated GOs do not intersect. For AIC, as  $\max(\text{deg}_A, \text{deg}_B)$  increases, the two IQRs remain well separated, despite their trend lines are getting closer. For neural network encoders, as  $\max(\text{deg}_A, \text{deg}_B)$  increases, the scores of child-parent pairs decrease, so that the IQRs overlap, making it harder to distinguish related GOs from unrelated ones. BERT4 is the only exception, where the two trend lines diverge. However, BERT4 has its IQRs for the two labels intersected at almost all degree values, making it the least desirable metric. Onto2vec and GCN have their IQRs first intersect at  $\max(\text{deg}_A, \text{deg}_B) > 2.5$  and  $\max(\text{deg}_A, \text{deg}_B) > 12.5$ , respectively; whereas, BiLSTM and BERT3 have their IQRs first intersect at  $\max(\text{deg}_A, \text{deg}_B) > 17.5$ . In this sense, BiLSTM and BERT3 are better than Onto2vec and GCN because they can adequately classify GO pairs containing terms with more neighbors. BERT1 and BERT2 perform best and second best, respectively; for both, the two IQRs first intersect at  $\max(\text{deg}_A, \text{deg}_B) > 32.5$  and  $\max(\text{deg}_A, \text{deg}_B) > 27.5$ .

In Figure 2, the performances of all the metrics are positively associated to the IC values of GO terms. As the  $\min(\text{IC}_A, \text{IC}_B)$  increases, that is when the GO terms annotate very few proteins, then the IQRs of scores for child-parent and unrelated pairs do not intersect, so that the methods can better identify the two labels. BERT4 and Onto2vec underperform; despite their two trend lines diverging as IC increases, the two IQRs overlap even for large IC values. For GCN, BiLSTM, BERT1, BERT2 and BERT3, the two IQRs entirely separate when  $\min(\text{IC}_A, \text{IC}_B)$  is over 6.25, 4.25, 4.75, 4.25, and 5.75, respectively. Here, BiLSTM and BERT2 are the two best metrics because they can adequately classify GO pairs containing more generic terms which are annotating more proteins.

Figure 2 indicates three points. First, encoding a GO term via its definition appears to be better than encoding a GO term based on its position on the GO tree. Second, within BERT, BERT1 and BERT2 are better than BERT3 and BERT4. Third, neural network encoders would perform well only for specialized GO terms with low degrees and/or high ICs. To achieve the best result for all GOs, we must integrate the newer methods with IC based models as noted in earlier works [7, 28].

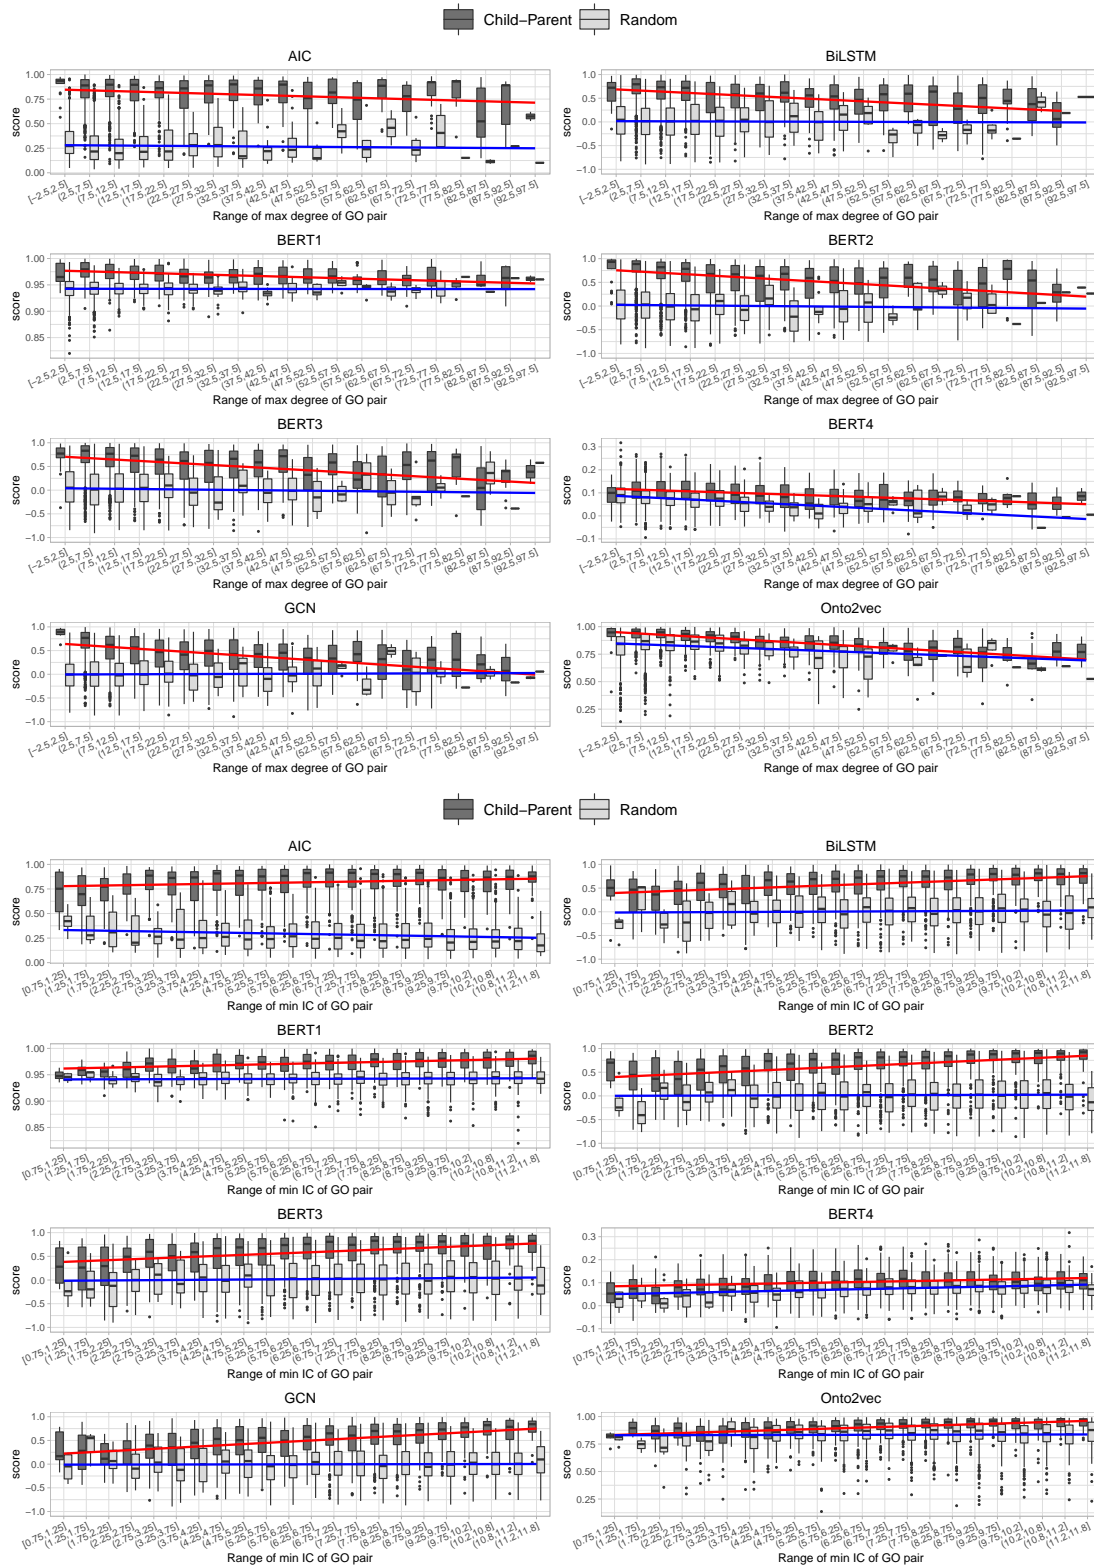


Figure 2: Encoder's ability to classify GO pairs is inversely correlated to the degrees of GO terms and is positively correlated to the ICs of GO terms.

## 3.2 Set comparison task

Because genes and proteins are annotated by GO terms, good GO metrics should well differentiate similar genes and proteins from unrelated ones. To compare the GO encoders, we conduct two experiments (1) classifying orthologs in Human, Mouse, and Fly and (2) classifying true protein-protein interactions in Human and Yeast.

### 3.2.1 Orthologs

We download the orthologs datasets and GO annotations in [7]. This data retains orthologs annotated by at least one GO term in each GO category and removed GOs with tag IEA, NAS, NA, and NR [15]. We test the following species: Human-Mouse (HM), Human-Fly (HF), and Mouse-Fly (MF). For each dataset, the positive set contains orthologs from the two species; whereas, the negative set contains randomly-matched genes. We set the sizes of the positive set and negative set to be equal. The HM, HF, and MF data has 10235, 4880, and 5091 pairs for each set, respectively. Here, comparing two genes is equivalent to comparing their two sets of GO annotations [15]. We use the best max average distance for the GOs in the two annotation sets [7, 15]. For this experiment, we use the entire GO annotations and compare terms across different ontology as in [7, 21].

Table 1 shows the summary statistics for the annotation in each ortholog dataset. Table 2 shows the Area Under the Curve (AUC) for each method. When compared to Resnik and AIC, performances of all GO encoders drop for orthologs data having less informative GO terms; here, results for GO position encoders decrease the most. This observation agrees with Fig. 2, where GO encoders do not perform well for GO terms with low ICs and/or high degrees. Here, encoding GO definitions often yields better accuracy than encoding GO positions on the GO tree. For the three datasets, BERT1 ranks first twice among the definition encoders; whereas BERT2 ranks first once. For the position encoders, there is not one best method.

### 3.2.2 Protein interaction network

Following [7], we download the Human data in [12] and Yeast data in [14]. These data have 6031 and 3938 positive Human and Yeast protein-protein interactions, respectively. For the negative set, we follow the same sampling procedure in [12]. We randomly assign edges between proteins that do not interact in the real PPI network. The real and random PPI network have the same proteins; we only require that they have different interacting partners. Table 2 shows the AUC for each method. Here, the AUC is computed using the exact process in section 3.2.1. For this experiment, we also include SimDef which uses Term Frequency – Inverse Document Frequency to compare the GO definitions [14]. Among the definition encoders, BERT1 and BERT2 rank best for Yeast and Human data, respectively. For position encoders, GCN is the best method for the two datasets.

Table 1: GO annotation summary statistics in each ortholog dataset. GO frequency counts the number of times the GO terms appear (including duplication); for example, if GO *A* and *B* both annotate protein *C* and *D*, then the total GO frequency is 4. Median degree (Deg) and Information Content (IC) for GO terms (including duplication) are proxies for how well annotated the datasets are.

	Negative set			Positive set		
	GO freq.	Deg	IC	GO freq.	Deg	IC
<b>Dataset 1</b>						
Human	123400	8	6.063	128666	8	6.142
Mouse	102775	8	5.907	147012	7	6.314
<b>Dataset 2</b>						
Human	57678	8	6.221	61708	8	6.158
Fly	28169	10	5.437	41481	8	6.290
<b>Dataset 3</b>						
Mouse	48520	9	6.012	67449	8	6.366
Fly	29026	10	5.404	39558	8	6.248

Table 2: AUC for classifying true ortholog pairs in Human, Mouse and Fly, and interacting proteins in Human and Yeast.

	Orthologs data			PPI data	
	Human-Mouse	Human-Fly	Mouse-Fly	Human	Yeast
<b>Info Content</b>					
Resnik	95.27700	94.19600	89.74800	86.97300	90.67000
AIC	95.79100	93.90600	89.84100	87.88900	87.77500
<b>TF-IDF</b>					
SimDef	NA	NA	NA	87.03400	88.22300
<b>GO definition</b>					
BiLSTM	95.19232	91.49026	80.28224	86.60700	88.23649
BERT1	96.72200	92.94441	79.61700	88.15400	88.99700
BERT2	95.94186	92.49684	81.64635	88.33334	89.95763
BERT3	96.05200	90.80000	78.99000	86.94500	89.26698
<b>GO position</b>					
BERT4	96.27481	85.39845	70.48564	83.92643	82.66531
GCN	94.98700	85.53500	72.99300	85.45300	86.74929
Onto2vec	91.79953	82.83256	74.41180	79.71700	83.98171

### 3.3 Annotation task

#### 3.3.1 DeepGO

For this task, we do not aim to design a completely new model that is better than existing baselines for prediction GO annotations. Rather, our purpose is to determine how much can the GO vectors affect the prediction results. For this purpose, we use the data and existing framework of DeepGO [10]. DeepGO data consists of 3 sets BP, MF, and CC. BP, MF and CC terms in this data (combining train, development, and test set) annotate at least 250, 50, 50 proteins. For this data, the parents of all the GOs annotating one protein are



also added into the label set predicted. In total, the number of GO terms and proteins for each BP, MF and CC training dataset are 932 | 36375, 589 | 25199, and 436 | 35546. During training and testing, we use the whole label set size 932, 589 and 436.

Next, we briefly describe the neural network in DeepGO. Given an amino acid sequence, for example  $p = \text{MARS} \dots$ , DeepGO converts  $p$  into a list of 3-mer as MAR ARS  $\dots$ . Each 3-mer is assigned a vector of dimension 128, so that if  $p$  has length 1002 amino acids, then the matrix representing  $p$  is  $E_p \in \mathbb{R}^{128 \times 1000}$ . A 1D-convolution layer and 1D-maxpooling are then applied to  $E_p$ . Flatten layer is applied to get a vector  $v_p$  representing  $p$ ; loosely speaking, we have  $v_p = \text{flatten}(\text{maxpool}(\text{conv1d}(E_p)))$ . DeepGO includes information from a protein-protein interaction network by concatenating  $c_p = [n_p \ v_p]$ , where  $n_p$  is a vector for protein  $p$  in the interaction network produced in [1]. To predict if  $\text{GO}_i$  is assigned to  $p$ , DeepGO fits a logistic regression layer  $\text{sigmoid}(B_i^T c_p + b_i)$ , where  $B_i$  and  $b_i$  are parameter specific to  $\text{GO}_i$ . The loss function is binary cross entropy. DeepGO can be applied with only the protein sequences and without the additional protein network in [1]; in Table 3, we use the name *DeepGOSeq* to refer to this simple implementation, and *Baseline1* as the version of DeepGO having the protein network data.

To add GO encoders into DeepGO, we make one minor change to avoid significantly altering the original DeepGO model. Let  $g_i$  be the vector of  $\text{GO}_i$ , for example  $g_i = \text{BERT3}(\text{definition of } \text{GO}_i)$ . We concatenate  $\hat{c}_{pi} = [c_p \ g_i]$ , and apply one linear transformation  $\tilde{c}_{pi} = \text{relu } W \hat{c}_{pi}$ .  $\tilde{c}_{pi}$  captures the interaction of the protein and GO vectors. To predict if  $\text{GO}_i$  is assigned to  $p$ , we fit  $\text{sigmoid}(B_i^T [\tilde{c}_{pi} \ c_p] + b_i)$  where  $[\tilde{c}_{pi} \ c_p]$  is the concatenation of the two vectors. For this experiment, we freeze the GO vectors and train only the DeepGO parameters. In this paper, our intention is to determine which GO encoders can work best out-of-the-box for predicting functions of unknown sequences. In future research, we will consider jointly training both GO-to-GO and GO-to-protein relationships.

The transformation  $\tilde{c}_{pi} = \text{relu } W [c_p \ g_i]$  may capture only interactions of values from the protein vector; in other words, the values of  $W$  corresponding to any values in  $g_i$  can be all zeros. Thus, we create one extra baseline (*Baseline2*) for this experiment where we remove  $g_i$  and let  $\tilde{c}_{pi} = \text{relu } W c_p$ . The rest of the layers follows exactly as in the previous paragraph. In *Baseline2*,  $\tilde{c}_{pi}$  represents the interaction of the protein vector from [1] and the encoded amino sequence  $v_p$  without any GO vectors.

We compute three metrics Fmax score, macro and micro-AUC which do not require the prediction probabilities to be rounded at a specific threshold. Fmax and micro-AUC put more weights on frequently occurring GO terms, so that mislabeling infrequent GO terms do not greatly affect the outcome; whereas macro-AUC treats all the GO labels equally so that mislabeling infrequent GO terms can significantly affect the outcome [3, 25]. Table 3 shows that our interaction layer  $\text{relu } W c_p$  alone (*Baseline2*) improves upon original model (*Baseline1*) in [10] for all metrics in the three ontology. More importantly, all the GO encoders increase the evaluation metrics with respect to *Baseline2* in all three ontology.

Because the frequencies of GO terms in training data affect their prediction accuracy [10, 29], we partition the GO terms based on their frequencies to further understand how each GO encoder performs. The 25%, 50%, and 75% quantile frequency for GO terms in the BP, MF, and CC training dataset are 233 | 365 | 860, 49 | 88 | 227, and 59 | 111 | 293 respectively. The number of GO terms in the <25% and >75% quantile groups are 232 | 232, 143 | 147, and

110 | 110 for BP, MF and CC respectively. We compute recall-at-k (R@k) and precision-at-k (P@k) for each group (Figure 3). For discovering unknown functions of protein sequences, having high recall rate is important so that we do not miss any annotations. However, by also observing precision rate, we can determine which GO encoders are the most well-balanced.

For this experiment, we use only Baseline2 because it is the most competitive against GO encoders. When evaluating all the BP terms, every encoder attains higher recall and precision, where the increment stems mostly from the encoders more accurately predicting common terms (Fig. 3). The recall and precision curve for terms with low and medium frequencies are not well separated from the baseline. For MF terms, at the overall level, all the encoders achieve better recall and precision than baselines, where the best improvement happens for rare GO terms. For more common MF terms, our encoders are at least equivalent to baseline. For CC terms, our encoders improve the overall recall but not precision curve, where the best increment happens for rare terms.

In conclusion, when introduced into the DeepGO architecture, our encoders boost the overall recall for the three ontology. However, the overall precision increases only for BP and MF terms. Conditioned on the performance over the complete dataset, BERT2 ranks the best because it is always in the group of encoders having the highest Fmax, macro and micro-AUC, recall, and precision for the 3 ontology.

Table 3: Evaluating how much can GO vectors boost DeepGO result. DeepGOSeq uses only protein sequences to predict GO annotation. Baseline1 adds protein network data from [1] into DeepGOSeq. Baseline2 improves Baseline1 by adding one linear layer to convolve vector from protein network and the representation of the amino acid sequence without any GO vectors.

Method	BP			MF			CC		
	Fmax	AUC		Fmax	AUC		Fmax	AUC	
		Macro	Micro		Macro	Micro		Macro	Micro
<b>Baseline</b>									
DeepGOSeq	34.05	62.67	81.89	38.63	72.42	86.73	57.20	67.01	92.39
+Baseline1	41.80	81.87	89.70	46.77	83.73	90.60	62.70	87.80	96.47
+Baseline2	42.51	82.97	90.27	48.62	87.25	93.14	63.77	89.79	97.10
<b>GO definition</b>									
BiLSTM	43.48	83.49	90.73	50.20	88.02	93.84	65.47	90.48	97.25
BERT1	43.45	83.11	90.55	49.24	87.43	93.52	65.39	90.27	97.27
BERT2	43.92	83.49	90.77	50.30	87.68	93.66	65.83	90.51	97.30
BERT3	43.15	83.41	90.59	49.93	87.29	93.39	65.24	90.57	97.26
<b>GO entity</b>									
BERT4	43.35	83.35	90.26	49.13	87.97	93.07	65.66	90.27	97.24
GCN	43.65	83.46	90.63	49.74	87.78	93.50	65.20	90.39	97.20
Onto2vec	42.91	83.48	90.57	49.35	87.92	93.83	64.82	90.50	97.21

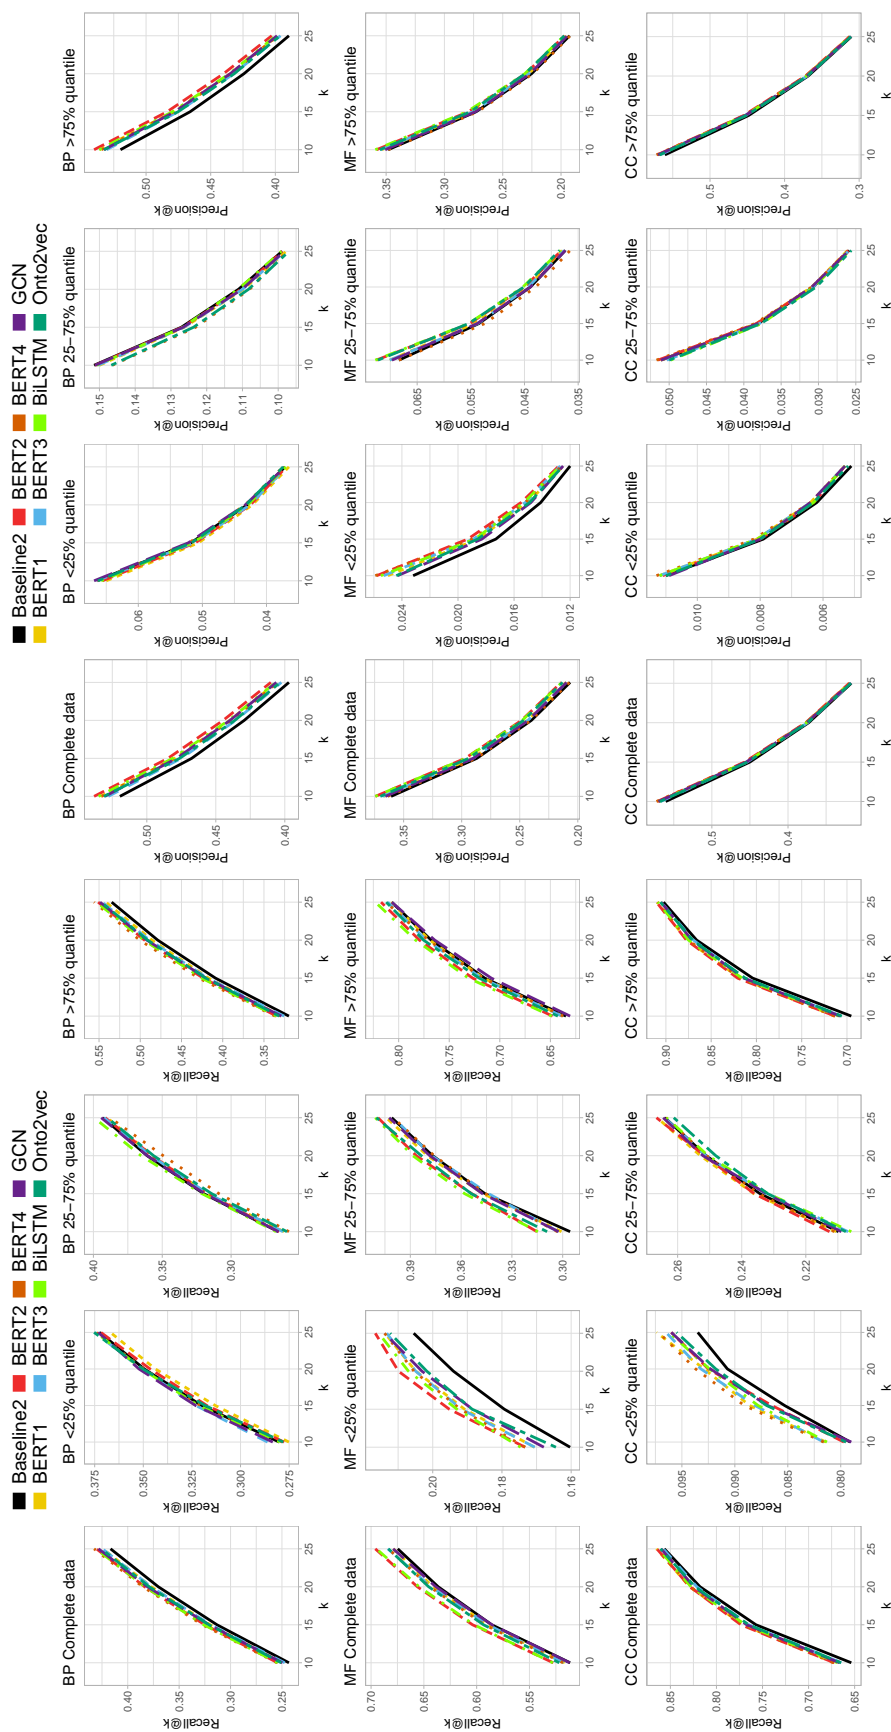


Figure 3: Recall and precision at-k for annotations of protein sequences. For each ontology, we split the GO terms into 3 sets based on their frequencies in the training data.

### 3.3.2 Expand DeepGO dataset

We repeat the experiment above, but extend the number of GO terms to be predicted to have more rare terms. Having a good prediction for rare GO terms is important because these terms are closest to the true functionality of the proteins. Using the same annotation file in [10], we expand the GO sets in their original data; we include BP, MF, and CC terms that have at least 50, 10, and 10 annotations ( $5\times$  less than the original criteria). We ensure that all the GO terms in the original dataset are included into this larger dataset; hence, most of these original terms will have a much larger occurrence frequencies. The BP training dataset now has 2980 terms; the 25%, 50% and 75% quantile occurrence frequency are 62 | 113 | 276. Here, our new BP dataset is harder to predict, because 75% of the GO terms occur less than 276 times (these GO terms barely pass the original cutoff criteria at 250); whereas, the original data has about 75% of terms occurring more than 276 times.

For MF and CC, the new number of GO terms to be predicted are 1677 and 979. The 25%, 50% and 75% quantile frequency are 13 | 22 | 66 and 15 | 34 | 117.5 respectively. Here, the new MF and CC data contain about 75% and 50% of the terms that would barely make the frequency cutoff at 50 in DeepGO original data [10]. CC data becomes harder to predict but not as much as BP and MF data. The number of GO terms in the  $<25\%$  and  $>75\%$  quantile groups are 736 | 742, 417 | 417, and 232 | 245 for BP, MF and CC respectively.

Table 4 and Figure 4 show the prediction outcome. Here, we compare the best BERT method in section 3.3.1 BERT2 against the often used biLSTM, and include GCN which shows significant success in other datasets with many rare labels [19]. The NN methods biLSTM, BERT2 and GCN increase the Fmax but not always macro and micro AUC (Table 4) for the three ontology. All NN methods however boost the recall and precision-at-k for rare GO terms (count frequency less than 25% quantile). For CC terms in medium and high frequency groups, NN methods do not necessary increase recall and precision-at-k; this behavior is also seen for the original DeepGO data (Fig. 3) and other datasets in [19]. More importantly, for our new BP and MF dataset which are harder to predict, biLSTM, BERT2 and GCN attain recall and precision that are better than or equal to Baseline2 for the terms in all frequency groups (Fig. 4). For all the BP terms, biLSTM, BERT2 and GCN appear equivalent. For all the MF terms, BERT2 ranks best. For CC, when considering rare GO terms where NN methods are better than Baseline2, then BERT2 ranks first.

Table 4: Evaluating how much can GO vectors improve Baseline2 for the expanded DeegGO dataset, where we lower the inclusion criteria to have GO terms with occurrence frequency from 250, 50 and 50 to 50, 10, and 10 for BP MF and CC.

Method	BP			MF			CC		
	Fmax	AUC		Fmax	AUC		Fmax	AUC	
		Macro	Micro		Macro	Micro		Macro	Micro
Baseline2	40.03	86.06	93.94	46.19	87.01	95.01	63.87	90.52	98.02
BiLSTM	40.58	85.68	93.85	47.46	85.68	95.01	64.17	90.99	98.12
BERT2	40.34	86.08	93.90	48.17	88.09	95.50	64.02	90.20	98.05
GCN	41.10	86.07	93.86	48.02	86.49	95.31	64.64	90.24	97.99

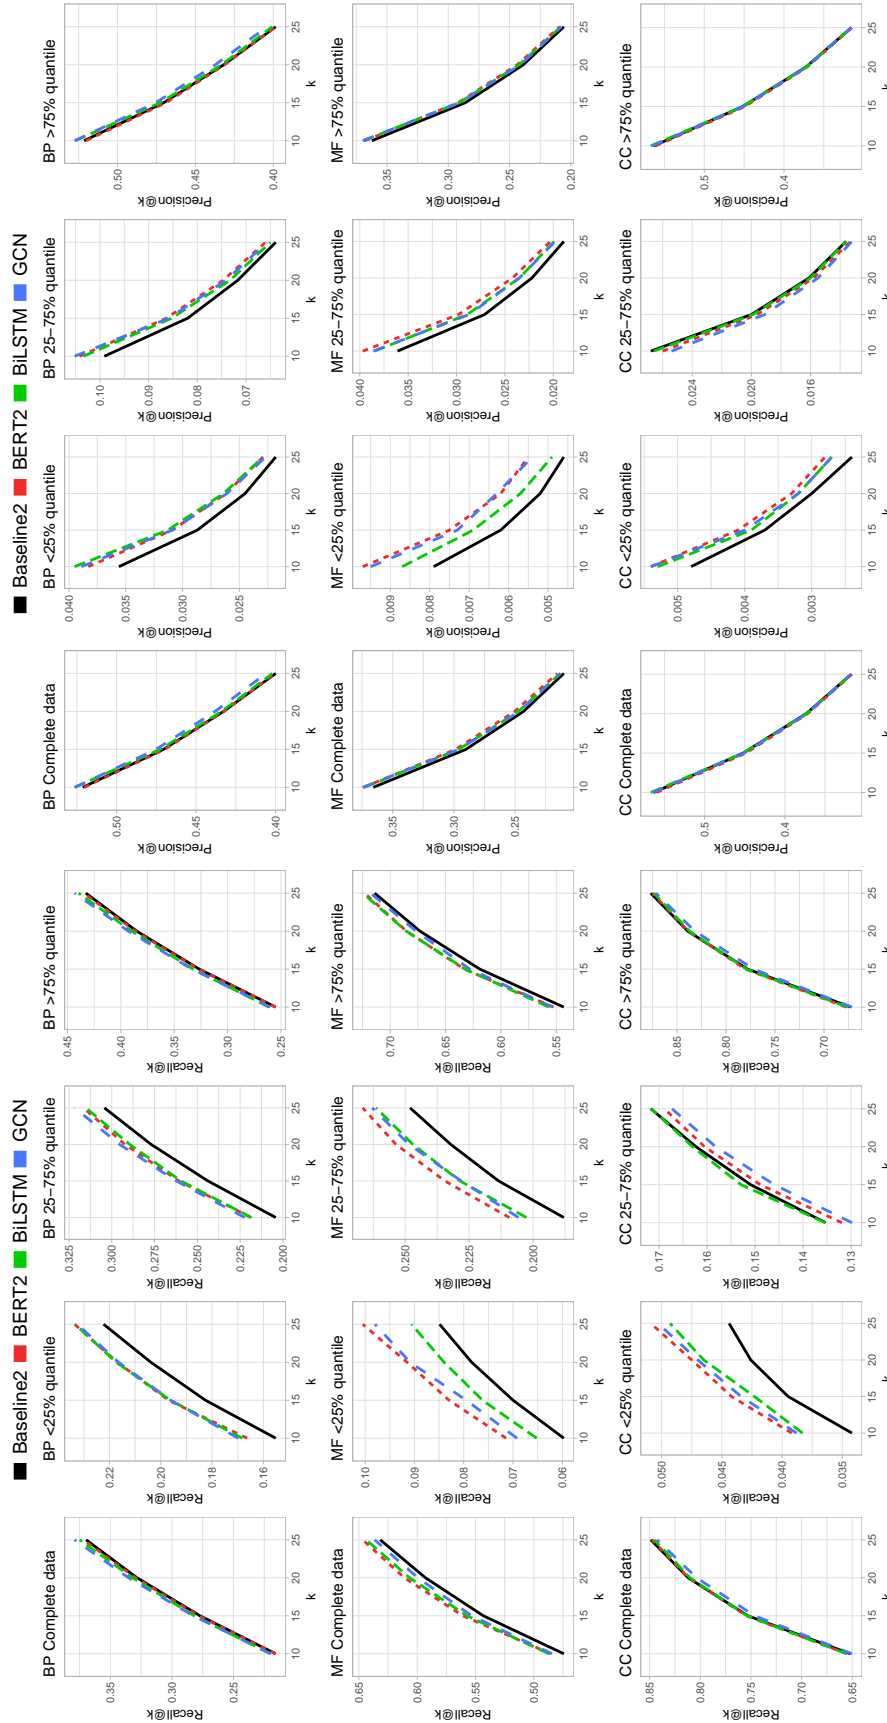


Figure 4: Methods are tested on the expanded DeepGO dataset, where we lower the inclusion criteria to have GO terms with occurrence frequency above 50, 10, and 10 for BP MF and CC. Recall and precision-at-k for annotations of protein sequences. For each ontology, we split the GO terms into 3 sets based on their frequencies in the training data.



## 4 Discussion

In task 1, when comparing the similarity scores of two GO terms against their ICs and degrees, the NN models and AIC can only well differentiate child-parent terms versus unrelated ones if the terms have high ICs and/or low degrees. AIC however outperforms NN models. In task 2, when asserting relationship between genes/proteins based on their GO annotations, NN models can outperform Resnik and AIC for dataset with detail GO annotations, but their accuracy drop compared to Resnik and AIC for less well annotated datasets. The task 1 and 2 results and previous works [7, 10, 14, 28] indicate that NN methods must be integrated with IC-based models to attain the best performance. In future work, we plan to use the ICs and degrees of GO terms as explicit features for the neural network approaches.

In task 3, when integrating GO vectors into DeepGO to predict function annotations of proteins, all encoders help DeepGO attaining higher classification accuracy in the original DeepGO dataset. We further expand original DeepGO dataset to include more rare GO terms where 25% of the BP, MF and CC terms occur about 50, 10 and 10 times. For these rare terms, GO encoders significantly increase the recall-at-k, which is useful for discovering true function annotations. In general, infrequent GO terms benefit the most from having GO vectors; whereas, common terms do not benefit much. However, being able to improve the prediction for rare GO terms is more important because these terms are closest to the true functionality of the proteins.

For all three tasks, GO definition encoders are often better than GO position encoders, with BERT2 always being in the group of best performers. We hope that BERT2 can provide the basis for more complex GO encoding techniques. In future work, we will continue developing BERT2, and integrate it with other models besides DeepGO to better predict function annotations.

## References

- [1] Alshahrani, M., Khan, M.A., Maddouri, O., Kinjo, A.R., Queralt-Rosinach, N. and Hoehndorf, R. (2017). Neuro-symbolic representation learning on biological knowledge graphs. *Bioinformatics*, **33**(17), 2723–2730.
- [2] Belanger, D. and McCallum, A. (2016). Structured prediction energy networks. In *International Conference on Machine Learning*, pages 983–992.
- [3] Chase Lipton, Z., Elkan, C. and Narayanaswamy, B. (2014). Thresholding classifiers to maximize f1 score. *arXiv preprint arXiv:1402.1892*.
- [4] Chen, M., Ju, C.J.T., Zhou, G., Chen, X., Zhang, T., Chang, K.W. et al (2019). Multifaceted protein–protein interaction prediction based on siamese residual rcnn. *Bioinformatics*, **35**(14), i305–i314.
- [5] Conneau, A., Kiela, D., Schwenk, H., Barrault, L. and Bordes, A. (2017). Supervised learning of universal sentence representations from natural language inference data. *arXiv preprint arXiv:1705.02364*.
- [6] Devlin, J., Chang, M.W., Lee, K. and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- [7] Duong, D., Ahmad, W.U., Eskin, E., Chang, K.W. and Li, J.J. (2018). Word and sentence embedding tools to measure semantic similarity of gene ontology terms by their definitions. *Journal of Computational Biology*, **26**(1), 38–52.
- [8] Gene Ontology Consortium (2017). Expansion of the gene ontology knowledgebase and resources. *Nucleic acids research*, **45**(D1), D331–D338.
- [9] Kipf, T.N. and Welling, M. (2016). Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- [10] Kulmanov, M., Khan, M.A. and Hoehndorf, R. (2017). Deepgo: predicting protein functions from sequence and interactions using a deep ontology-aware classifier. *Bioinformatics*, **34**(4), 660–668.
- [11] Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C.H. et al (2019). Biobert: pre-trained biomedical language representation model for biomedical text mining. *arXiv preprint arXiv:1901.08746*.
- [12] Mazandu, G.K. and Mulder, N.J. (2014). Information content-based gene ontology functional similarity measures: Which one to use for a given biological data type? *PLoS ONE*, **9**(12), e113859.
- [13] Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S. and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- [14] Pesaranghader, A., Matwin, S., Sokolova, M. and Beiko, R.G. (2015). simdef: definition-based semantic similarity measure of gene ontology terms for functional similarity analysis of genes. *Bioinformatics*, **32**(9), 1380–1387.

- [15] Pesquita, C., Faria, D., Bastos, H., Ferreira, A.E., Falcão, A.O. and Couto, F.M. (2008). Metrics for go based protein semantic similarity: a systematic evaluation. *BMC bioinformatics*, **9**(5), S4.
- [16] Pesquita, C., Pessoa, D., Faria, D. and Couto, F. (2009). Cessm: Collaborative evaluation of semantic similarity measures. *JB2009: Challenges in Bioinformatics*, **157**, 190.
- [17] Profiti, G., Martelli, P.L. and Casadio, R. (2017). The bologna annotation resource (bar 3.0): improving protein functional annotation. *Nucleic acids research*, **45**(W1), W285–W290.
- [18] Resnik, P. (1999). Semantic similarity in a taxonomy: An information-based measure and its application to problems of ambiguity in natural language. *J. Artif. Intell. Res.(JAIR)*, **11**, 95–130.
- [19] Rios, A. and Kavuluru, R. (2018). Few-shot and zero-shot multi-label learning for structured label spaces. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, volume 2018, page 3132. NIH Public Access.
- [20] Rong, X. (2014). word2vec parameter learning explained. *arXiv preprint arXiv:1411.2738*.
- [21] Smaili, F.Z., Gao, X. and Hoehndorf, R. (2018). Onto2vec: joint vector-based representation of biological entities and their ontology-based annotations. *Bioinformatics*, **34**(13), i52–i60.
- [22] Song, X., Li, L., Srimani, P.K., Yu, P.S. and Wang, J.Z. (2014). Measure the semantic similarity of GO terms using aggregate information content. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, **11**(3), 468–476.
- [23] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N. et al (2017). Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- [24] Wang, S., Cho, H., Zhai, C., Berger, B. and Peng, J. (2015). Exploiting ontology graph for predicting sparsely annotated gene function. *Bioinformatics*, **31**(12), i357–i364.
- [25] Wu, X.Z. and Zhou, Z.H. (2017). A unified view of multi-label performance measures. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3780–3788. JMLR. org.
- [26] Xiao, H. (2018). bert-as-service. [github.com/hanxiao/bert-as-service](https://github.com/hanxiao/bert-as-service).
- [27] Xiong, W., Yu, M., Chang, S., Guo, X. and Wang, W.Y. (2018). One-shot relational learning for knowledge graphs. *CoRR*, **abs/1808.09040**.
- [28] Yang, H., Nepusz, T. and Paccanaro, A. (2012). Improving go semantic similarity measures by exploring the ontology beneath the terms and modelling uncertainty. *Bioinformatics*, **28**(10), 1383–1389.
- [29] Zhang, C., Zheng, W., Freddolino, P.L. and Zhang, Y. (2018). Metago: Predicting gene ontology of non-homologous proteins through low-resolution protein structure prediction and protein–protein network mapping. *Journal of molecular biology*, **430**(15), 2256–2265.
- [30] Zhang, Z., Zhang, J., Fan, C., Tang, Y. and Deng, L. (2017). Katzlgo: large-scale prediction of lncrna functions by using the katz measure based on multiple networks. *IEEE/ACM transactions on computational biology and bioinformatics*, **16**(2), 407–416.