

Bird's Eye: Probing for Linguistic Graph Structures with a Simple Information-Theoretic Approach

Yifan Hou

Department of Computer Science
ETH Zürich
yifan.hou@inf.ethz.ch

Mrinmaya Sachan

Department of Computer Science
ETH Zürich
mrinmaya.sachan@inf.ethz.ch

Abstract

NLP has a rich history of representing our prior understanding of language in the form of graphs. Recent work on analyzing contextualized text representations has focused on hand-designed probe models to understand how and to what extent do these representations encode a particular linguistic phenomenon. However, due to the inter-dependence of various phenomena and randomness of training probe models, detecting how these representations encode the rich information in these linguistic graphs remains a challenging problem. In this paper, we propose a new information-theoretic probe, *Bird's Eye*, which is a fairly simple probe method for detecting if and how these representations encode the information in these linguistic graphs. Instead of using classifier performance, our probe takes an information-theoretic view of probing and estimates the mutual information between the linguistic graph embedded in a continuous space and the contextualized word representations. Furthermore, we also propose an approach to use our probe to investigate localized linguistic information in the linguistic graphs using perturbation analysis. We call this probing setup *Worm's Eye*. Using these probes, we analyze BERT models on their ability to encode a syntactic and a semantic graph structure, and find that these models encode to some degree both syntactic as well as semantic information; albeit syntactic information to a greater extent. Our implementation is available in https://github.com/yifan-h/Graph_Probe-Birds_Eye.

1 Introduction

Graphs have served as a predominant representation for various linguistic phenomena in natural language (Marcus et al., 1993; De Marneffe et al., 2006; Hockenmaier and Steedman, 2007; Hajic et al., 2012; Abend and Rappoport, 2013; Banarescu et al., 2013; Bos, 2013). These graph based

representations have served our intuition for representing both language structure (Chomsky, 1957) as well as meaning (Koller et al., 2019).

With the growing popularity of pretrained language models that build contextualized text representations (Reid et al., 2020; Devlin et al., 2019, inter alia), various probing models have been introduced to understand if and how our linguistic intuitions are encoded in these representations. These probes train supervised models to predict pieces of linguistic information such as POS (part-of-speech), morphology, syntactic and semantic relations, and other local or long-range phenomena in language (Belinkov et al., 2017; Conneau et al., 2018; Hewitt and Manning, 2019; Tenney et al., 2019b; Jawahar et al., 2019). However, it is still an open question if these representations somehow encode entire linguistic graph structures such as dependency and constituency parse trees or graph structured meaning representations such as AMR (Abstract Meaning Representation), UCCA (Universal Conceptual Cognitive Annotation), etc.

A popular recent work, the *structural probe* (Hewitt and Manning, 2019), has investigated how contextualized representations encode syntax trees. They tested if a linear transformation of the network's word representation space can predict *particular features* of the syntax tree, namely, the *distance between words* and *depth of words* in the tree. Thus, the structural probe cannot by itself answer the question if these representations encode entire linguistic graph structures. Moreover, the structural probe is only designed for tree structures and cannot be extended to general graphs.

In this work, we introduce a new probing approach, *Bird's Eye*, which can be used to detect if contextualized text representations encode entire linguistic graphs. *Bird's Eye* is a simple information-theoretic probe (Pimentel et al., 2020b) which first encodes the linguistic graph into a con-

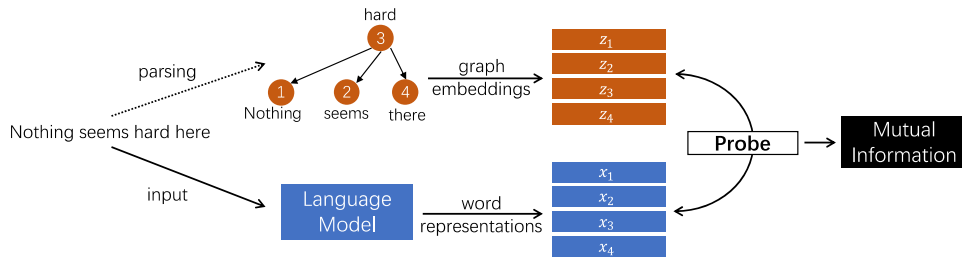


Figure 1: Methodology of *Bird's Eye*: To probe pretrained language models, linguistic graphs are embedded in a continuous space and the mutual information between graph embeddings and word representations is calculated.

tinuous representation using graph embedding approaches (Cai et al., 2018) and then, estimates the mutual information between the linguistic graph representation space and the contextualized word representation space. An illustration of the probe approach is given in Figure 1. The information theoretic approach is more reliable than training a probe and using accuracy for probing as it is debatable if the classifier-based probe is probing or trying to solve the task (Hewitt and Liang, 2019; Pimentel et al., 2020b). We further extend *Bird's Eye* to probe for localized linguistic information in the linguistic graphs such as POS or dependency arc labels in dependency parses. We call this probe, *Worm's Eye*.

In our experiments, we first illustrate the reliability of our probe methods and show the randomness of previous probe methods that use accuracy. Then, we use *Bird's Eye* to detect syntactic and semantic structures in BERT, showing that much syntactic and some semantic structure are encoded in BERT. Besides, we also use *Worm's Eye* to probe for specific linguistic information in syntactic trees and semantic graphs respectively to see which kinds of localized linguistic information is encoded in BERT. Our probing results are consistent with previous probe methods (Hewitt and Manning, 2019; Reif et al., 2019; Liu et al., 2019; Tenney et al., 2019a,b; Wu et al., 2021). We also discuss limitations of our probe and how future work can build upon our foundation.

2 *Bird's Eye* Probe

In this section, we introduce our information-theoretic approach for probing linguistic graph structures in word representations. The MI estimate is used to understand how much of the information in the linguistic graph structure has been learnt by the pretrained models.

Let $X = \{x_1, \dots, x_T\}$ denote an input sentence (each x_i is the contextual embedding of a token in the given vocabulary \mathcal{V}) and G denote the cor-

responding linguistic graph. Furthermore, let \mathcal{X} denote a random variable that takes values ranging over all possible token sequences in \mathcal{V} . Correspondingly, let \mathcal{G} denote a random variable that ranges over all possible corresponding linguistic graphs. We use $I(\mathcal{X}; \mathcal{G})$ to denote the linguistic structure information that is included in the given word representations. Note that the MI value $I(\mathcal{X}; \mathcal{G})$ is always non-negative, and a large MI implies that more of the structure information is encoded in the word representations. In order to make the MI computation easier, we additionally assume alignments between the nodes V in the graph G and the words in X . This alignment is one to one, for example, in dependency parsing (Marcus et al., 1993) but an aligner might be needed in some cases (Banarescu et al., 2013).

There are three main challenges in estimating MI in our setting. First, the MI estimation of discrete graphs and continuous features has been an elusive problem (Ross, 2014; Kraskov et al., 2004; Escolano et al., 2017), since there is no widely accepted definition of mutual information in this setting. Second, the dimensionality of the contextualized word representations is very high. Traditional methods (Moon et al., 1995; Steuer et al., 2002; Paninski, 2003) for MI estimation do not scale well with large sample size or dimension (Gao et al., 2015). Getting accurate estimates of mutual information in the high dimension is not easy. Third, graphs across different linguistic formalisms could have different entropy values, and thus the MI value $I(\mathcal{X}; \mathcal{G})$ may be uncomparable across the different linguistic graph formalisms. For example, if syntactic trees \mathcal{G} and semantic graphs \mathcal{G}' have the same MI value with \mathcal{X} i.e. $I(\mathcal{X}; \mathcal{G}) = I(\mathcal{X}; \mathcal{G}')$ while the entropy values are fairly different i.e. $I(\mathcal{X}; \mathcal{G}) \approx H(\mathcal{G}) \ll H(\mathcal{G}')$, it is not proper to conclude that \mathcal{X} contains the same amount of information from structures \mathcal{G} and \mathcal{G}' , since they correspond to different percentages of the amount of uncertainty. Thus, the MI values must be inter-

preted carefully.

`Bird's Eye` tackles the aforementioned difficulties by transforming the linguistic graphs into a continuous space using a graph embedding approach. Then the MI between graph embeddings and word representations is estimated using a recently proposed method (Belghazi et al., 2018) which performs well even in high dimensions. Finally, we also estimate a lower and upper bound of the MI, which is used to interpret the MI value. We describe various stages of `Bird's Eye` below:

2.1 Graph Embedding

The provided linguistic graphs can typically be represented as an adjacency matrix. Directly calculating MI with the adjacency matrix is not useful due to the sparsity and discreteness of the adjacency matrix representation. Thus, we transform the graphs into a continuous space where each node is represented by a continuous representation of same dimensionality.

Theoretically, if the graph embedding approach is perfect, we can use the invariant property of mutual information (Kraskov et al., 2004). This property states that under some fairly strong conditions, there exists an invertible function $f(\cdot)$ that satisfies $\mathcal{G} = f^{-1}(f(\mathcal{G}))$, where the graph embeddings are $\mathcal{Z} = f(\mathcal{G})$. Thus, we can transform \mathcal{G} into graph embeddings \mathcal{Z} , and:

$$I(\mathcal{X}; \mathcal{Z}) \approx I(\mathcal{X}; \mathcal{G}) \quad (1)$$

In this paper, we use DeepWalk (Perozzi et al., 2014), which is based on the skip-gram model (Huang et al., 1993; Mikolov et al., 2013) for graph embeddings¹. Specifically, given a node $v \in V$ encoded as the one-hot vector $\mathbb{1}_v$, the model tries to predict its neighbor's vector $\mathbb{1}_{v'}$ where $v' \in N_v$. The graph $G = \{V, E\}$ is first sampled to generate a set of random walks. Then the graph neighborhood relationship is represented by the co-occurrence of nodes in the walk paths. Finally, for all the walks, Word2vec (Mikolov et al., 2013) with skip-gram is used to maximize the co-occurrence likelihood²:

$$\mathcal{L}(\theta) = \prod_{v \in V} \prod_{v' \in N_v} \mathbb{P}(\mathbb{1}_{v'} | \mathbb{1}_v; \theta). \quad (2)$$

¹Note that the `Bird's Eye` probe is a general probe. Other graph embedding approaches can also be selected for the transformation under specific conditions.

²Note that in Word2vec, the window size is a hyperparameter that need to be selected by users. Here, for simplicity, we set the window size as 1.

Let $Z = \oplus\{z_v | v \in V\}$ denote the learnt graph embedding where z_v is the embedding of node v . Here \oplus denotes the concatenation operation.

In our experiments, we also explore to what extent the original linguistic graphs can be restored by the graph embeddings, which tests the extent to which eq. 1 holds and if we can use $I(\mathcal{X}; \mathcal{Z})$ instead of $I(\mathcal{X}; \mathcal{G})$ to estimate MI. More details can be found in Appendix A

2.2 Mutual Information Estimation

To estimate $I(\mathcal{X}; \mathcal{Z})$ in high dimensions, we maximize the compression lemma lower bound (Banerjee, 2006) as mentioned in Belghazi et al. (2018). Specifically, for a pair of random variables \mathcal{X} and \mathcal{Z} , the mutual information is equivalent to the Kullback-Leibler (KL) divergence between the joint distribution $\mathbb{P}_{\mathcal{X}\mathcal{Z}}$ and the product of the marginal distributions $\mathbb{P}_{\mathcal{X}} \otimes \mathbb{P}_{\mathcal{Z}}$:

$$I(\mathcal{X}; \mathcal{Z}) = D_{KL}(\mathbb{P}_{\mathcal{X}\mathcal{Z}} || \mathbb{P}_{\mathcal{X}} \otimes \mathbb{P}_{\mathcal{Z}}). \quad (3)$$

From the compression lemma lower bound (Banerjee, 2006), the KL divergence $D_{KL}(\mathbb{P} || \mathbb{Q})$ can be bounded as:

$$D_{KL}(\mathbb{P} || \mathbb{Q}) \geq \sup_{T \in \mathcal{F}} \mathbb{E}_{\mathbb{P}}[T] - \log(\mathbb{E}_{\mathbb{Q}}[e^T]), \quad (4)$$

where \mathcal{F} can be any class of functions $T : \Omega \rightarrow \mathbb{R}$ satisfying certain integrability constraints. Thus, in the inequality 4, the lower bound can be obtained by finding a function in the set \mathcal{F} :

$$I(\mathcal{X}; \mathcal{Z}) \geq \sup_{T \in \mathcal{F}} \mathbb{E}_{\mathbb{P}_{\mathcal{X}\mathcal{Z}}}[T] - \log(\mathbb{E}_{\mathbb{P}_{\mathcal{X}} \otimes \mathbb{P}_{\mathcal{Z}}}[e^T]).$$

To get a tight estimate of $I(\mathcal{X}; \mathcal{Z})$, we need the lower bound to be as high as possible. Thus, the MI estimation problem turns into an optimization problem to maximize the compression lemma lower bound. To ensure that, similar to Belghazi et al. (2018), we let $\mathcal{F} = \{T_{\theta}\}_{\theta \in \Theta}$ be the set of functions parametrized by a neural network, and optimize the neural network using stochastic gradient descent. Formally, the objective function is:

$$\max_{\theta \in \Theta} (\mathbb{E}_{\mathbb{P}_{\mathcal{X}\mathcal{Z}}^{(n)}}[T_{\theta}] - \log(\mathbb{E}_{\mathbb{P}_{\mathcal{X}}^{(n)} \otimes \mathbb{P}_{\mathcal{Z}}^{(n)}}[e^{T_{\theta}}])). \quad (5)$$

Here, $\mathbb{P}_{\mathcal{X}\mathcal{Z}}^{(n)}$, $\mathbb{P}_{\mathcal{X}}^{(n)}$ and $\mathbb{P}_{\mathcal{Z}}^{(n)}$ are empirical joint and marginal distributions over a sample of n (sentence, graph) pairs.

We calculate graph embeddings for each sentence independently, and regard one sentence as

a mini-batch to optimize the neural network iteratively for MI estimation. Note that different from existing probe models, our objective of the neural network is to find an optimal function in \mathcal{F} and estimate MI, rather than use prediction accuracy. Besides, the neural network is very simple (MLP). Therefore, there is no need to split dataset into training and test to test generalization in MI estimation³. The negative of the training loss as eq. 5 can be taken as MI estimation directly (Belghazi et al., 2018; Cristiani et al., 2020). In our experiments, we verify the effectiveness of the MI estimation method to prove that the probe is stable. More technical details of the MI estimation model and how it is trained are given in Appendix B.

2.3 Control Bounds

Next, we introduce two control bounds to interpret the MI value, whose functions are similar to the control task introduced by Hewitt and Liang (2019). As mentioned, comparing MI alone across different types of structures is not useful, since the entropy values of graph embeddings can also be different. Thus, we calculate an upper and a lower bound of the MI value based on the graph structures. Instead of using the MI value alone, we interpret it by its relative value in terms of the two control bounds. Formally, for the MI between graph embeddings and word representations, we have:

$$I(\mathcal{R}; \mathcal{Z}) \leq I(\mathcal{X}; \mathcal{Z}) \leq I(\mathcal{Z}; \mathcal{Z}). \quad (6)$$

The lower bound is the MI between a truly random variable \mathcal{R} (i.e., independent of the graph) and the graph embedding \mathcal{Z} . Thus, $I(\mathcal{R}; \mathcal{Z}) = 0$. The upper bound telescopes to the graph structure’s self-entropy⁴.

Using these two control bounds, we interpret the structure information by the relative MI value⁵:

$$MIG(\mathcal{G}) = \frac{\hat{I}(\mathcal{X}; \mathcal{Z}) - \hat{I}(\mathcal{R}; \mathcal{Z})}{\hat{I}(\mathcal{Z}; \mathcal{Z}) - \hat{I}(\mathcal{R}; \mathcal{Z})}, \quad (7)$$

The MI estimates $\hat{I}(\mathcal{Z}; \mathcal{Z})$ and $\hat{I}(\mathcal{R}; \mathcal{Z})$ can be obtained in the same way as $\hat{I}(\mathcal{X}; \mathcal{Z})$ (using the MI estimation method mentioned above). MIG (eq. 7)

³Alternatively, the dataset can be divided evenly into training and test for MI estimation.

⁴Note that for continuous random variables \mathcal{Z} , the number of values that \mathcal{Z} can take is infinite. In this condition, $I(\mathcal{Z}; \mathcal{Z})$ tends to infinity. Thus, we use a small noise ϵ and approximate it as $I(\mathcal{Z} + \epsilon, \mathcal{Z})$.

⁵The definition is similar to the uncertainty coefficient.

scales the MI value for graph embeddings with different self-entropy values into the same range: $MIG(\mathcal{G}) \in [0, 1]$. Intuitively, MIG captures what percentage of the structure information is encoded in word representations.

Since $MIG(\mathcal{G})$ is scaled using $I(\mathcal{R}; \mathcal{Z})$, it also helps reduce the error in MI estimation. As mentioned, we maximize compression lemma lower bound 5 as the MI estimate. However, there could be a gap between it and the ground-truth MI value. Based on the fact that the ground-truth $I(\mathcal{R}; \mathcal{Z}) = 0$, we can know that the gap $I(\mathcal{R}; \mathcal{Z}) - \hat{I}(\mathcal{R}; \mathcal{Z})$ is equal to $-\hat{I}(\mathcal{R}; \mathcal{Z})$. In MIG (eq. 7), the gap is added for both numerator and denominator, which reduces the error brought by MI estimation⁶.

2.4 Worm’s Eye Probe for Localized Linguistic Structure

Bird’s Eye allows us to probe for entire linguistic structures. However, for us to have a complete understanding, we might also want to probe for some localized information in the linguistic graphs. For example, we may want to know if BERT knows about POS tags or certain dependency relations in the syntax parse. We formulate probing for localized linguistic information as probing for a subgraph of the linguistic graph and reuse our Bird’s Eye probe for it. We call this setting Worm’s Eye as we are now analyzing if these representations capture local sub-structures.

To probe localized linguistic information $G_s = \{V_s, E_s\}$, we use perturbation of the original structure for analysis. Specifically, we add a perturbation to the original graph embedding \mathcal{Z} based on the subgraph G_s . For all the nodes in V_s or nodes connected by edges in E_s , we add a noise on their corresponding node representations in \mathcal{Z} . Let \mathcal{Z}' denote the corrupted graph embedding. Then, we define the following:

$$MIL(\mathcal{G}_s) = 1 - \frac{\hat{I}(\mathcal{X}; \mathcal{Z}') - \hat{I}(\mathcal{R}; \mathcal{Z})}{\hat{I}(\mathcal{X}; \mathcal{Z}) - \hat{I}(\mathcal{R}; \mathcal{Z})}, \quad (8)$$

MIL describes how much MI is contributed by the local structure \mathcal{G}_s . When the local structure is the whole graph, \mathcal{Z}' is completely noisy and $MIL(\mathcal{G}_s)$ equals to 1, which means the entire MI value $I(\mathcal{X}; \mathcal{Z})$ is contributed by the local structure. If the local structure is an empty set, we have

⁶In our experiment, we show that the estimated values satisfy $|\hat{I}(\mathcal{R}; \mathcal{Z})| < 10^{-3} \times \hat{I}(\mathcal{Z}; \mathcal{Z})$, which is small enough to be ignored.

$\mathcal{Z}' = \mathcal{Z}$. Then we can get $MIL(\mathcal{G}_s) = 0$, representing that the local structure does not contribute anything to the MI value.

If we control the perturbation of different types of local structures at the same level, we can compare how well they are captured by the word representations relative to each other using eq. 8. Specifically, for relations with labels, e.g., types of dependency relations in syntax trees, we set the same perturbation on the graph embeddings. Then, we test and compare $MIL(\mathcal{G}_s)$ for different types of relations. Larger $MIL(\mathcal{G}_s)$ for a particular relation type implies that more information about this relation type is encoded in the word representations.

3 Probing for Syntactic and Semantic Graph Structures

We use our Bird's Eye probe to detect two linguistic structures in the pretrained models, namely, dependency syntax (Marcus et al., 1993; De Marneffe et al., 2006) and a more semantic formalism, AMR (Banarescu et al., 2013).

We first use our model to probe for Stanford dependencies (de Marneffe et al., 2006). For a sentence X with tokens $\{x_1, x_2, \dots, x_T\}$, the syntax tree defines a directed labelled tree where tokens x_i are represented as nodes and relations among them as labeled edges. We ignore the edge direction and labels for simplicity in our work⁷. Future work can consider incorporating edge direction and labels. We embed the given syntax tree into a continuous space as mentioned before. Then, we calculate the MIG (eq. 7) as described before to determine how much syntax information is captured in the given contextualized representations.

Next, we test if contextualized representations capture a semantic graph representation – the Abstract Meaning Representation (AMR) (Banarescu et al., 2013). Different from syntactic trees, semantic graphs are not tree structured, and there can be loops or reentrancies. In the AMR annotation, plurality, articles and tenses were dropped and thus, there is no one-to-one corresponding between words in the sentence and nodes in the AMR graph. Thus, we use an off-the-shelf aligner (Pourdamghani et al., 2014) and calculate MI between the AMR graph embedding and the representations of those words that are aligned with a node in the AMR graph. For simplicity, edge directions and

⁷The Stanford dependency tree also contains one empty root node, which is also ignored

labels are also ignored in this setting.

4 Experiments

Our experiments mainly comprise of two parts:

1. Verification of the probe: The first part is for verification of the probing methodology and ensuring that the graph embeddings retain information about the linguistic graphs i.e eq. 1 holds. We do this by testing if the graph embeddings can be used to restore the original graph.

2. Probing for graph structures: The second part is about using the probe to detect syntactic and semantic graph structures in BERT. Importantly, we probe if pretrained BERT captures entire graph structures as well as specific relational information in these linguistic graphs. To contrast with previous accuracy and training based probes, we also train a group of simple MLP models with different number of hidden layers and use accuracy for probing. We show that designing and training a model to probe entire or localized linguistic structures is not as reliable as our information-theoretic approach.

We use gold annotations from the Penn Treebank and the AMR Bank for all our experiments. For the contextualized word representations, we select pretrained BERT models, specifically BERT-base (uncased) and BERT-large (uncased). Since BERT generates word-piece embeddings, to align them with gold word-level tokens, we represent each token as the average of its word-piece embeddings as in Hewitt and Manning (2019). We also use two non-contextual word embeddings as baselines: GloVe embeddings (Pennington et al., 2014) and ELMo-0, character-level word embeddings with no contextual information generated by pretrained ELMo (Reid et al., 2020).

4.1 Evaluation of Graph Embeddings

We first evaluate how well the graph embeddings can capture the linguistic graph structures by predicting the original graphs with them. We use simple MLPs of 6 different settings with varying number of hidden layers. More details can be found in Appendix C. We use AUC score as the metric to evaluate the graph prediction performance, which is a common metric in link prediction that computes area under the ROC curve (Fawcett, 2006).

The results are presented in Table 1. We can see that for both syntax trees and semantic graphs, MLPs can achieve good performance in restoring the original graph using graph embeddings where

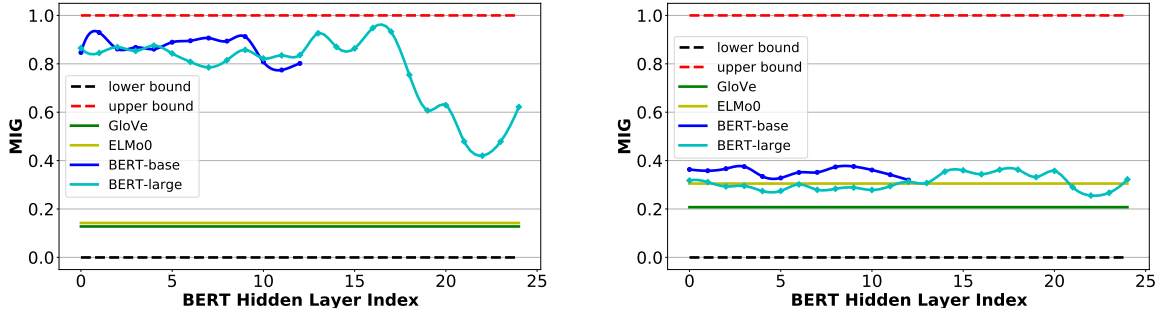


Figure 2: *MIG* scores with syntactic and semantic structures, respectively for word representations in BERT models (BERT-base with 12 layers and BERT-large with 24 layers). Note that results at the input layer are also reported, where the BERT Hidden Layer Index is 0).

Table 1: Performance (AUC score for link prediction) on restoring graphs with graph embeddings

# of hidden layers	Syntax tree	AMR graph
0	0.5620	0.5620
1	0.6330	0.5494
2	0.9637	0.8804
3	0.9780	0.9263
4	0.9806	0.9162
5	0.9791	0.9192

the AUC score is quite high. Thus, we can be confident that equation 1 holds, and we can calculate MI based on the graph embeddings. Future work can explore better graph embedding approaches. We also evaluate our probe by adding noisy representations to the graph embeddings to prove that it is capable of teasing out different levels of dependencies. Details can be referred to in Appendix D.

4.2 Probing Entire Structures

We first used the `Bird's Eye` probe to detect if entire linguistic structures are encoded in hidden representations of BERT⁸. We also include two non-contextual word representations – GloVe and ELMo-0 as baselines. We report *MIG* as the results of our probe on the two graph structures in Figures 2(a) and 2(b).

The *MIG* estimations for syntactic structure probing of both BERT-base and BERT-large are quite high, which implies that BERT encodes much syntactic information. However, for the semantic structure, the *MIG* scores of BERT models are lower, suggesting that BERT does not encode the semantic structures as well. These two conclusions are consistent with previous works (Liu et al., 2019;

⁸For all MI estimation experiments, we repeat the experiment 20 times and take the average to get stable results.

Tenney et al., 2019b; Wu et al., 2021) which have found that unlike syntax, semantics is not captured well by the pretrained models.

We also observe an interesting trend when comparing *MIG* across layers. We find that for syntax, *MIG* starts to decrease in the upper layers, especially for the BERT-large. This is consistent with previous works which report that BERT models syntax more in the lower and middle layers (Tenney et al., 2019a). For semantic graphs, *MIG* is steady across all layers. It means that semantic information is spread across the entire model. The results are consistent with existing work (Rogers et al., 2020). For the two non-contextual baselines, GloVe and ELMo-0, we can see that their *MIG* scores are lower compared with contextualized representations, especially for syntax. Previous work (Hewitt and Manning, 2019) has drawn similar conclusions. While for the semantic graphs, the gap is not significant.

4.3 Probing Localized Information

In this section, we show how we can use the `Worm's Eye` probe to understand if the contextualized representations capture localized linguistic information in the dependency parses such as POS information or relational dependency information. As described before, we design various perturbation experiments using our `Worm's Eye` probe. For probing POS information or a dependency relation type, we add noise to the graph embeddings of the corresponding node(s). After that, we calculate the *MIL* ratio (eq. 8) to show how much particular linguistic information (POS or relation type information) is contained in the word representations. We repeat the experiment 20 times and

use boxplots to present all the results.

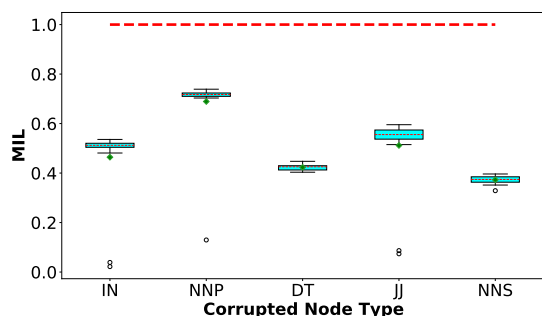


Figure 3: *MIL* scores of probing 5 types of POS tags (localized syntactic structure) for word representations in BERT-base (output layer). The local structure is decided by the POS tags attached on nodes.

First we use `Worm's Eye` to test for POS information, which is tagged as node labels in the dependency tree. We select 5 POS tags: *IN*, *NNP*, *DT*, *JJ*, and *NNS*, which have high and roughly the same frequencies in the Penn Treebank dataset. Complete statistics about the POS tag frequencies can be found in Appendix E. We ensure that the amount of perturbation of the graph embeddings is the same for each type. Figure 3 presents the results. We find that *NNP* achieves the highest *MIL* score, while *NNS* achieves the lowest. This implies that BERT encodes syntactic information for singular proper nouns (*NNP*) and adjectives (*JJ*) more than plural nouns (*NNS*).

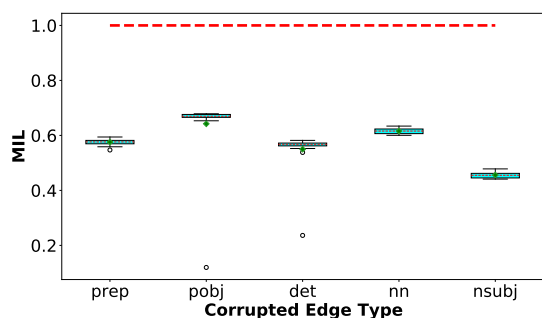


Figure 4: *MIL* scores of probing 5 types of localized syntactic structure for word representations in BERT-base (output layer). The local structure is decided by the dependency relation labels attached on edges.

Next, we probe 5 types of universal dependency relations in the Penn Treebank dataset (PTB). These are *prep*, *pobj*, *det*, *nn* and *nsubj*. These 5 relations also roughly occur the same number of times in PTB. Complete statistics about the number of occurrences of these relation types can be found in Appendix E. Similarly, for each type of

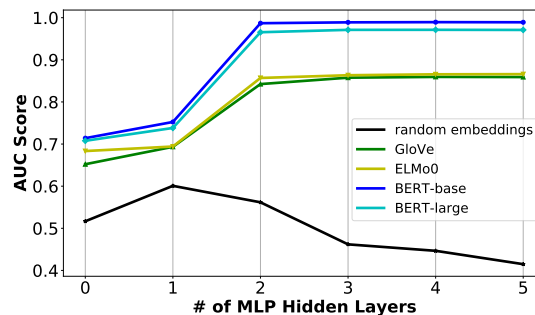


Figure 5: AUC scores of predicting syntactic trees by various word representations.

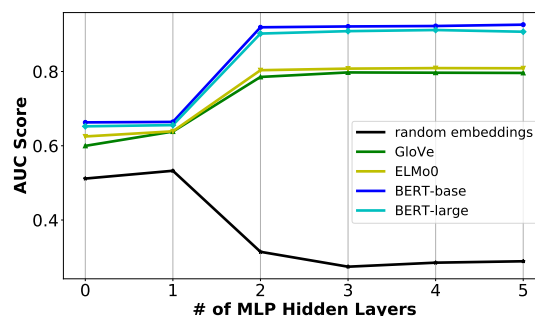


Figure 6: AUC scores of predicting semantic graphs by various word representations.

relation, we add same amount of perturbation to graph embeddings of nodes connected by the specific relations. Figure 4 shows the results, where *nsubj* relations have the lowest *MIL* score compared with other 4 types. This means that BERT encodes more syntactic structure for prepositional modifiers (*prep*), object of a preposition (*pobj*), and noun compound modifier (*nn*) than nominal subject (*nsubj*). Reif et al. (2019) have drawn similar conclusions while probing for dependency arc labels. Similar experiment for semantic structure can be found in Appendix F.

4.4 On Accuracy-Based Probing

In contrast to our information-theoretic approach to probing, we train a group of MLP models to probe entire and local structures in BERT-base. We show that these probe results mainly depend on the model complexity rather than the structure itself.

Probing entire graph structures. A group of MLPs are trained to predict entire syntactic and semantic structures with word representations. Figure 5 and Figure 6 show the results. Their trends are similar. Shallow MLPs perform the worst and deep ones perform much better. Previous work on structural probing (Hewitt and Manning, 2019) argues that powerful models could parse the word repre-

sentations, thus a simple model should be designed. However, in Table 1, we find that linear model even could not restore the graph by its embeddings. Obviously, its performance cannot indicate how much structure information is included in the graph embeddings. Thus, there is no reasonable principle to decide the complexity of the probe model. Given this, designing and training a model is not suitable to probe entire structures. A similar argument has been placed by previous works (Pimentel et al., 2020a,b; Lovering et al., 2021).

Table 2: Performance (AUC score) of predicting graph structures with MLPs of different complexity

Relation Type	linear	1	2	3	4	5
prep	0.6966	0.7232	0.9838	0.9858	0.9863	0.9866
pobj	0.6224	0.6619	0.9874	0.9888	0.9891	0.9894
det	0.7016	0.7222	0.9928	0.9938	0.9940	0.9943
nn	0.6877	0.7221	0.9884	0.9899	0.9904	0.9903
nsubj	0.6754	0.6938	0.9841	0.9859	0.9868	0.9869
arg	0.6686	0.6652	0.9185	0.9217	0.9189	0.9215
general	0.6621	0.6574	0.9192	0.9223	0.9221	0.9221
op	0.6500	0.6500	0.9098	0.9177	0.9130	0.9160

Probing localized graph structures. To prove that accuracies of probe models for localized structure also mainly depend on the model’s complexity rather than the local structure, we train the group of MLPs to predict the entire syntactic structure by word representations, and calculate the AUC scores for each type of relations in test set as probing results. Table 2 shows the AUC score of predicting specific type of relations. For syntactic structure, same 5 types of relations are selected, and for semantic graphs, we select 3 groups of relations to probe: *arg*, *general* and *op*. Complete statistics of AMR Bank are in Appendix E. From the results, we can find that for MLP models with different number of hidden layers, the ranks of AUC scores of relation prediction are quite different. For both syntax trees and semantic graphs, there is no consistent interpretation of the results to conclude which types of relations are encoded in BERT. We also run the experiment in the perturbation settings, which can be referred to Appendix G.

Combining the results of probing with accuracy in Figure 5, Figure 6, and Table 2, we can find that the prediction decisions are not based purely on the structure but rather on spurious heuristics. This has also been concluded and discussed in some recent works (Hewitt and Liang, 2019; Lovering et al., 2021). Thus, training models is not feasible to probe structures. For our probe methods, the randomness of models such as complexity is not an issue, since the one with highest estimation should

be selected for tighter compression lemma bound 4 as introduced by Pimentel et al. (2020b).

4.5 Hyperparameter and Efficiency

Information-theoretic approaches sacrifice simplicity and efficiency to achieve reliable probing results compared to accuracy-based probes. Even though our probes are quite simple, there are more hyperparameters that need to be selected by users compared to accuracy-based probes. To help users implement our methods in their setting, we briefly describe some guiding principles to help them select hyperparameters, and point out several potential ways to make our probing approach more efficient.

Our probes are composed of two steps: (a) computation of the graph embedding, and (b) estimation of the mutual information. The guiding principle in the graph embedding step is to retain as much linguistic graph information as possible. In our experiments, we used default hyperparameters in DeepWalk (Perozzi et al., 2014) for simplicity. Details can be found in Appendix A. However, users may also use other graph embedding approaches that incorporate edge labels, etc. to improve our model. As the mutual information estimation procedure is estimating a lower bound to the true mutual information, the guiding principle for hyperparameter selection in this step should be to let the MI estimation values be as large as possible. In particular, model size is worth noting. Deeper models can achieve a tighter lower bound. However, these are less efficient than shallow ones. Thus, the selection of MI estimator’s complexity is a tradeoff. According to our empirical experience, a relatively good choice is to use a two-layer MLPs. More details can be found in Appendix D. Note that it might also be harder to achieve convergence with deeper models as training of MI estimators is notoriously difficult. We leave a better exploration of this to future work.

Potential users might also resort to other solutions to make the probes more efficient. If the bottleneck is in the graph embedding step, some fast approaches (Hamilton et al., 2017; Tang et al., 2015) can be chosen instead. If the mutual information estimation step is the bottleneck, some sampling strategies can be used. A simple way is to sample a subset of the dataset, and optimize eq. 5 based on that subset. Alternatively, potential users can use more sophisticated sampling strategies in training as in Recht and Ré (2012). These

approaches achieve a much better convergence rate for MI estimation.

5 Related Work

Syntax and Semantics Probing. Many existing works probe language models directly or indirectly showing how much syntactic and semantic information is encoded in them. [Belinkov et al. \(2017\)](#) tested NMT models and found that higher layers encode semantic information while lower layers perform better at POS tagging. Similarly, [Jawahar et al. \(2019\)](#) tested various BERT layers and found that it encodes a rich hierarchy of linguistic information in the intermediate layers. [Tenney et al. \(2019b\)](#); [Wu et al. \(2021\)](#) compared the syntactic and semantic information in BERT and its variants, and found that more syntactic information is encoded than semantic information. [Conneau et al. \(2018\)](#) focused on probing various linguistic features with 10 different designed tasks. [Hewitt and Manning \(2019\)](#) designed a tree distance and depth prediction task to probe syntax tree structures.

Information Theoretic Probe. With the popularity of probe methods, limitations of previous methods have also been found. Information theoretic methods have been proposed as an alternative. To avoid the randomness of performance brought by the varying sizes of the probe models, [Pimentel et al. \(2020b\)](#) proposed an information-theoretic probe with control functions, which used mutual information instead of model performance for probing. [Voita and Titov \(2020\)](#) restricted the probe model size by Minimum Description Length. Training a model is recast as teaching it to effectively transmit the data. [Lovering et al. \(2021\)](#) pointed out that if we train a model to probe, the decisions are often not based on information itself, but rather on spurious heuristics specific to the training set.

Mutual Information Estimation. Mutual information estimation is a well-known difficult problem, especially when the feature vectors are in a high dimensional space ([Chow and Huang, 2005](#); [Peng et al., 2005](#)). There are many traditional ways to estimate MI, such as the well-known histogram approach ([Steuer et al., 2002](#); [Paninski, 2003](#)), density estimations using a kernel ([Moon et al., 1995](#)), and nearest-neighbor distance ([Kraskov et al., 2004](#)). [Belghazi et al. \(2018\)](#) was recently proposed as a way to estimate MI using neural networks, which showed marked improvement over previous methods for feature vec-

tors in high-dimensional space.

6 Limitations and Future Work

In this paper we propose a general information-theoretic probe method, which is capable of probing for linguistic graph structures and avoids the randomness of training a model. In the experiments, we use our probe method to show the extent to which syntax trees and semantic graphs are encoded in pretrained BERT models. Further, we perform a simple perturbation analysis to show that with small modifications, the probe can also be used to probe for specific linguistic sub-structures. There are some limitations of our probe. First, a graph embedding is used, and some structure information could be lost in this process. We provide simple ways to test this. Second, training a MI estimation model is difficult. Future work can consider building on our framework by exploring better graph embedding and MI estimation techniques.

Broader Impact and Discussion of Ethics

In recent years, deep learning approaches have been the main models for state-of-the-art systems in natural language processing. However, understanding the decision making in these systems has been hard, and has challenges when these systems are used in human contexts. Probing helps us gain interpretability and hence is useful in deploying these black-box models. Our work introduces a simple and general way for understanding how linguistic properties represented as graph structures are encoded in large pretrained language models which are being applied to a wide range of structures in NLP. The methodology and probing results can be helpful to the development of future NLP models.

While our model is not tuned for any specific real-world application domain, our methods could be used in sensitive contexts such as legal or healthcare settings, and it is essential that any work using our probe method undertake extensive quality-assurance and robustness testing before using it in their setting. The datasets used in our work do not contain any sensitive information to the best of our knowledge.

Acknowledgments

We would like to thank reviewers for the constructive comments and providing suggestions for future work. This work was funded by SNF project #201009.

References

- Omri Abend and Ari Rappoport. 2013. **Universal conceptual cognitive annotation (UCCA)**. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, ACL 2013, 4-9 August 2013, Sofia, Bulgaria, Volume 1: Long Papers*, pages 228–238. The Association for Computer Linguistics.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. **Abstract meaning representation for sembanking**. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse, LAW-ID@ACL 2013, August 8-9, 2013, Sofia, Bulgaria*, pages 178–186. The Association for Computer Linguistics.
- Arindam Banerjee. 2006. **On bayesian bounds**. In *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, volume 148 of *ACM International Conference Proceeding Series*, pages 81–88. ACM.
- Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, R. Devon Hjelm, and Aaron C. Courville. 2018. **Mutual information neural estimation**. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 530–539. PMLR.
- Yonatan Belinkov, Lluís Màrquez, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James R. Glass. 2017. **Evaluating layers of representation in neural machine translation on part-of-speech and semantic tagging tasks**. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing, IJCNLP 2017, Taipei, Taiwan, November 27 - December 1, 2017 - Volume 1: Long Papers*, pages 1–10. Asian Federation of Natural Language Processing.
- Johan Bos. 2013. **The groningen meaning bank**. In *Proceedings of the Joint Symposium on Semantic Processing, Textual Inference and Structures in Corpora, JSSP 2013, Trento, Italy, November 20-22, 2013*, page 2. Association for Computational Linguistics.
- Hongyun Cai, Vincent W. Zheng, and Kevin Chen-Chuan Chang. 2018. **A comprehensive survey of graph embedding: Problems, techniques, and applications**. *IEEE Trans. Knowl. Data Eng.*, 30(9):1616–1637.
- Noam Chomsky. 1957. Logical structures in language. *American Documentation (pre-1986)*, 8(4):284.
- Tommy W. S. Chow and Di Huang. 2005. **Estimating optimal feature subsets using efficient estimation of high-dimensional mutual information**. *IEEE Trans. Neural Networks*, 16(1):213–224.
- Alexis Conneau, Germán Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018. **What you can cram into a single $\{\&!#\}$ vector: Probing sentence embeddings for linguistic properties**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2126–2136. Association for Computational Linguistics.
- Valence Cristiani, Maxime Lecomte, and Philippe Maurine. 2020. **Leakage assessment through neural estimation of the mutual information**. In *Applied Cryptography and Network Security Workshops - ACNS 2020 Satellite Workshops, AIBlock, AIHWS, AIoTS, Cloud S&P, SCI, SecMT, and SiMLA, Rome, Italy, October 19-22, 2020, Proceedings*, volume 12418 of *Lecture Notes in Computer Science*, pages 144–162. Springer.
- Marie-Catherine De Marneffe, Bill MacCartney, Christopher D Manning, et al. 2006. **Generating typed dependency parses from phrase structure parses**. In *Lrec*, volume 6, pages 449–454.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Francisco Escolano, Edwin R. Hancock, Miguel Angel Lozano, and Manuel Curado. 2017. **The mutual information between graphs**. *Pattern Recognit. Lett.*, 87:12–19.
- Tom Fawcett. 2006. **An introduction to ROC analysis**. *Pattern Recognit. Lett.*, 27(8):861–874.
- Shuyang Gao, Greg Ver Steeg, and Aram Galstyan. 2015. **Efficient estimation of mutual information for strongly dependent variables**. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2015, San Diego, California, USA, May 9-12, 2015*, volume 38 of *JMLR Workshop and Conference Proceedings*. JMLR.org.
- Jan Hajic, Eva Hajicová, Jarmila Panevová, Petr Sgall, Ondrej Bojar, Silvie Cinková, Eva Fucíková, Marie Mikulová, Petr Pajas, Jan Popelka, Jirí Semecký, Jana Sindlerová, Jan Stepánek, Josef Toman, Zdenka Uresová, and Zdenek Zabokrtský. 2012. **Announcing prague czech-english dependency treebank 2.0**. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC 2012, Istanbul, Turkey, May 23-25, 2012*, pages

- 3153–3160. European Language Resources Association (ELRA).
- William L. Hamilton, Zhitao Ying, and Jure Leskovec. 2017. [Inductive representation learning on large graphs](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 1024–1034.
- John Hewitt and Percy Liang. 2019. [Designing and interpreting probes with control tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 2733–2743. Association for Computational Linguistics.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4129–4138. Association for Computational Linguistics.
- Julia Hockenmaier and Mark Steedman. 2007. [Ccg-bank: A corpus of CCG derivations and dependency structures extracted from the penn treebank](#). *Comput. Linguistics*, 33(3):355–396.
- Xuedong Huang, Fileno A. Alleva, Hsiao-Wuen Hon, Mei-Yuh Hwang, Kai-Fu Lee, and Ronald Rosenfeld. 1993. [The SPHINX-II speech recognition system: an overview](#). *Comput. Speech Lang.*, 7(2):137–148.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What does BERT learn about the structure of language?](#) In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3651–3657. Association for Computational Linguistics.
- Alexander Koller, Stephan Open, and Weiwei Sun. 2019. [Graph-based meaning representations: Design and processing](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics: Tutorial Abstracts, ACL 2019, Florence, Italy, July 28, 2019, Volume 4: Tutorial Abstracts*, pages 6–11. Association for Computational Linguistics.
- Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. 2004. Estimating mutual information. *Physical review E*, 69(6):066138.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019. [Linguistic knowledge and transferability of contextual representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 1073–1094. Association for Computational Linguistics.
- Charles Lovering, Rohan Jha, Tal Linzen, and Ellie Pavlick. 2021. [Predicting inductive biases of pre-trained models](#). In *International Conference on Learning Representations*.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of english: The penn treebank. *Comput. Linguistics*, 19(2):313–330.
- Marie-Catherine de Marneffe, Bill MacCartney, and Christopher D. Manning. 2006. [Generating typed dependency parses from phrase structure parses](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation, LREC 2006, Genoa, Italy, May 22-28, 2006*, pages 449–454. European Language Resources Association (ELRA).
- Tomás Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 3111–3119.
- Young-Il Moon, Balaji Rajagopalan, and Upmanu Lall. 1995. Estimation of mutual information using kernel density estimators. *Physical Review E*, 52(3):2318.
- Liam Paninski. 2003. [Estimation of entropy and mutual information](#). *Neural Comput.*, 15(6):1191–1253.
- Hanchuan Peng, Fuhui Long, and Chris H. Q. Ding. 2005. [Feature selection based on mutual information: Criteria of max-dependency, max-relevance, and min-redundancy](#). *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(8):1226–1238.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543. ACL.
- Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. [Deepwalk: online learning of social representations](#). In *The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014*, pages 701–710. ACM.

- Tiago Pimentel, Naomi Saphra, Adina Williams, and Ryan Cotterell. 2020a. [Pareto probing: Trading off accuracy for complexity](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 3138–3153. Association for Computational Linguistics.
- Tiago Pimentel, Josef Valvoda, Rowan Hall Maudslay, Ran Zmigrod, Adina Williams, and Ryan Cotterell. 2020b. [Information-theoretic probing for linguistic structure](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 4609–4622. Association for Computational Linguistics.
- Nima Pourdamghani, Yang Gao, Ulf Hermjakob, and Kevin Knight. 2014. [Aligning english strings with abstract meaning representation graphs](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIG-DAT, a Special Interest Group of the ACL*, pages 425–429. ACL.
- Benjamin Recht and Christopher Ré. 2012. [Beneath the valley of the noncommutative arithmetic-geometric mean inequality: conjectures, case-studies, and consequences](#). *arXiv preprint arXiv:1202.4184*.
- Machel Reid, Edison Marrese-Taylor, and Yutaka Matsuo. 2020. [VCDM: leveraging variational bi-encoding and deep contextualized word representations for improved definition modeling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 6331–6344. Association for Computational Linguistics.
- Emily Reif, Ann Yuan, Martin Wattenberg, Fernanda B. Viégas, Andy Coenen, Adam Pearce, and Been Kim. 2019. [Visualizing and measuring the geometry of BERT](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 8592–8600.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A primer in bertology: What we know about how BERT works](#). *Trans. Assoc. Comput. Linguistics*, 8:842–866.
- Brian C Ross. 2014. [Mutual information between discrete and continuous data sets](#). *PloS one*, 9(2):e87357.
- Ralph E. Steuer, Jürgen Kurths, Carsten O. Daub, Janko Weise, and Joachim Selbig. 2002. [The mutual information: Detecting and evaluating dependencies between variables](#). In *Proceedings of the European Conference on Computational Biology (ECCB 2002), October 6-9, 2002, Saarbrücken, Germany*, pages 231–240.
- Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015. [LINE: large-scale information network embedding](#). In *Proceedings of the 24th International Conference on World Wide Web, WWW 2015, Florence, Italy, May 18-22, 2015*, pages 1067–1077. ACM.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019a. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4593–4601. Association for Computational Linguistics.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R. Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R. Bowman, Dipanjan Das, and Ellie Pavlick. 2019b. [What do you learn from context? probing for sentence structure in contextualized word representations](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Elena Voita and Ivan Titov. 2020. [Information-theoretic probing with minimum description length](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 183–196. Association for Computational Linguistics.
- Zhaofeng Wu, Hao Peng, and Noah A. Smith. 2021. [Infusing finetuning with semantic dependencies](#). *Trans. Assoc. Comput. Linguistics*, 9:226–242.

A Details of Graph Embedding

In this section, we present the technique details of the graph embedding approach, as well as parameters. Given a graph such as syntax tree and semantic graph, we first run random walk algorithm on it to sample walk paths. The random walk strategy is simple, each time a neighbor of current node is selected from its neighbor set based on uniform distribution. For each node, the length of random walk path is 10. And the repeat time is 100. In general, each node has 100 different walk paths with length 10. Then we put those paths into Word2vec model (Mikolov et al., 2013), with window size equal to 2, since we only want graph embeddings to capture the one-hop neighborhood relationships. The hidden states, in other words graph embeddings are vectors with 128 dimensions.

B Details of MI Estimation

In this section we present technique details about MI estimation, such as the neural network model design and parameters. There are two terms in the objective function 5: one is about joint distribution

$$\mathbb{E}_{\mathbb{P}_{\mathcal{X}\mathcal{Z}}^{(n)}}[T_\theta]$$

and another is about marginal distribution

$$\log(\mathbb{E}_{\mathbb{P}_{\mathcal{X}}^{(n)} \otimes \mathbb{P}_{\mathcal{Z}}^{(n)}}[e^{T_\theta}]).$$

For the joint distribution part, we concatenate the graph embeddings \mathcal{Z} and word representations \mathcal{X} first, and then put them into our designed neural network to compute a scalar. Then the average value of the scalar is computed. For the marginal part, we randomly shuffle the representations \mathcal{X} . After random shuffle, there is no dependency between \mathcal{X} and \mathcal{Z} anymore. Then, we put the concatenation of the shuffled representations and graph embeddings into the neural network to get another scalar. We take the exponential value of that scalar and take the average value for the whole dataset.

As aforementioned, the selection of model size is a tradeoff. In our experiments, we design an MLP model with two layers for MI estimation. The first layer is linear without nonlinear activation function, to encode graph embeddings and word representations into same space, with 64 dimensions. Then we concatenate those two hidden states and put them through a nonlinear layer to get a scalar. For example, we have one sentence with 10 words.

Assume we can get graph embeddings with size $10 * 128$, and word representations of BERT with size $10 * 768$. Then we use a linear function to map those two vectors into hidden space, say with 32 dimensions. Then we concatenate those two embeddings as a $10 * 64$ matrix, and use one extra linear function with nonlinear activation functions to map it as $10 * 1$ matrix. Then we can get the compression lemma lower bound as the mean value of the $10 * 1$ matrix, which is the mutual information estimation that we want.

The loss is defined directly as the minus value of objective function. With stochastic gradient decent, we can maximize the lower bound to get the estimation. About the mini-batch, since the document contains many sentences, we select one sentence as one min-batch to optimize the neural network.

The reason why we treat one sentence as one minibatch is that we get word representations of BERT and graph embeddings in that way. One sentence has a complete syntax tree structure, and getting word representations with one sentence in BERT can make attention computed within the sentence. Another reason is that using two sentences as input may exceed the maximum BERT input size: 512 tokens sometimes. However, if we use mini-batch to estimate the mutual information, it brings errors. The reason is that if we want to estimate the mutual information between \mathcal{X} and \mathcal{Z} , the expectation should be all the data points that we know. But here we use minibatch to calculate the expectation for one batch only. To alleviate this error, as introduced in Belghazi et al. (2018), we select small learning rate to keep the error small.

C Details of MLP Models

This section introduces how to use MLP to do link prediction task, as well as the details about MLPs. For one sentence, given its graph embeddings, we simply use MLP to calculate a score for all node pairs, and then compare with the ground-truth graph with the predicted distribution vector. AUC score is computed based on the distribution vector and ground-truth vector. Note that since the graph is very sparse, it makes the task very difficult. Generally, the task can be regarded as a binary classification task with an extremely unbalanced data distribution.

For the details, linear MLP simply predicts the graph by the concatenation of two input vectors to decide whether there is an edge between them.

For MLPs with hidden layers, The concatenated vectors are through non-linear layers first, then the final output layer is linear. The dimension of all hidden states is 128. And the learning rate is 10^{-4} .

D Reliability of MI Estimation

MI estimation for features in high-dimensional space is difficult. To prove that our estimated MI values are quite accurate, we test the MI estimation method (Belghazi et al., 2018) on sets of graph embeddings with different levels of dependencies. To have that, we add noise on graph embeddings as \mathcal{Z}' . \mathcal{Z} and \mathcal{Z}' can have different dependencies based on the added noise rate. Noise vectors are sampled from a standard Gaussian independently. We test the estimation for various levels of noise, from original graph embeddings \mathcal{Z} to the condition that 100% signals are noise σ . For example, 40% means that for each graph embedding $\mathcal{Z}' = 60\% \times \mathcal{Z} + 40\% \times \sigma$. Then, we calculate $\hat{I}(\mathcal{Z}'; \mathcal{Z})$ to see whether the values is small with large noise added.

Table 3: MI value with different level of noise

MI percentage (%)	Structure	
	Syntactic tree	Semantic graph
Noise ratio (%)		
0	100.00	100.00
10	91.84	80.40
20	74.10	63.29
30	62.23	51.88
40	49.53	32.70
50	38.61	20.44
60	29.46	9.36
70	22.09	3.74
80	12.25	0.94
90	2.07	0.02
100	-0.03	-0.06

Table 3 presents the exact values. To make it more readable, we report the MI percentage of $\hat{I}(\mathcal{Z}'; \mathcal{Z})/\hat{I}(\mathcal{Z}; \mathcal{Z})$. From the results, we find that for the two structures, results are not very similar. But the general tendencies are consistent, where less dependencies caused by larger noise have smaller MI values. Note that when the noise rate is 100%, $\hat{I}(\mathcal{Z}', \mathcal{Z})$ degenerates to the lower bound $\hat{I}(\mathcal{R}, \mathcal{Z})$. As mentioned before, the absolute value of it represents the gap between the estimated MI and ground-truth MI, which is very small (less than $10^{-3} \times \hat{I}(\mathcal{Z}; \mathcal{Z})$). It also proves that our MI estimations are reliable.

E Statistics of Penn Treebank and AMR Bank

We provide the relation number and connected word number of Penn Treebank dataset and AMR Bank in this section. The word number for POS tags of Penn Treebank is also provided. The statistics are for the whole dataset. We only report de-

Table 4: Penn Treebank statistics of syntax relations

Statistic	# of Relations	% of Relations	# of Connected Words
Relation			
punct	121,395	11.60	193,648
prep	100,997	9.648	189,783
pobj	98,586	9.418	197,164
det	86,228	8.237	172,446
nn	81,381	7.774	143,248
nsubj	73,802	7.050	147,498
amod	66,381	6.341	125,119
root	43,948	4.198	87,896
dobj	43,054	4.113	85,997
aux	37,267	3.560	73,436
# others	26,1791	25.01	509,641

tails of relations ranked top 10. However, there are other 35 types of relations, which are categorized together in type # others. For the POS tagging

Table 5: Penn Treebank statistics of POS tags

Statistic	# of Words	% of Words
Tags		
NN	146,228	15.89
IN	108,434	11.78
NNP	101,427	11.02
DT	90,158	9.798
JJ	67,396	7.324
NNS	65,867	7.158
CD	40,337	4.384
RB	34,331	3.731
VBD	33,430	3.633
VB	29,001	3.152
# others	203,578	22.12

statistics, we also only present tags with word number ranked top 10. There are still 28 types of POS tags that are catergorized into one type # others.

Table 6: AMR bank statistics of relations

Statistic	# of Relations	% of Relations
Relation		
arg	409,322	58.11
general	208,287	29.57
op	67,307	9.556
quantities	13,092	1.859
others	5,216	0.7406
date	1,114	0.1582

The AMR graphs are different from syntax trees. The relations of AMR graphs can be classified into 6 groups, and each group contains many types. Specifically, the group *general* includes: “accompanier”, “age”, “beneficiary”, “concession”, “condition”, “consist”, “degree”, “destination”, “direction”, “domain”, “duration”, “example”, “extent”, “frequency”, “instrument”, “location”, “manner”, “medium”, “mod”, “name”, “part”, “path”, “polarity”, “poss”, “purpose”, “source”, “subevent”, “subset”, “time”, “topic”, “value”, “ord”, and “range”. And the *quantities* group includes “quant”, “scale”, and “unit”.

For the description, *arg* represents frame arguments, following PropBank conventions. *general* are composed of a set of general semantic relations. *op* means the relations for lists. Similarly, *quantities* are relations for quantities. And *date* contains relations for date-entities.

F Probe Specific Semantic Relations

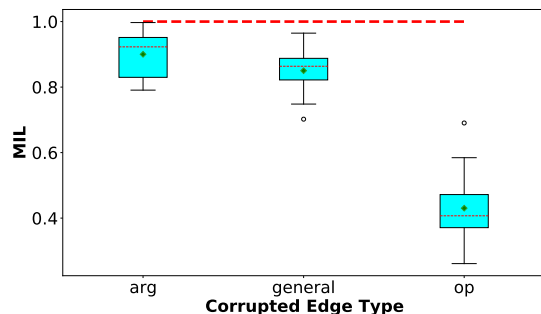


Figure 7: *MIL* scores of probing 5 types of localized semantic structure for word representations in BERT-base (output layer). The local structure is decided by the labels attached on edges in AMR graph.

For semantic structure, we run the perturbation experiment in a similar way. Different from syntax trees, the relations of AMR graphs can be grouped into 6 types. And number distribution is not very even. Thus, we corrupt the graph embeddings with 50% noise. Other settings are similar to that of syntactic structures. From the results, we can find that BERT encodes more structure information about *arg* and *general* relations. While for the *op* relations, which represents the relations for lists, are not well encoded.

G Probing Localized Information with Accuracy

Similar to our localized probing experiments, we add perturbations on word representations. Specifically, we corrupt word representations equally and with same number for each relation type. Then we train an MLP with 5 hidden layers to predict entire structures with the corrupted word representations. AUC scores of all relation types are calculated. As in the *Worm’s Eye*, we only report 5 types of relations for syntactic and 3 types for semantic structures. Results are shown in Figure 8(a) and Figure 8(b). We can find that the probe accuracies are very unusual. First, corrupt one type of relations, the accuracies for other types of relations change significantly. Besides, the MLP is trained with corrupted relations e.g., *nsubj* while predicts *prep* with worst AUC score. The results also prove the point that prediction decisions are not based purely on the structure but rather on spurious heuristics (Lovering et al., 2021).

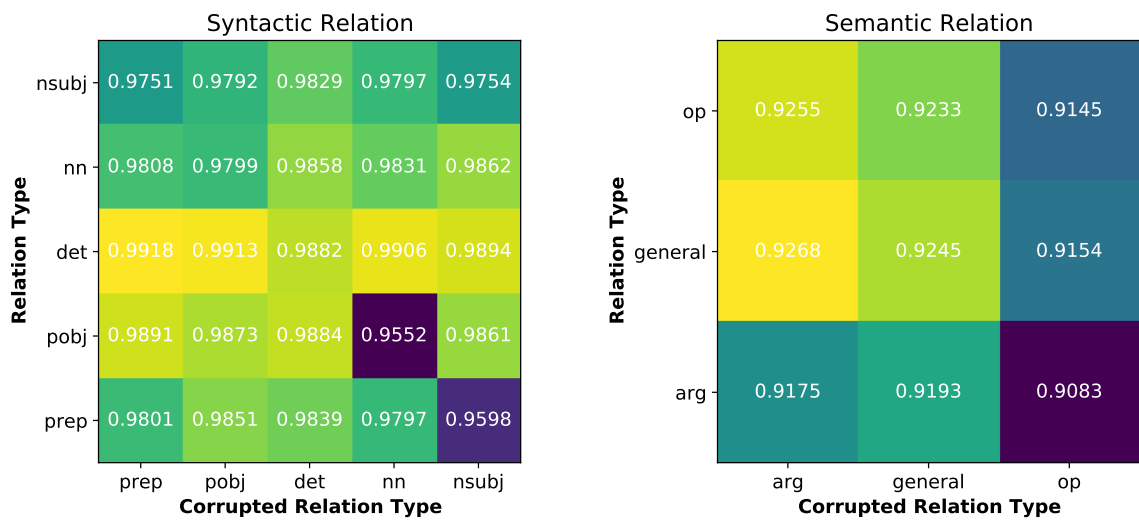


Figure 8: Probe specific syntactic and semantic relations