

Picking BERT’s Brain: Probing for Linguistic Dependencies in Contextualized Embeddings Using Representational Similarity Analysis

Michael A. Lepori

Department of Computer Science
Johns Hopkins University
mlepori19@gmail.com

R. Thomas McCoy

Department of Cognitive Science
Johns Hopkins University
tom.mccoy@jhu.edu

Abstract

As the name implies, contextualized representations of language are typically motivated by their ability to encode context. Which aspects of context are captured by such representations? We introduce an approach to address this question using Representational Similarity Analysis (RSA). As case studies, we investigate the degree to which a verb embedding encodes the verb’s subject, a pronoun embedding encodes the pronoun’s antecedent, and a full-sentence representation encodes the sentence’s head word (as determined by a dependency parse). In all cases, we show that BERT’s contextualized embeddings reflect the linguistic dependency being studied, and that BERT encodes these dependencies to a greater degree than it encodes less linguistically-salient controls. These results demonstrate the ability of our approach to adjudicate between hypotheses about which aspects of context are encoded in representations of language.

1 Introduction

Contextualized word embeddings (Devlin et al., 2019; Peters et al., 2018), which are vector representations of words in context, enable neural models of language to achieve dramatic performance improvements over models whose word embeddings do not have access to context (Pennington et al., 2014; Mikolov et al., 2013). The most obvious explanation for the success of these models is that contextualized word embeddings can incorporate contextual information, whereas other embeddings cannot. Contextual information provides clues about the semantic and syntactic roles that a word plays in a sentence. For example, a verb might be understood differently when placed in different contexts: In sentence (1a), the verb *charged* means “ran towards,” where in sentence (1b) the verb means “formally accused”.

- (1) a. The bull **charged** the man.
- b. The prosecutor **charged** the man.

There are many aspects of context that could conceivably be captured in contextualized embeddings, from linear context (e.g., what word precedes this one?) to syntactic context (e.g., what is the parent of this word in a dependency tree?). In this work, we investigate which aspects of context are captured in these embeddings. We do so by studying the embeddings’ representational geometry, which is the spatial relationship between representations of stimuli. In our case, the stimuli are contextualized embeddings of words.

We study this geometry by applying representational similarity analysis (Kriegeskorte et al., 2008, RSA) to the contextualized word embeddings given by BERT (Devlin et al., 2019). This allows us to ask fine-grained questions about which context words are encoded, and to what degree.¹

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

¹For brevity, we say that a model M encodes an aspect of context if the representational similarity between M ’s embeddings and a hypothesis model encoding that aspect of context is greater than the representational similarity between M ’s embeddings and a null hypothesis model. We also say that M encodes a set of context words A more than another set B if the representational similarity between M ’s embeddings and the hypothesis model for A is greater than the representational similarity between M ’s embeddings and the hypothesis model for B .

We find that the representational geometries of BERT’s word and sentence embeddings reflect several linguistic dependencies. In particular, we find that BERT’s embeddings of verbs encode the subject of those verbs more than they encode nouns that are not arguments of those verbs. We also find that the contextualized embeddings of pronouns encode the pronouns’ antecedents more than they encode other nouns. Finally, we find that BERT’s sentence embeddings encode the main verbs of sentences more than any other content words. This is consistent with the standard assumption in dependency parsing that the main verb of a sentence is the sentence’s head. These results demonstrate the ability of our approach to illuminate which aspects of context are encoded in contextualized embeddings.²

2 Background: Representational Geometry

A representational geometry is the spatial arrangement of a set of vector representations. Representational geometries are typically formed by taking the pairwise dissimilarities between those representations. For example, the representational geometry of the BERT embeddings for the subjects of a set of sentences is given by the pairwise dissimilarities between the contextualized word embeddings corresponding to those subjects. When the set of representations is poorly understood, one can gain insight into it by comparing it to a set of representations that is well-understood. If the two sets have similar representational geometries, then one can infer that the sets encode similar information. Representational similarity analysis (RSA) allows for comparisons between two different representational geometries. In cognitive neuroscience, this technique is used to analyze distributed activity patterns in the brain (see Kriegeskorte and Kievit (2013) for a review); we use it to analyze the representations of artificial neural networks.

We compare contextualized BERT embeddings to representations that we construct to instantiate specific linguistic hypotheses, which we refer to as *hypothesis models*. These models represent specific linguistic information and abstract away from all other information (Kriegeskorte et al., 2008). We compare the similarity among the representations of each hypothesis model to the similarity among BERT’s embeddings; if the representational geometry of BERT’s embeddings is better matched by the representational geometry of hypothesis model A than by that of hypothesis model B , we conclude that the hypothesis instantiated by A is a better description of the content of BERT’s embeddings. Each hypothesis model represents a specific type of context word (e.g., a verb’s subject), while ignoring other context words. Such a hypothesis model instantiates the hypothesis that this aspect of context is the only aspect represented in BERT’s embeddings. Of course, such extreme hypotheses are almost certainly wrong; our goal is not to find a perfect hypothesis but rather to find which of two hypotheses is closer to the truth.

3 Probing Contextualized Embeddings

Our approach works as follows: First, we create a corpus C of N sentences that contain the syntactic structures we wish to study (Figure 1, Step 1). Next, we define a reference model M_{Ref} , which consists of the representations that are being investigated (e.g., the embeddings of the main verbs from every sentence in C). Then, we define two hypothesis models, M_{Hyp1} and M_{Hyp2} . These hypothesis models instantiate hypotheses about the representational geometry of the contextualized word embeddings (Figure 1, Step 2). We then draw a sample c of n sentences from our corpus, and calculate the $n \times n$ representational geometries g of each model by applying a dissimilarity metric D to the relevant word embeddings from this sample (Figure 1, Step 3). $D(M, c)$ finds the dissimilarity between the representations generated by model M for each pair of sentences in the sample c .

$$g_{Ref} = D(M_{Ref}, c) \tag{1}$$

$$g_{Hyp1} = D(M_{Hyp1}, c) \tag{2}$$

$$g_{Hyp2} = D(M_{Hyp2}, c) \tag{3}$$

Finally, we calculate the similarity s between the representational geometries of our hypothesis models and our reference model using a similarity metric sim . Because the g matrices are symmetric, sim only

²Code and data can be found at https://github.com/mleporil/Picking_BERTs_Brain.

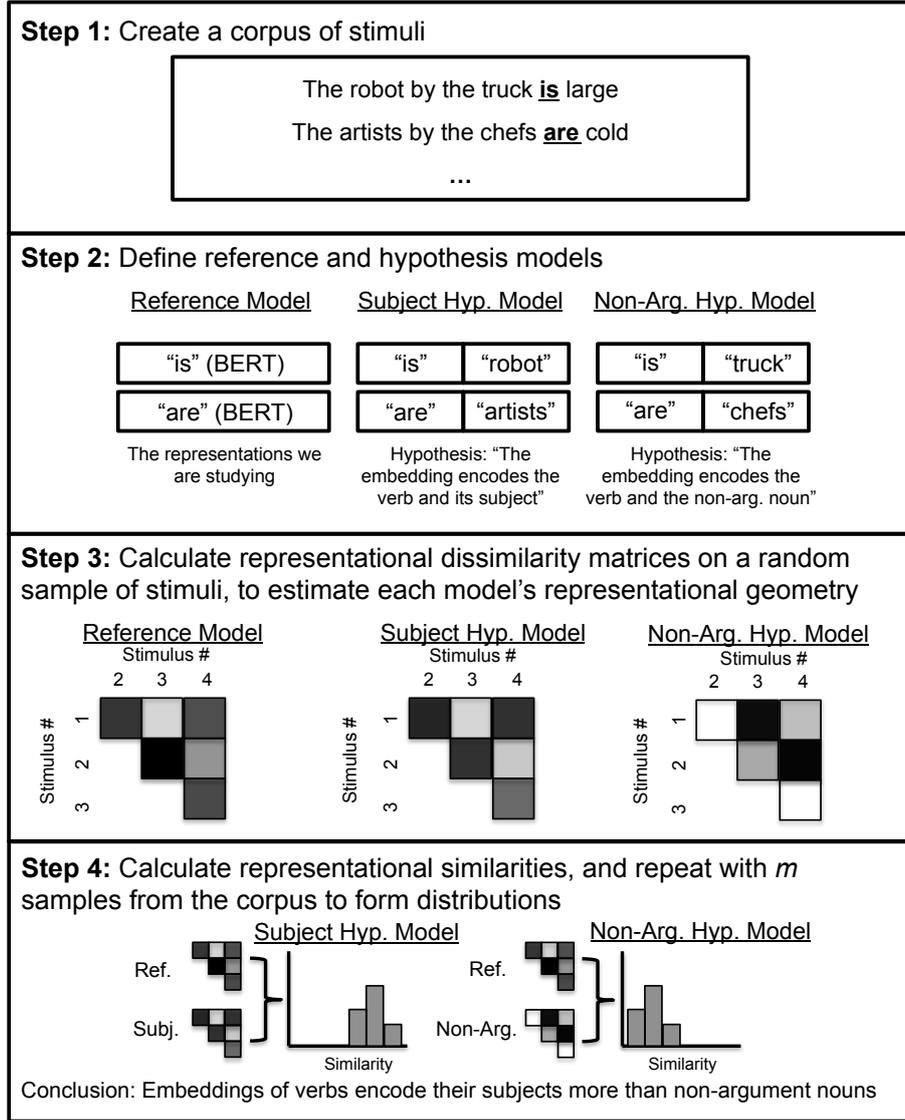


Figure 1: A summary of our approach. From this example, we would conclude that the contextualized embeddings of verbs encode their subjects more than they encode non-argument nouns.

operates on the upper triangle of each g matrix.

$$s_{Hyp1} = sim(g_{Ref}, g_{Hyp1}) \quad (4)$$

$$s_{Hyp2} = sim(g_{Ref}, g_{Hyp2}) \quad (5)$$

We then repeat the process on m samples from our corpus in order to create two m -length vectors of representational similarities, S_{Hyp1} and S_{Hyp2} (Figure 1, Step 4). Finally, we apply a nonparametric sign test to the difference of these vectors, $S_{Hyp1} - S_{Hyp2}$, to test whether there is a consistent difference between measurements of s_{Hyp1} and s_{Hyp2} . In the following section, we will walk through this approach with a concrete example.

4 Experiment 1: Subject-Sensitivity of Verb Embeddings

Most verbs in English require arguments, such as subjects and direct objects. Here we focus on subjects, which are required by nearly all verbs. Specifically, we consider whether the contextualized word embedding of the main verb encodes the subject of a sentence to a greater degree than it encodes other nouns in the sentence that are not arguments of the verb. In particular, we study the verbs *is* and *are*.

Corpus Generation: We generate two corpora using a probabilistic context free grammar (PCFG): the Prepositional Phrase Corpus, containing sentences of the form (2a), and the Relative Clause Corpus, containing sentences of the form (2b).

- (2) a. The [NOUN₁] [PREPOSITION] the [NOUN₂] [is/are] [ADJECTIVE].
- b. The [NOUN₁] that [VERB] the [NOUN₂] [is/are] [ADJECTIVE].

The following is an example sentence from the Prepositional Phrase Corpus:

- (3) The doctors by the cars are ugly.

Importantly, any of the nouns in our grammar can occupy the [NOUN₁] position or the [NOUN₂] position. For example, both sentence (3) and the following sentence appear in the Prepositional Phrase Corpus:

- (4) The cars by the doctors are ugly.

Additionally, the two nouns in any given sentence are of the same syntactic number (singular vs. plural) in order to eliminate cues from number features. We generated 2,000 sentences for each corpus and then removed repeated sentences and sentences for which NOUN₁ and NOUN₂ are identical, leaving 1,863 sentences in the Prepositional Phrase corpus and 1,869 in the Relative Clause Corpus. Finally, we chose the vocabulary such that every sentence is semantically plausible.

Representational Models: In this experiment, we analyze the embeddings of verbs. Thus, our *reference model* is the set of contextualized word embeddings of the verb for each sentence, taken from the last layer of BERT (Devlin et al., 2019).³ This study uses the BERT-base-uncased variant of BERT, which creates 768-dimensional embeddings. We compare the representational geometry of these embeddings to that of three hypothesis models. These models instantiate the following hypotheses:

Subject Hypothesis: The representational geometry of the contextualized embeddings of verbs reflects information about their subjects.

Non-Argument Hypothesis: The representational geometry of the contextualized embeddings of verbs reflects information about nouns that are not arguments of the verb. These non-argument nouns are either the objects of prepositions (in the Prepositional Phrase Corpus) or are the nouns within relative clauses (in the Relative Clause Corpus).

Null Hypothesis: The representational geometry of the contextualized embeddings of verbs does not reflect information about subjects or non-argument nouns.

We create the Subject Hypothesis Model by taking the 300-dimensional (noncontextualized) GloVe embedding (Pennington et al., 2014) of the verb and concatenating it with the 300-dimensional GloVe embedding of the subject noun. We use GloVe embeddings that are pretrained on the Wikipedia 2014 + Gigaword 5 corpus. The Non-Argument Hypothesis Model concatenates the GloVe embeddings of the verb and non-argument noun. The Null Hypothesis Model concatenates the GloVe embedding of the verb and the embedding of a random noun from our grammar that does not appear in the sentence.⁴ We assume that, if BERT’s representations of the verbs in our corpora encode the verbs’ subjects more saliently than non-argument nouns, then the representational geometry of BERT’s representations will be more similar to the representational geometry of the hypothesis model containing only the subject and verb than to the representational geometry of the hypothesis model containing only the non-argument noun and verb.

³We use representations from the last layer because Ethayarajh (2019) found that BERT representations are more context-specific in higher layers than lower layers. Because we are looking for the effects of contextualization, we study the layer that is most contextualized.

⁴This means that the Null Hypothesis Model represents information about the verb and *some* noun. If the contextualized word embedding only encoded the verb and the fact that it refers to a noun (without encoding anything specific about that noun), then we would expect the Null Hypothesis Model to achieve approximately equal representational similarity as the other Hypothesis Models.

Note that GloVe embeddings play no role in BERT, and are likely to be very different from BERT embeddings. One clear difference is that they are of different dimensionalities. However, such differences are immaterial in applying RSA, which illustrates one particular strength of this approach: All that is required is for the hypothesis models to have the hypothesized *representational geometry* (which these concatenated GloVe embeddings do), allowing us to abstract away from superficial differences between the models.

This method of creating hypothesis models has two advantages over other plausible approaches. First, it is very likely that BERT’s embedding of the verb will be strongly influenced by the identity of the verb (*is* vs. *are*). Because we are just interested in the effect of *context* on the representation, we want to control for the effect of the verb’s identity on the representational geometries of our hypothesis models. The approach we have chosen allows us to control for this factor by including each verb’s GloVe embedding in each hypothesis model. An additional advantage of our approach is that the use of GloVe embeddings allows our similarity measures to be more granular than other plausible approaches (such as one-hot encodings of the verb and relevant noun) would allow.

Finally, we note that these hypothesis models will likely distort the effect of context compared to the BERT verb representation. By construction, 50% of each vector in all three hypothesis models consists of a noun and 50% of the vector consists of the verb itself. It would be surprising to learn that the contextualized embedding encoded only this context information, and in exactly this proportion. Thus, we do not expect the absolute fit of the hypothesis models to be very good. However, each hypothesis model makes the same exaggeration, so comparisons between the models are valid even if each model has a poor absolute fit. Thus, we consider only the *differences* in representational similarity between hypothesis models, which indicates which hypothesis model provides a better fit to the reference model.

Pitting these models against each other allows us to determine which aspect of context is encoded to a greater degree. If syntactic structure dominates BERT’s representations, then we would expect the subject to be encoded to a greater degree than the non-argument noun. However, if BERT has learned to rely on surface heuristics based on linear distance, then we would expect the opposite. In addition to comparing the hypothesis models to each other, we can also compare each hypothesis model to the null model to determine whether the two nouns that appear in the sentence influence the representational geometry of the BERT embedding more than a random noun.

Applying RSA: We now perform RSA on our models, as described in Section 3. We specify the sample size $n = 200$, the number of samples $m = 100$, the dissimilarity metric $D = 1 - Spearman's \rho$ and similarity metric $sim = Spearman's \rho$.

Zhelezniak et al. (2019) show that Spearman’s ρ is the most appropriate measurement of (dis)similarity for GloVe embeddings, as these embeddings violate the assumptions underlying other common metrics. We perform a similar analysis to show that it is also the most appropriate dissimilarity measurement for BERT embeddings (See Appendix B). Finally, we compare representational geometries using Spearman’s ρ because it is robust and makes few assumptions (Diedrichsen and Kriegeskorte, 2017).

Results: We summarize our results in Table 1a. For both corpora, both the Subject and Non-Argument Hypothesis Models exhibit significantly greater representational similarity to the Reference Model than the Null Hypothesis Model does. This shows that the contextualized representations of verbs encode both of the noun categories in the sentence. Furthermore, these two syntactic categories are not encoded to the same degree, as the Subject Model reliably exhibits greater representational similarity to the Reference Model than the Non-Argument Model does for both corpora (Figure 2). From this, we can infer that BERT’s embeddings of these verbs encode the verbs’ subjects to a greater degree than they encode the non-argument nouns, despite the fact that the non-argument nouns exhibit much less surface-level distance from the verb. This result corroborates other evidence—from behavioral evaluations (Goldberg, 2019), analyses of attention heads (Clark et al., 2019), and probing classifier tests (Klafka and Ettinger, 2020)—that BERT is sensitive to subject-verb dependencies.

	Subject	Non-Argument	Null
Prepositional Phrases	.524	.478	.440
Relative Clauses	.667	.651	.600

(a) Representational similarities from Experiment 1.

	Antecedent	Non-Antecedent	Null
Reflexive	.717	.696	.663
Pronominal	.726	.715	.691

(b) Representational similarities from Experiment 2.

Table 1: Representational similarities from Experiments 1 and 2. All differences across columns are statistically significant ($p < .001$). We do not compare across rows because differences in sentence structure make it difficult to compare these numbers.

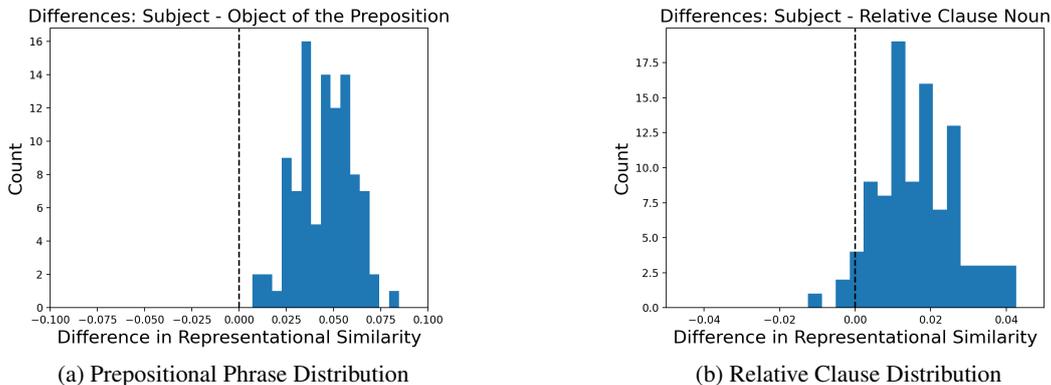


Figure 2: Distributions of differences of representational similarity for Experiment 1. Values greater than 0 denote that the reference model exhibits greater representational similarity to the Subject Hypothesis Model than the Non-Argument Hypothesis Model. For both sentence structures, the Subject Model consistently yields higher representational similarity ($p < .001$).

5 Experiment 2: Pronoun Coreference

Like the meanings of verbs, the meanings of pronouns can also be affected by context. Pronouns typically refer to some contextually-understood noun, which is called its *antecedent*. For example, in *I love New York - it's my favorite city*, *it* refers to *New York*; whereas in *The Death Star is a threat to the galaxy, so it must be destroyed*, *it* refers to *the Death Star*. Here we investigate whether the contextualized representations of pronouns encode information about the pronouns' antecedents.

We focus on two types of pronouns, *reflexives* and *pronominals*. Reflexives must refer to a locally-occurring (i.e., in the same clause) noun. Pronominals must *not* refer to a locally-occurring noun. Sentence (5a) contains a reflexive (*himself*), which refers to the noun *politician*, while sentence (5b) contains a pronominal (*him*), referring to the noun *person*. Note that, in both cases, the pronoun cannot refer to the other noun (underlined).

- (5) a. The person believes that the **politician** loves **himself**.
 b. The **person** believes that the politician loves **him**.

We study both types of pronoun separately. We thus generate one corpus containing sentences of the same form as Example (5a), and another containing sentences of the same form as Example (5b), using a probabilistic context-free grammar. We exclude any sentence where both nouns are the same. This yields 1,828 valid sentences in our reflexive corpus, and 1,826 valid sentences in our pronominal corpus.

For both corpora, we consider the same reference model: the set of BERT embeddings of the pronouns. We also consider two hypothesis models: the Antecedent Model concatenates the GloVe embedding of the pronoun with the GloVe embedding of its antecedent for every sentence. Note that the antecedent is the first noun in the sentence for the pronominal corpus, and the second noun in the sentence for the reflexive corpus (bolded in Example (5a) and Example (5b)). The Non-Antecedent Model concatenates the GloVe embedding of the pronoun with the GloVe embedding of the non-antecedent noun. The non-

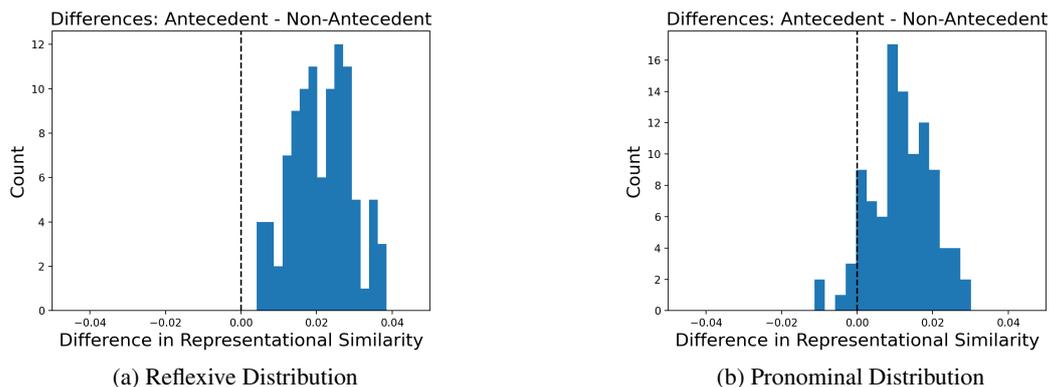


Figure 3: Distributions of differences of representational similarity for Experiment 2. Values greater than 0 denote that the reference model exhibits greater representational similarity to the Antecedent Hypothesis Model than the Non-Antecedent Hypothesis Model. For both sentence structures, the Antecedent Model consistently yields higher representational similarity ($p < .001$).

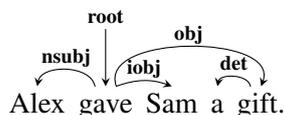


Figure 4: Example dependency parse

antecedent noun is the first noun in the sentence for the reflexive corpus, and the second noun in the sentence for the pronominal corpus (underlined in Example (5a) and Example (5b)). We also include a Null Hypothesis Model, which concatenates the GloVe embedding of the pronoun with the embedding of a random noun from our grammar that does not appear in the sentence. These hypothesis models allow us to test the hypothesis that the BERT embeddings of pronouns encode their antecedents more than any other noun in the sentence.

Results: We find that, for both corpora, the representational geometry of the set of BERT embeddings of the pronouns is significantly more similar to the hypothesis model that represents antecedent nouns than the hypothesis model that represents non-antecedent nouns (see Table 1b, as well as Figure 3). This experiment controls for the absolute and relative linear positions of the words, as the linear position of the antecedent and the non-antecedent swap when the pronoun is a reflexive as opposed to a pronominal. Thus, the representational geometry of BERT pronoun embeddings is also sensitive to syntactic dependencies, as opposed to strictly surface-level cues. BERT’s sensitivity to the relationship between pronouns and their antecedents—which was also observed by Clark et al. (2019) through analysis of BERT’s attention heads—may explain BERT’s strong performance on coreference resolution, a task that relies on identifying pronoun-antecedent relationships (Joshi et al., 2019).

6 Experiment 3: Heads of Sentences

Our approach can be applied to any type of representation that incorporates information from multiple words. So far we have applied it to contextualized word representations, but it can also be applied to full-sentence representations; here we give an example of such a usage inspired by dependency parsing.

In standard approaches to dependency parsing (de Marneffe et al., 2006; Nivre et al., 2016), the main verb of a sentence acts as the head of the entire sentence, as in Figure 4. We study whether verbs exhibit the same primacy in BERT embeddings as they do in dependency parses. To do so, we use the embedding of the [CLS] token as an embedding of the full sentence, as is standard when a sentence embedding is required from BERT. We create four corpora, each containing a different type of sentence structure. These corpora are denoted the Intransitive Corpus, Intransitive + Adjective Corpus, Transitive Corpus,

Corpus	Example
Intransitive	The painter swims.
Intransitive + Adjective	The happy politician talks.
Transitive	The person moves a lamp.
Transitive + Adjective	A scary lawyer likes the red chair.

Table 2: Corpora used in Experiment 3.

and Transitive + Adjective Corpus, respectively. An example from each corpus is shown in Table 2.

For each corpus, we use the set of BERT embeddings for each sentence’s [CLS] token as our reference model. We then consider one hypothesis model for each type of content word (i.e., verb, subject noun, direct object noun, adjective modifying the subject, or adjective modifying the direct object), where each hypothesis model consists of the GloVe embeddings for the instances of the relevant type of content word. We specify 10 unique words that can appear in the position of every content word. For example, in the Transitive + Adjective Corpus the first and second adjectives are randomly selected from separate, non-overlapping vocabularies of 10 words, as are the two nouns. Thus, the number of words that can fill each slot is matched across conditions. We also include a Null Hypothesis Model, which consists of a GloVe embedding for a verb in our vocabulary that does not appear in the sentence.

These hypothesis models allow us to determine which individual content word is encoded to the greatest degree by the full sentence embedding. Standard dependency-parsing frameworks suggest that the main verb is the most salient word in a sentence. If BERT’s full sentence embeddings reflect this intuition, then we would expect the representational geometry of the Verb Model to exhibit the greatest similarity to the representational geometry of our reference model.

When we perform RSA on the models using the Intransitive Corpus, we use a sample size n of 50, as there are only 200 sentences in that corpus. The remaining corpora are larger, allowing for sample sizes n of 200 (the Intransitive + Adjective Corpus has 1,273 sentences, Transitive Corpus has 1,572 sentences, and Transitive + Adjective Corpus has 1,995 sentences, after generating 2,000 sentences from a PCFG for each one and then excluding duplicate sentences).

Results: Across all sentence types, the Verb Model exhibits the greatest representational similarity to our reference model (Table 3). In addition, the Verb Model exhibits a significantly higher representational similarity to the Reference Model than the Null Hypothesis Model does. Thus, the type of content word that is encoded to the greatest degree in the representational geometry of BERT’s full sentence embeddings is the verb, a result that aligns with the primacy of verbs in standard dependency parsing formalisms.

We also find that, for both the Transitive Corpus and Transitive + Adjective Corpus, the Object Model exhibits a greater representational similarity to the reference model than the Subject Model does. We do not have an explanation for why the direct object would play a greater role in the sentence representation than the subject does.

	Verb	Subject	Subj. Adj.	Object	Obj. Adj.	Null
Intransitive	.221	.104	-	-	-	-.034
Intransitive + Adj	.218	.113	.148	-	-	.010
Transitive	.301	.113	-	.185	-	.017
Transitive + Adj	.313	.094	.083	.132	.055	.014

Table 3: Results of Experiment 3. All differences between the Verb Model and any other model are statistically significant, as are both differences between the Subject Model and Object Model ($p < .001$). ‘-’ denotes that the corpus did not contain content words of that type.

7 Comparisons to Other Approaches

One popular approach for analyzing neural network models is the *behavioral* approach, in which a model’s performance is evaluated on some challenge set designed to highlight particular linguistic phenomena. These methods have been used to determine whether various language models track dependencies based on syntactic number (Linzen et al., 2016; Gulordava et al., 2018; Goldberg, 2019). Additionally, natural language inference tasks directly test a model’s ability to infer semantic relationships (McCoy et al., 2019; Bowman et al., 2015). For more examples, see Belinkov and Glass (2019). Behavioral methods are well-suited for holistic questions about a model’s overall handling of language, but here our goal is to analyze specific internal representations independently of behavior. For this question, the behavioral approach is less well-suited because it can only give an indirect window into the structure of the representational space; analyzing the internal representations themselves is much more direct.

Another popular approach, the *diagnostic model* approach, permits analyses that directly investigate internal representations. Diagnostic models are simple (usually linear) classification or regression models that take in activation patterns and predict some information of interest. If the diagnostic model performs well, then the activation pattern is typically said to encode this information. Prior work has used this method to determine that BERT’s representations are sensitive to the hierarchical structure of language (Lin et al., 2019), contain information relevant to a variety of tagging tasks (Liu et al., 2019), contain semantic information (Tenney et al., 2019), and contain much other linguistically-useful information (Rogers et al., 2020; Belinkov and Glass, 2019).

These usages of diagnostic models are generally focused on testing whether an embedding *does or does not* represent certain information, but we instead focus on the *degree* to which the embedding represents information from particular words, and standard diagnostic-model approaches make it difficult to obtain relative results of this sort, because it is possible that information from all words in the sentence is represented in the embedding to some degree. This claim is supported by recent work, which finds that several salient linguistic features can be recovered from *every* word in a sentence using diagnostic classifiers (Klafka and Ettinger, 2020). Thus, it seems that a diagnostic classifier approach may be insufficient for answering the questions that we address here. Indeed, we attempted to devise a diagnostic model approach and applied it to our pronoun coreference and verb subject-sensitivity studies, but found little success (See Appendix A). A modification to the diagnostic model approach which makes this approach more amenable to illuminating relative strengths of encoding is minimum description length (MDL) probing, introduced in the concurrent work of Voita and Titov (2020). MDL probing characterizes how regularly specific information is encoded in vectors, while our approach investigates the representational geometry of a set of encodings. These concepts of regularity and representational geometry are likely related, but we leave for future work an investigation of the precise relationship between them.

One final approach that has been used to study whether linguistic dependencies are captured in BERT is to analyze whether specific attention heads track particular dependency relationships (Clark et al., 2019; Htut et al., 2019). Though this approach enables analysis of the same linguistic phenomena that we study, it does not directly demonstrate whether or how this information is encoded in BERT’s vector representations. Instead, it illustrates which inputs are most salient in generating these vector representations, without revealing what information from those inputs is encoded in the final representation. In addition, because BERT uses many attention heads, the behavior of any one attention head gives only a partial picture of the inputs for a given vector.

8 Related Work

Other Applications of RSA in NLP: Previous work has also explored the application of Representational Similarity Analysis to neural networks. Indeed, one of the first applications of this type of analysis was performed on artificial neural systems (Laakso and Cottrell, 2000). Abnar et al. (2019) use RSA to characterize how the representational geometries of various models change when given different amounts of context. Other work has used RSA to compare the semantic representations of CNN object detection models with that of word vectors (Dharmaretnam and Fyshe, 2018), to compare the representations of utterances in a spoken-word encoder to the representations of those same words in the text or

visual domain (Chrupała, 2019), and to compare the representational spaces of two agents in an emergent language game (Bouchacourt and Baroni, 2018).

Chrupała and Alishahi (2019) also employ RSA to study the correspondence between the representational geometry of various language models (including BERT) and a hypothesis model based on gold syntax trees. They introduce a new technique, *RSA_{regress}*, which merges RSA with a diagnostic model approach. The present work is complementary to their study. Whereas their study compares different neural models against one hypothesis model, our study compares multiple hypothesis models against a single neural model, in order to adjudicate between specific linguistic hypotheses.

Geometric Analyses of NLP Systems: Though we study the representational geometry of contextualized word embeddings, other work has been done to characterize the geometry of the network activations themselves. These are distinct concepts, as our current work can be thought of as analyzing the second-order geometry (comparing the relationships between representations of one model to the relationships between representations of another), while this prior work is analyzing first-order geometry (comparing the representations of one model to the representations of another). This first-order approach has also been used to analyze models’ linguistic capabilities (Wu et al., 2020; Reif et al., 2019; Kim and Linzen, 2019; Hewitt and Manning, 2019). Most directly related, Lin et al. (2019) uses a first-order approach to analyze BERT’s representation of subject-verb agreement and pronoun coreference.

9 Conclusion and Future Work

We have introduced a framework for using representational similarity analysis to adjudicate between hypotheses about the representational geometry of a neural network’s embeddings. We then applied this procedure to BERT embeddings, and demonstrated that they are sensitive to linguistic dependencies. In particular, we showed that BERT’s embeddings of pronouns encode the pronouns’ antecedents more than they encode other nouns, that the embeddings of verbs encode the verbs’ subjects more than they encode non-argument nouns, and that BERT’s sentence embeddings most saliently encode the sentences’ main verbs, as predicted by standard dependency frameworks. Not only do these studies reveal a sensitivity to linguistic dependencies, two of them (the subject-sensitivity and pronoun-coreference studies) demonstrate that these dependencies are *more* salient than relationships between words at the surface level.

This framework can enable the investigation of many linguistically-motivated questions. As long as hypothesis models can be defined to instantiate the relevant linguistic hypotheses, this framework allows researchers to study which of the hypotheses is closer to the truth. Additionally, future work can focus on making more complete hypothesis models. Our studies only required us to account for the *differences* between the representational similarity of different hypothesis models, but not the *magnitude* of that similarity. In future work, we plan to focus on creating more complex hypothesis models that exhibit greater absolute representational similarity to the reference model.

Acknowledgments

We are grateful to Tal Linzen for helpful comments. Any errors are our own. This work was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1746891. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- Samira Abnar, Lisa Beinborn, Rochelle Choenni, and Willem Zuidema. 2019. Blackbox meets blackbox: Representational similarity & stability analysis of neural language models and brains. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 191–203, Florence, Italy, August. Association for Computational Linguistics.
- Yonatan Belinkov and James Glass. 2019. Analysis methods in neural language processing: A survey. *Transactions of the Association for Computational Linguistics*, 7:49–72.

- Diane Bouchacourt and Marco Baroni. 2018. How agents see things: On visual representations in an emergent language game. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 981–985, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal, September. Association for Computational Linguistics.
- Grzegorz Chrupała and Afra Alishahi. 2019. Correlating neural and symbolic representations of language. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2952–2962.
- Grzegorz Chrupała. 2019. Symbolic inductive bias for visually grounded learning of spoken language. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6452–6462, Florence, Italy, July. Association for Computational Linguistics.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D Manning. 2019. What does BERT look at? an analysis of BERT’s attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286.
- 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’06)*, Genoa, Italy, May. European Language Resources Association (ELRA).
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Dhanush Dharmaretnam and Alona Fyshe. 2018. The emergence of semantics in neural network representations of visual information. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 776–780.
- Jörn Diedrichsen and Nikolaus Kriegeskorte. 2017. Representational models: A common framework for understanding encoding, pattern-component, and representational-similarity analysis. *PLoS computational biology*, 13(4):e1005508.
- Kawin Ethayarajh. 2019. How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings. 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)*, pages 55–65, Hong Kong, China, November. Association for Computational Linguistics.
- Yoav Goldberg. 2019. Assessing BERT’s syntactic abilities. *arXiv preprint arXiv:1901.05287*.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. Colorless green recurrent networks dream hierarchically. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1195–1205, New Orleans, Louisiana, June. 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, Volume 1 (Long and Short Papers)*, pages 4129–4138.
- Phu Mon Htut, Jason Phang, Shikha Bordia, and Samuel R Bowman. 2019. Do attention heads in BERT track syntactic dependencies? *arXiv preprint arXiv:1911.12246*.
- Mandar Joshi, Omer Levy, Luke Zettlemoyer, and Daniel Weld. 2019. BERT for coreference resolution: Baselines and analysis. 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)*, pages 5803–5808, Hong Kong, China, November. Association for Computational Linguistics.
- Najoung Kim and Tal Linzen. 2019. Compositionality as directional consistency in sequential neural networks. In *Workshop on Context and Compositionality in Biological and Artificial Neural Systems*.
- Josef Klafka and Allyson Ettinger. 2020. Spying on your neighbors: Fine-grained probing of contextual embeddings for information about surrounding words. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4801–4811, Online, July. Association for Computational Linguistics.

- Nikolaus Kriegeskorte and Rogier A Kievit. 2013. Representational geometry: integrating cognition, computation, and the brain. *Trends in cognitive sciences*, 17(8):401–412.
- Nikolaus Kriegeskorte, Marieke Mur, and Peter A Bandettini. 2008. Representational similarity analysis—connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:4.
- Aarre Laakso and Garrison Cottrell. 2000. Content and cluster analysis: assessing representational similarity in neural systems. *Philosophical psychology*, 13(1):47–76.
- Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. Open sesame: Getting inside BERT’s linguistic knowledge. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 241–253, Florence, Italy, August. Association for Computational Linguistics.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. Assessing the ability of LSTMs to learn syntax-sensitive dependencies. *Transactions of the Association for Computational Linguistics*, 4:521–535.
- 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, Volume 1 (Long and Short Papers)*, pages 1073–1094.
- R. Thomas McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy, July. Association for Computational Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. Universal Dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 1659–1666, Portorož, Slovenia, May. European Language Resources Association (ELRA).
- 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)*, pages 1532–1543.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of NAACL-HLT*, pages 2227–2237.
- Emily Reif, Ann Yuan, Martin Wattenberg, Fernanda B Viegas, Andy Coenen, Adam Pearce, and Been Kim. 2019. Visualizing and measuring the geometry of BERT. In *Advances in Neural Information Processing Systems*, pages 8592–8600.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. A primer in BERTology: What we know about how BERT works. *arXiv preprint arXiv:2002.12327*.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, and Ellie Pavlick. 2019. What do you learn from context? probing for sentence structure in contextualized word representations. In *International Conference on Learning Representations*.
- Elena Voita and Ivan Titov. 2020. Information-theoretic probing with minimum description length. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- John Wu, Yonatan Belinkov, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James Glass. 2020. Similarity analysis of contextual word representation models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4638–4655, Online, July. Association for Computational Linguistics.
- Vitalii Zhelezniak, Aleksandar Savkov, April Shen, and Nils Hammerla. 2019. Correlation coefficients and semantic textual similarity. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 951–962, Minneapolis, Minnesota, June. Association for Computational Linguistics.

Appendix A: Diagnostic Classifier Comparison

We attempted to investigate the argument-sensitivity of BERT’s embeddings of verbs using diagnostic models. To do so, we trained two separate logistic regression classifiers.⁵ The input for both classifiers was the GloVe embedding of any word in the sentence *except the main verb* concatenated with the BERT embedding of the verb. The Subject Classifier was trained to classify whether the GloVe embedding corresponded to the subject of the verb or not. The Non-Argument Noun Classifier was trained to classify whether the GloVe embedding of a word corresponded to the non-argument noun or not. We train these classifiers for both the Prepositional Phrase Corpus and the Relative Clause Corpus. For both classifiers, we use 80% of the data for training and 20% for testing. Looking at Table 4, we see that all classifiers failed to do better than majority-class performance.

We performed a similar investigation of pronoun coreference. The input for our classifiers was the GloVe embedding of any word in the sentence *except the pronoun*, concatenated with the BERT embedding of the pronoun. For both the pronominal and reflexive corpora, we created an Antecedent Classifier (which was trained to identify the antecedent of a sentence), and a Non-Antecedent Classifier (which was trained to identify the non-antecedent noun of a sentence). For both classifiers, we use 80% of the data for training and 20% for testing. Looking at Table 5, we see that all classifiers failed to do better than majority class performance.

	Accuracy	Precision	Recall
PP Subject	.832	.023	.003
PP Non-Argument	.831	.074	.013
RC Subject	.866	0.0	0.0
RC Non-Argument	.870	.224	.037

Table 4: Results of the Diagnostic Classifier approach to Experiment 1. Note that the majority class accuracy for the prepositional phrase corpus classifiers is .83, as 5 out of 6 tested words in every sentence are not subjects/non-argument nouns. The majority class accuracy for the relative clause corpus classifiers is .87, as 6 out of the 7 tested words in every sentence are not subject/non-argument nouns

	Accuracy	Precision	Recall
Reflexive Antecedent	.865	.045	.003
Reflexive Non-Antecedent	.875	.128	.015
Pronominal Antecedent	.867	.042	.003
Pronominal Non-Antecedent	.870	.162	.017

Table 5: Results of the Diagnostic Classifier approach to Experiment 2. Note that the majority class accuracy for the Antecedent and Non-Antecedent classifiers is .87, as 6 out of 7 tested words in every sentence are not antecedents/non-antecedent nouns.

Appendix B: Spearman’s ρ for Embeddings

Zhelezniak et al. (2019) show that it is most appropriate to use Spearman’s ρ as a (dis)similarity metric when embeddings are non-normal. They go on to show that GloVe embeddings violate the assumption of normality by using Shapiro-Wilk tests and analyzing Q-Q plots of some GloVe embeddings. Here, we show that the same result holds for BERT embeddings, and thus show that it is most correct to use

⁵All logistic regression classifiers used the default Scikit-Learn hyperparameters.

Spearman’s ρ (as opposed to cosine similarity or other metrics) as a (dis)similarity metric for BERT embeddings as well.

We analyze our Prepositional Phrase, Relative Clause, Reflexive, and Pronominal Corpora. For these analyses, we only exclude repeated sentences. The resulting Prepositional Phrase Corpus contains 1,935 sentences, the Relative Clause Corpus contains 1,944 sentences, the Reflexive Corpus contains 1,909 sentences, and the Pronominal Corpus contains 1,899 sentences.

Following Zhelezniak et al. (2019), we treat each embedding as a ‘sample of observations from a scalar random variable’. First, we Z-normalize all embeddings, such that they have a mean of 0 and standard deviation of 1. We then apply a Shapiro-Wilk test to the BERT embedding corresponding to every word in the corpus (i.e. not the [CLS] or [SEP] token embeddings). We first apply these tests to the full BERT embeddings, and then repeat the procedure on a random subsample (without replacement) of 300 values from each of these embeddings. We subsample in order to show that a large number of embeddings are still found to be non-normal even with a smaller sample size. See Table 6 for the results.

Corpus	Full Embedding % Non-Normal	Sampled Embedding % Non-Normal	# of Embeddings
Reflexive	99.98	48.24	15,272
Pronominal	100	49.01	15,192
Prepositional Phrase	99.92	48.23	13,545
Relative Clause	99.89	46.59	15,552

Table 6: Results from Shapiro-Wilk tests performed on BERT embeddings for every word in each corpus. Embeddings are considered non-normal if the Shapiro-Wilk test returns a P-value less than .05.

Furthermore, we analyze one BERT embedding for each corpus using Q-Q plots. We see from Figure 5 that these embeddings contain significant outliers, indicating that they are not normally distributed.

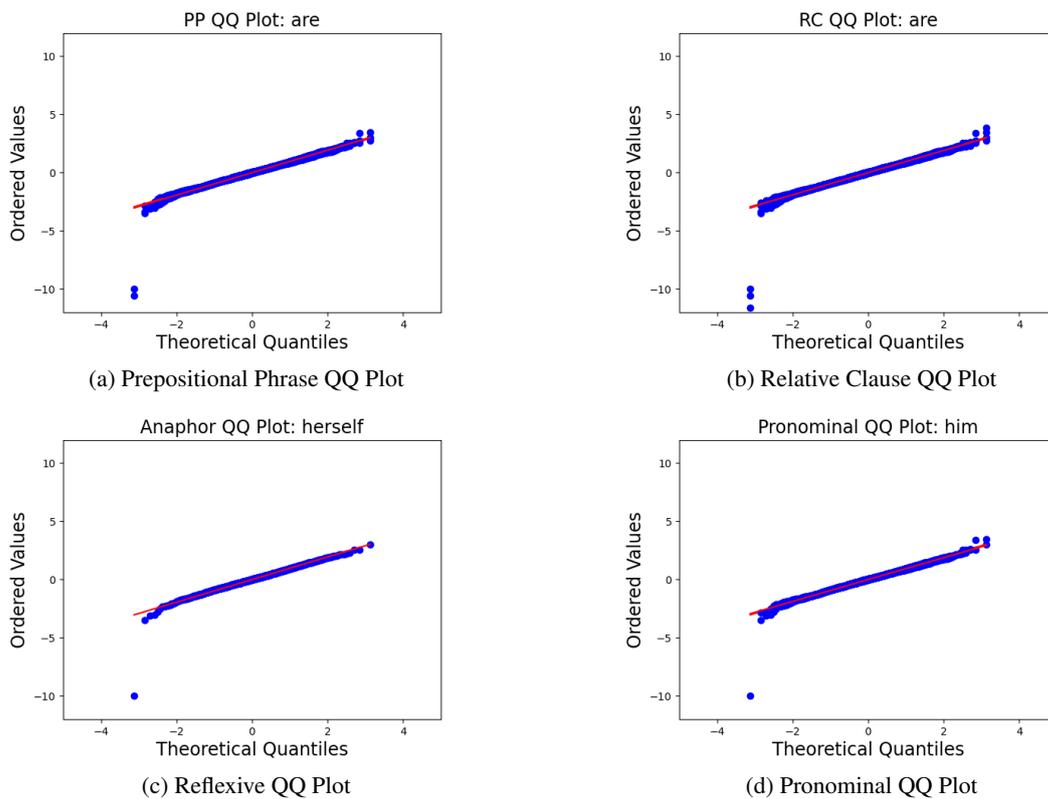


Figure 5: QQ Plots for BERT embeddings from each corpus. We choose embeddings that we use to construct our reference models in Experiments 1 and 2. All plots demonstrate the presence of at least one large outlier, and several other smaller outliers.