

Compositional and Lexical Semantics in RoBERTa, BERT and DistilBERT: A Case Study on CoQA

Ieva Staliūnaitė and Ignacio Iacobacci

Huawei Noah's Ark Lab, London, United Kingdom

{ieva.staliunaite, ignacio.iacobacci}@huawei.com

Abstract

Many NLP tasks have benefited from transferring knowledge from contextualized word embeddings, however the picture of what type of knowledge is transferred is incomplete. This paper studies the types of linguistic phenomena accounted for by language models in the context of a Conversational Question Answering (CoQA) task. We identify the problematic areas for the finetuned RoBERTa, BERT and DistilBERT models through systematic error analysis - basic arithmetic (counting phrases), compositional semantics (negation and Semantic Role Labeling), and lexical semantics (surprisal and antonymy). When enhanced with the relevant linguistic knowledge through multitask learning, the models improve in performance. Ensembles of the enhanced models yield a boost between 2.2 and 2.7 points in F1 score overall, and up to 42.1 points in F1 on the hardest question classes. The results show differences in ability to represent compositional and lexical information between RoBERTa, BERT and DistilBERT.

1 Introduction

It has recently been recognized in the research community that neural network models generally do not exploit the compositionality of language, often relying on superficial features¹. Compositionality refers to the fact that linguistic constituents combine into phrases hierarchically to compose meaning. Contextualized word embeddings (BERT, Devlin et al., 2018; RoBERTa, Liu et al., 2018; DistilBERT, Sanh et al., 2019; etc.) can be expected to be limited in their ability to learn such complex aspects of language, since the models are usually trained with cloze filling and next sentence pre-

¹<https://2020.ieeeicassp.org/program/plenary-speakers/deep-representation-learning/>

diction objectives, and are not directly exposed to semantic relations between phrases.

While larger models yield higher performance, they still lack generalization ability (Talmor et al., 2019) and are computationally expensive, which has led to an increasing interest in reducing model size. For instance, DistilBERT is built using the knowledge distillation technique (Bucilă et al., 2006; Hinton et al., 2015) on BERT, which leads to a lighter and faster model that does not lose much in performance on the majority of the tested tasks. On the other hand, state-of-the-art models use vast amounts of training data - 16GB for BERT, 10 times more for RoBERTa.

This study tackles the question of what type of linguistic knowledge is missing from contextualized word embeddings, comparing models on the basis of their training set size (BERT vs. RoBERTa) as well as their model size (BERT vs. DistilBERT).

The tasks of machine reading comprehension (MRC) and dialogue are particularly fitting for this purpose, due to the fact that they require a system to interpret language within a context and perform semantic and pragmatic inference between sentences. The task in the Conversational Question Answering dataset (Reddy et al., 2019, CoQA) is MRC combined with dialogue - the input to the system is a context document and a dialogue of questions and answers about that text, which lead up to the question that the system is required to answer. An example from CoQA follows.

Background: [...] At the time, the name did not describe a single political entity or a *distinct population of people* [...]

Question n-1: Did the name describe a political body?

Answer n-1: No

Question n: Did it describe a people group?

Answer n: No

In order to answer question n without relying on superficial features, one needs to be able to interpret the logical operators “*and*” and “*or*” (underlined), and their scopes, as well as determine that the italicized phrases are synonymous. This study tackles such cases with linguistically enhanced models. Our assumption in this paper is that if a model performs poorly on the classes of question-answer (QA) pairs that require certain linguistic knowledge X (e.g. negation, disjunction and synonymy in the example above) for their solutions, and if its performance boosts when it explicitly learns X (e.g. through a multitask setting with an auxiliary linguistic task), it can be considered evidence for the original model’s lack of linguistic representations of X . The main contributions of this paper can be summarized as:

1. A qualitative error analysis of the models in the BERT family on a conversational question answering task;
2. An automatic feature extraction method for quantitative evaluation of the lexical and compositional semantics learnt by the models;
3. An improvement of the model performance by injecting linguistic knowledge into the pre-trained models through a multi-task approach.

2 Related Work

There has recently been much interest in diagnostic analysis of BERT, studying what type of linguistic representations it learns. [Rogers et al. \(2020\)](#) provide an overview of such papers under the name of BERTology, showing what types of linguistic phenomena cause difficulty for BERT.

Some such studies focus on compositional semantics in the form of negation, Negative Polarity Items (NPIs) and Semantic Role Labeling (SRL). In testing NPI licensing, [Warstadt et al. \(2019\)](#) perform a cloze task and compare whether BERT predicts a higher probability for an NPI in a licensed context or outside of such a span. They show that while BERT is capable of detecting NPI licensors (e.g. “*don’t*”) and NPIs themselves (e.g. “*ever*”), it only does so successfully in cases where the span of the NPI appears in the canonical position with regard to its licensor. This suggests that the model relies on word order instead of parsing the syntactic dependencies.

When it comes to interpreting negation, [Ettinger \(2020\)](#) analyzes whether BERT predicts a higher

probability for sentences such as “*Robin is not a tree*” and “*Robin is a bird*” than “*Robin is a tree*” and “*Robin is not a bird*”, and conclude that BERT is not very sensitive to negation. This test, however, relies strongly on BERT’s ability to represent lexical semantics of the nouns and the lack of intersection of their typical meanings. It could be argued that it therefore is not a reliable test to tell whether BERT can make logical deductions based on negation. Instead, we argue that it should be tested in a context where “*Robin*” can be anything, including a name of a tree, so that it can be determined whether BERT can infer that in such a case “*Robin*” would not be a bird.

Similarly, [Ettinger \(2020\)](#) has also addressed the question of whether BERT has the knowledge required to infer semantic roles from a text. For example, she tests BERT’s ability to assign a higher probability to the word “*served*” in a statement such as “*the restaurant owner forgot which customer the waitress had served*” than in “*the restaurant owner forgot which waitress the customer had served*”. This test is analogous to the previous example with a Robin, in the sense that it tests the model’s ability to learn biases of common semantic roles that certain nouns manifest. The model can simply rely on the fact that one of these two word orders is more likely than the other, however this does not provide evidence that BERT can make inferences about semantic roles. Thus, it remains to be verified whether BERT can abstract the semantic roles from any range of naturally occurring sentences, some of which exhibit uncommon semantic role occurrences (e.g. customers serving waiters).

In a similar vein, previous research has shown that the embeddings of antonyms in models such as BERT are not clearly distinguishable ([Talmor et al., 2019](#)). This, similarly to issues with negation, shows that BERT is not good at representing non-intersecting denotations.

What is more, [Richardson et al. \(2019\)](#) show that BERT performs poorly on artificially constructed diagnostic items which test the model’s ability to perform logical inference. Nonetheless, they demonstrate that it is possible for BERT to extrapolate the relevant linguistic phenomena quickly by finetuning the model on the same artificial data.

Pragmatics also plays a role in determining relations between sentences, however this field has been less explored with regard to contextualized word embeddings. Some research has probed

BERT on its capabilities to infer pragmatic phenomena related to negation, such as factives, conditionals and questions (Jiang and de Marneffe, 2019). For instance, one has to make the pragmatic inference that a speaker would only utter “*not a one of them realized I was not human*” if their lack of humanity was already established as common knowledge. Jiang and de Marneffe (2019) show that BERT takes longer to learn such complex reasoning than negation.

In contrast, some studies research what common sense knowledge and abstract reasoning that BERT and other language models learn. Talmor et al. (2019) show that there is a large gap between BERT and RoBERTa with regard to their inference abilities. For instance, since RoBERTa is trained on significantly more data, it can determine which person is older based on their ages or dates of birth, while BERT cannot. Interestingly, however, even RoBERTa is shown to rely on the range of examples seen at training time, as it is not able to generalize to the ages of people who are not born between 1920s and 2000s. This suggests that there is a need for a more abstract reasoning ability in models such as BERT, which does not seem to be solved by an increased size of the training set. Finally, Ju et al. (2019) show that even for the RoBERTa-based abstractive model, which reached state-of-the-art results on CoQA at the time, questions with numerical answers account for a disproportionately large fraction of errors.

Based on the studies conducted so far, one general trend appears to be pertinent. That is, while many studies have explored the linguistic knowledge of BERT, it is still not clear whether BERT is able to infer compositional structures from text as opposed to relying on biases. In addition, to the best of our knowledge, no probing tasks have been performed on BERT in the conversational question domain, which is fitting for analyzing BERT’s behaviour in complex reasoning and inference. Finally, while larger models such as RoBERTa yield gains in performance, they still lack in generalization ability. Thus, this paper aims to shed light on the less scrutinized aspects of BERT’s linguistic capabilities.

3 Dataset

The CoQA dataset² is used as a case study in this paper. It covers several domains and amounts to

²<https://stanfordnlp.github.io/coqa/>

127,000+ samples including a story, a QA pair and the dialogue history. The answers to the questions are based on the context document, however they can be paraphrases. The training data also contains *rationales*, which are the spans of the background text containing both the answer and the context required to determine the answer. The test set is composed of the Reddit and Science domains, while the rest of the domains are split between train, development and test (see Table 1). Covering various domains makes CoQA diverse with regard to style and content of the dataset, whereas the addition of the dialogue history makes the dataset interesting in that it combines different language modes - a written paragraph and a conversation. Such diversity allows for a robust analysis of linguistic relations since it gives access to negation in questions as well as statements, fictional settings of unusually flipped semantic roles, counting of any abstract or concrete objects, etc. The state-of-the-art models on this dataset (Ju et al., 2019) use RoBERTa, while the dataset has not received much attention with smaller or distilled models such as DistilBERT.

Domain	#Passages	#QA pairs	Passage Length	#Turns per passage
Childrens Stories	750	10.5k	211	14.0
Literature	1,815	25.5k	284	15.6
School exams	1,911	28.6k	306	15.0
News	1,902	28.7k	268	15.1
Wikipedia	1,821	28.0k	245	15.4
Reddit	100	1.7k	361	16.6
Science	100	1.5k	251	15.3
Total	8,399	127k	271	15.2

Table 1: CoQA dataset details (Reddy et al., 2019).

4 Baseline Models

The input to the model is a concatenation of the background story, the latest dialogue history of 64 tokens, and the current question. The length of the input is limited to 512 tokens. We build the baseline RoBERTa, BERT and DistilBERT base models for CoQA as extractive models, within the framework of Wolf et al. (2019) and following Wu et al. (2019), who produce the highest results with a BERT-based extractive model on CoQA. An extractive model does not generate the answer as an abstractive model would, but selects the span in the document that best matches the gold answer. In order to train our extractive models, the substrings

of the rationales which are most similar to the gold answers (as measured by F1) are selected as the training labels.

Following a standard procedure, a linear classifier head is added on top of BERT with ReLU activation which classifies every token in the input sequence as start or end of the answer span. Another linear classifier predicts whether each token in the input span falls within the rationale span or outside of it. Finally, one more classifier predicts whether the example is FREEFORM, has a YES/NO answer or is UNANSWERABLE. YES/NO/UNANSWERABLE answers are used instead of the predicted span if the model predicts the latter classes with higher confidence than the start and end tokens of the answer. Models are trained for 4 epochs (taking a few hours on a single GeForce GTX 1080 Ti GPU) with a learning rate of $3e-5$ and AdamW optimizer (Loshchilov and Hutter, 2017). The same hyperparameters are used for all models trained in this work.

5 Baseline Results and Error Analysis

On the development set³, the RoBERTa model gets 81.2 points F1⁴, BERT scores 76.9 F1 and falls two points short of the Wu et al. (2019) implementation, and DistilBERT scores 66.6 F1, which establishes a baseline as this is the first work using DistilBERT on CoQA (see Table 2).

To resolve what types of linguistic inference are the hardest for the baseline models, several potentially difficult QA classes are analyzed. They are defined based on the findings of previous research as well as the observations of a qualitative evaluation of the errors made by the BERT model. Then, a quantitative evaluation of how the baseline models perform on each class is performed. There is ample variation in how the models score on various example classes (see Table 2). Nonetheless, a noticeable trend appears across the three models which are failing in similar classes, with DistilBERT lagging behind BERT in most of the classes, by up to 15 points in F1 in some, and RoBERTa beating BERT by a smaller margin.

The first expected source of error for the base-

³Due to a limitation of at most 2 submissions per week for official evaluation on CoQA, the experiments on the test set are not included.

⁴The results do not compare to the 89.5 F1 score of the extractive model of Ju et al. (2019), who do not report which RoBERTa model they use. The difference would be accounted for if they use RoBERTa large.

line models is the inability to count listed phrases. Since the models are extractive, counting cannot fall within their capabilities. For example, the rationale below lists the characters in the background story:

Rationale: a poor man Ti, his son Dicky and their alien dog CJ7

Question: How many characters were there?

Answer: Three

In the current setting the models cannot chunk the text into noun phrases and then count the chunks to answer the question, as they are limited to extracting the answer from the background text. While the model performance is satisfactory on a wide range of questions with numerical answers (NUM), they fail consistently on questions with answers in the integer set between 1 and 5 (1-5). The NUM class is defined using a state-of-the-art rule-based question classification system from Madabushi and Lee (2016), which evaluates each QA based on the question alone. The contrast between the scores on the two classes can be explained by the fact that while extracting numerical answers such as dates is easy for the models, they struggle on the task of counting linguistic objects, which are usually manifested in low value integers.

The second expected problematic area is negation. Negation cues are words (“*not*”, “*without*”) or morphemes (“*dis-*”, “*un-*”, “*ab-*”), which negate the context that they span over. The example below illustrates two ways that the model can fail in face of negation cues.

Rationale: Something looked like a bird’s belly [...] it was not a bird’s belly [...] a bottle floated there

Question: What looked like a bird’s belly?

Answer: A bottle

Wrong answer 1: A bird’s belly

Wrong answer 2: Not a bird’s belly

The most general type of error is neglecting the negation cue altogether and answering the question with **Wrong answer 1**. This reflects on the model’s inability to determine that the noun phrase “*a bird’s belly*” falls under the scope of the negation cue “*not*”. **Wrong answer 1** would be the correct answer if the phrase was not negated. The second and more rarely observed type of error reflects a lack of pragmatic knowledge as opposed to semantic or syntactic. In **Wrong answer 2** the model could be argued to have answered correctly as it

Model	class															
	overall	NUM	1-5	NEG	YES	NO	SENT	ANT	ORD	SRL-	SRL+	HUM	LOC	ENT	SURP	
size	7983	972	128	436	790	682	2443	185	6817	3175	1242	1418	624	1089	433	
RoBERTa	81.2	76.4	40.1	72.9	89.5	85.9	83.0	82.3	81.8	80.7	83.5	83.2	83.6	83.1	78.6	
BERT	76.9	77.2	41.9	68.9	85.3	77.3	79.0	74.8	77.2	76.4	78.5	77.4	80.8	79.4	74.1	
DistilBERT	66.6	69.6	36.8	56.7	82.1	71.4	69.0	71.3	67.2	65.8	70.0	64.2	65.4	70.0	58.9	

Table 2: The results of the baseline models on the CoQA development set (F1 scores). The grayscale colors reflect the variation within models between QA classes.

is technically true that what looked like a bird’s belly was not a bird’s belly. However assuming Grice’s maxim of quantity, which states that one should be as informative as required (Grice, 1989), the answer is not satisfactory. **Wrong answer 2** is not informative at all as it has already been implied by the question. We define the NEG QA class as containing answers that are embedded under negation. For recognizing such answers, negation cues and their spans are detected with a BERT-based model following Khandelwal and Sawant (2020) and trained on the Sherlock dataset (Morante and Blanco, 2012). We reproduce the results on that dataset before using the model for detecting the negated spans in the background documents in the CoQA dataset to find the NEG type answers. Our baseline models perform worse on the NEG QA class than overall, and score much higher on questions with YES answers than NO answers, which suggests that the models do not interpret negation correctly. The effect of negation on performance is particularly stark in the case of DistilBERT.

Furthermore, the ANT class is composed of examples in which the rationale contains antonyms of the words in the question, using WordNet (Fellbaum, 1998). Here explicit negation is not necessarily involved, however the model’s ability to reason over semantic polarity is tested in this QA class. Our baseline results on class ANT are in line with previous conclusions stating that BERT is not good at representing antonymy, as it scores lower on this class than overall. Yet interestingly, DistilBERT as well as RoBERTa perform better on this subclass of questions than overall. We conjecture that lexical semantics is the strongest feat of BERT, therefore it is likely that DistilBERT retains most of the lexical information such as antonymy through the process of distillation. On the other hand, RoBERTa learns more about lexical features such as antonymy from the huge size of the training set.

In addition, SENT is a QA class in which the sen-

timent of the sentence containing the rationale is different from the sentiment of the question. The class items are determined by sentence splitting (Honnibal and Montani, 2017) and sentence-level sentiment classification (Wolf et al., 2019). This class is intended to capture examples where the polarity between the question and the answer can be expressed not only by negation or antonymy but also any other means, for example pragmatics. However, a qualitative analysis of the examples of the SENT class shows that the examples which contain contradictory sentiments between the question and the answer mostly do not require one to determine the sentiment in order to answer the question correctly. For instance, the question below has a slightly negative connotation about a long job hunt:

Rationale: Nearly four years later, as Obama seeks reelection, Casillas has finally landed his first full-time job, emerging out of the group known as the long-term unemployed.

Question: How much later did he get his next job?

Answer: Nearly four years later

In contrast, the rationale takes a more positive outlook. Nonetheless, the answer is “four years”, regardless of whether that is considered too long or not. There generally does not appear to be a relation between the sentiment values of the rationale and question and the model predictions. Accordingly, neither of the baseline models struggle to answer SENT questions.

Moreover, as stated in Reddy et al. (2019), the order of questions on the CoQA dataset follows the natural order of text. That is, later questions refer, generally, to information presented towards the end of the background story. Hence, the one answering the questions ought to make inferences about what has already been discussed and where in the story they are when a given question is posed. In some cases this knowledge can be crucial for reaching the correct answer. For instance, in the example

below, one has to determine what time within the story the question refers to.

Background: [...] The hospital had been bombed and Hans had made his way back into Western Germany on foot. [...]

Question: Was he in the Eastern or Western part?

Answer: East Germany

That is, the answer differs depending whether the question refers to Hans’s location prior or post his journey, which is also reflected in the different answers that the human annotators provide to this question. In this case the answer is “*East Germany*” even though that part of the country is never mentioned in the text, which makes the example very challenging with regard to pragmatic inference. In order to evaluate how our models perform with regard to following the dialogue flow, they are evaluated on items which do in fact follow the order of the document, so that the answer to question n in the text is subsequent to the answer to question $n - 1$ (ORD). It appears that all baseline models are able to infer this order to some extent and perform better on such questions than those that jump to previous passages in the text.

Furthermore, examples are classified with regard to whether the order of the semantic roles mentioned in the question is the same (SRL+) or different (SRL-) to the semantic role order in the sentence containing the rationale. SRL is performed employing an AllenNLP (Shi and Lin, 2019) model. To illustrate, Figure 1 shows an example where the roles of agent (Arg0) and patient (Arg1) are reversed in the question by means of a passive voice. All three models fail more on such examples with different word order between the question and the rationale, scoring lower on the SRL- class than overall or SRL+. The results of the experiments show that the models find the correct answer more frequently when they can rely on the word order, avoiding the need to reason over semantic roles.

Finally, some of the observed issues are induced by the model choosing prominent entities as answers regardless of their actual relation to the question at hand. For instance, a document tells a fictional children’s story wherein foods and utensils are anthropomorphised, as in the example below:

Rationale: cereal is winning the race in a bowl of milk

Question: Who is a good swimmer?

Answer: cereal

The baseline BERT model chooses a human entity that is mentioned by name at the beginning of the text instead of the inanimate entity that takes up the agent position in the rationale. In contrast, if “*cereal*” is substituted with a common name such as “*Mark*” in the background document, the model correctly chooses it as the answer. This suggests that the model relies on lexical semantics and biases about types of entities denoted by nouns more than analyze the semantic relations in the relevant sentence. Therefore, we define a QA class where the rationale contains entities that have high entropy and are thus surprising given the rest of the sentence (Hale, 2001; Levy, 2008; Smith and Levy, 2008), like “*cereal*” in the above example. In order to detect such entities, proper nouns (as tagged by spaCy, from Honnibal and Montani, 2017) are masked and BERT is used to evaluate the likelihood of the original word being the filler for that mask. Words that fall below the likelihood threshold of $5e-5$ are then deemed to be surprising entities⁵. All three models perform worse on questions about surprising entities (SURP) than overall, with DistilBERT exhibiting the largest margin.

Moreover, the classes of human (HUM), location (LOC) or general entities (ENT), as classified by Madabushi and Lee (2016), test the models’ ability to answer questions about entity roles. RoBERTa and BERT’s performance on these classes is higher than their overall performance. On the other hand, DistilBERT fails on HUM and LOC entity questions more than other QA types. For many HUM and LOC questions there are multiple entities in the text that fit the entity type. Together with the results on the SURP class, this is an indication that DistilBERT relies on entity type more than the larger models.

The baseline results on the various classes corroborate most of the results of previous research on BERT’s shortcomings. Moreover, the results show that DistilBERT mostly repeats the same mistakes and often more gravely, except for some cases of lexical semantics. DistilBERT appears to lose more of BERT’s already limited representations of the formal aspects of language and have stronger biases. Finally, RoBERTa also exhibits a lack of ability to perform compositional reasoning and reaches the highest scores on the more lexical QA types.

⁵The threshold is selected after finetuning the model augmented with surprising word substitution, as discussed in Section 6.




Figure 1: Illustration of the order of semantic roles in a CoQA example.

6 Model Enhancement

6.1 Auxiliary tasks

The methods for defining QA classes are also used as sources of linguistic knowledge which are incorporated in the baseline models to enhance their performance with regard to the respective classes. Firstly, besides the existing FREEFORM, YES/NO and UNKNOWN, five additional classifiers of integer answers between 1 and 5 are defined within the model, as it would be impossible for the models to answer counting questions extractively. This results in the *base#* model⁶. Then, four additional enhanced multitask models are built in order to tackle the issues observed in the previous section.

For every enhanced model (*negation#*, *order#*, *sentiment#*, *srl#*), the training data is tagged with annotations of the relevant linguistic information that was also used for defining the problematic classes. For *negation#*, tokens are labelled as under the scope of negation (1) or not (0); for *order#* they are labelled as occurring after the answer to question $n - 1$ (1) or not (0); for *sentiment#* they are labelled as part of a sentence with a negative sentiment (1) or not (0); while for *srl#* a multi-label setup is used where every token is labelled as either taking a particular semantic role (1) or not (0). Each of these sets of labels are then used as an additional training goal for the model. The loss of a given additional goal is added to the main loss in a weighted sum. Namely, the sum is calculated as:

$$\begin{aligned}
 loss = & \textit{start_of_answer_loss} / 2 + \\
 & \textit{end_of_answer_loss} / 2 + \\
 & \textit{rationale_loss} + \\
 & \alpha * \textit{additional_goal_loss}
 \end{aligned}$$

with *alpha* values from [0.01, 0.1, 0.2].

In addition to the multitask approach, other architectures were explored for adjoining the information from the four knowledge sources. These approaches include supplying the information as an additional input feature added or concatenated

⁶The # sign next to the model name marks that the model includes a classifier for low-valued integer answers.

to the BERT model at the level of BERT inputs themselves or the BERT model outputs. However, experiments with the latter methods showed no considerable increase in the model performance. Thus, the multitask approach was finally adopted for enhancing models. The multitask approach is also beneficial as the model can be applied to other test sets without the overhead of extracting the linguistic knowledge from the new set.

6.2 Surprising Word Substitution

One more enhanced model is produced by augmenting the training data by means of surprising word substitution (*surprisal#*). Supplementary data samples are produced by substituting surprising entities in the CoQA training set with entities that would be very likely to take their place, according to BERT. In order to ensure that the sentence structure is not affected and an entity is substituted with another entity, the substituting word was only selected if the new word was also tagged as a proper noun in the newly produced sentence. This procedure leads to 5880 additional samples for training. The reason behind adding these items on top of the training set instead of substituting the surprising examples is the intention to provide rare entities with better context instead of ignoring them. Many models in NLP suffer from strong social biases and therefore this approach attempts to level the playing field for rare entities by introducing them in the same contexts as the common entities. Such a method could potentially also be applied to larger datasets and more early stages such as pretraining.

Finally, the enhanced models are combined into an *ensemble* in order to combine the strongest points of each model. In order to use the specialized knowledge from each model where it is relevant, the ensemble is created by selecting the model with the highest confidence for each prediction from all the models, including *base#*.

7 Enhanced Model Results

The results of all the enhanced models on all QA classes on the development set are presented in Table 3. BERT and RoBERTa gain most in terms of

class	overall	NUM	1-5	NEG	YES	NO	SENT	ANT	ORD	SRL-	SRL+	HUM	LOC	ENT	SURP
model	RoBERTa														
<i>base</i>	81.2	76.4	40.1	72.9	89.5	85.9	83.0	82.3	81.8	80.7	83.5	83.2	83.6	83.1	78.6
<i>base#</i>	82.1	81.4	68.8	73.5	89.6	86.3	83.8	82.4	82.5	81.6	83.9	83.0	85.0	84.2	80.5
<i>negation#</i>	81.7	81.2	69.4	73.0	89.5	85.0	83.9	86.0	82.2	81.0	83.2	82.4	83.1	83.2	79.8
<i>sentiment#</i>	81.7	81.2	71.5	71.8	90.3	84.1	83.6	83.5	82.1	81.3	83.2	82.1	85.4	83.2	80.7
<i>order#</i>	82.3	82.7	76.0	74.3	90.0	85.1	84.2	85.8	82.8	82.3	84.4	83.6	85.4	83.8	80.3
<i>srl#</i>	82.2	82.6	73.3	72.9	88.6	85.7	84.5	83.5	82.6	81.6	83.6	83.4	85.2	83.6	80.4
<i>surprisal#</i>	82.1	81.7	73.2	72.7	89.3	86.0	83.6	84.5	82.4	81.8	83.9	82.3	84.2	83.7	79.4
<i>ensemble</i>	83.9	82.7	82.2	81.7	81.7	82.1	83.0	82.9	82.6	82.1	82.3	82.4	83.0	82.6	83.3
model	BERT														
<i>base</i>	76.9	77.2	41.9	68.9	85.3	77.3	79.0	74.8	77.2	76.4	78.5	77.4	80.8	79.4	74.1
<i>base#</i>	76.9	79.6	63.5	69.7	86.3	80.2	79.5	79.6	77.4	75.7	79.6	77.1	79.6	77.9	72.6
<i>negation#</i>	77.0	79.5	63.2	79.7	84.5	81.0	79.0	79.7	77.4	76.1	79.5	76.3	81.0	78.4	76.5
<i>sentiment#</i>	75.2	76.9	45.6	66.2	81.4	79.1	76.7	76.2	75.6	75.1	76.8	74.8	78.9	76.9	71.6
<i>order#</i>	76.3	76.6	45.2	69.0	83.2	81.1	78.0	78.3	76.9	75.9	78.9	77.3	80.4	78.5	72.8
<i>srl#</i>	77.2	80.8	67.0	68.5	84.4	81.2	79.5	86.1	77.5	78.6	75.8	76.0	81.5	79.2	74.6
<i>surprisal#</i>	76.7	80.5	70.4	68.5	84.1	81.6	78.6	82.7	77.0	75.9	78.8	76.2	78.8	77.1	73.7
<i>ensemble</i>	79.2	80.8	62.2	70.8	86.2	82.8	81.1	81.5	79.6	78.7	81.3	79.8	82.7	80.7	76.2
model	DistilBERT														
<i>base</i>	66.6	69.6	36.8	56.7	82.1	71.4	69.0	71.3	67.2	65.8	70.0	64.2	65.4	70.0	58.9
<i>base#</i>	66.9	69.5	37.0	58.0	82.8	72.4	69.5	70.9	67.4	65.9	69.7	65.0	66.2	70.2	60.0
<i>negation#</i>	65.3	70.1	45.4	58.3	80.2	73.2	68.4	68.9	66.0	64.6	69.6	62.2	63.3	67.2	58.7
<i>sentiment#</i>	64.9	68.6	35.4	54.9	80.7	72.4	68.1	66.8	65.4	64.6	67.5	62.0	61.5	67.6	60.9
<i>order#</i>	64.5	69.3	34.9	53.5	81.4	72.5	67.1	69.2	65.3	64.3	66.8	60.2	62.3	66.5	59.9
<i>srl#</i>	65.8	69.0	39.1	56.4	83.4	72.0	68.4	72.2	66.5	65.3	68.1	62.9	62.8	68.7	57.9
<i>surprisal#</i>	66.8	70.3	35.6	59.2	82.9	72.0	69.0	72.0	67.1	65.9	69.2	63.8	65.8	70.2	60.4
<i>ensemble</i>	68.8	65.7	65.8	65.3	64.9	66.8	67.4	66.1	65.7	66.9	64.5	67.2	67.8	67.3	67.5

Table 3: The results of the baseline and enhanced models on the CoQA development set (F1 scores). The heatmap colors reflect the variation within QA classes between models. The results of the *base#* models should be compared to the *base* results in gray to see the effect of adding the numerical answer classifier, whereas the remaining models should be compared to the *base#* results in gray in order to see the effects of the additional linguistic knowledge.

F1 with the counting model on counting questions (*base#* on 1-5), while DistilBERT only improves on these questions remarkably with the *ensemble* model, requiring more auxiliary resources than the larger models for this level of abstraction.

Moreover, BERT appears to learn formal aspects of semantics in the multitask setting. The *srl#* model improves on the QA class with semantic roles in a different order in the question and the answer (SRL-), while the *negation#* model improves the results on the answers that require interpreting negation (NEG and NO). To illustrate, while the *base#* model fails to take into account the negative morpheme “dis-” and answers “yes” to the question below, the *negation#* model avoids such errors.

Rationale: Some refer to Tolstoy’s disagreement with state-backed religion.

Question: Did he agree with the states view on Religion?

Answer: No

This shows that the *negation#* model has learnt to recognize the negation spanning over the “agree” root, which plays a deciding role for determining the correct answer to the question.

In the meantime, RoBERTa does not improve on the QAs in the negation or Semantic Role classes. One might say that RoBERTa has learnt the abstract linguistic representations already, however its *base* results show that it makes many of the same mistakes as BERT and DistilBERT on NEG and SRL-. In fact, BERT outperforms RoBERTa on the NEG class when enhanced with the explicit information about negation (the *negation#* model). This, combined with the fact that RoBERTa gets a big improvement on NEG only when various linguistic features are combined into an *ensemble*, suggests that RoBERTa mostly relies on better lexical representations for its higher scores, which is only outweighed when many compositional semantic cues are provided. Similarly, DistilBERT only gains a small boost over the baseline on NEG and

NO QA classes with *negation#* and also requires an *ensemble* to improve on SRL-.

On the other hand, BERT and DistilBERT improve on the HUM and LOC classes with the *ensemble* models, demonstrating an ability to improve its lexical representations. RoBERTa does not yield an improvement in this case, however even its *base* model performs relatively well on these classes.

Furthermore, BERT and DistilBERT do not get a boost in the cases of pragmatics, namely *sentiment#* and *order#*. In contrast, RoBERTa gets a boost over the ANT class from the *sentiment#*, and gains the largest increases across almost all classes from *order#*. It appears that RoBERTa can improve on its already high score on items containing antonymy relying on more pragmatic aspects of lexical semantics, and also is the most receptive to the pragmatic aspects of dialogue in CoQA.

Moreover, the model trained on the dataset which was augmented through surprising word substitution (*surprisal#*) improves over *base#* on the class with surprising entities (SURP) with BERT and DistilBERT. This shows that the method helps the models generalize better to cover new examples with surprising entities and get rid of some of the biases about entities. For example, while the *base#* BERT model answers the question below incorrectly, the model enhanced with *surprisal#* selects the name of the relevant character.

Background: [...] When Link Merwell went down again Dave looked at Nat Poole, thinking that lad might possibly attack him. [...] Merwell brought this on himself—you know that as well as I do. He’s pretty badly hurt, I fear. [...]

Question: Who is badly injured?

base# Answer: Nat Poole

surprisal# Answer: Link Merwell

The *base#* model simply selects the name that is more common. In contrast, the enhanced model detects the correct name of the person who is performing the role in the question. Substituting the name in the original background document with a mask token shows that BERT assigns a larger probability to the name “*Nat Poole*” than “*Link Merwell*”. Such examples stand to show that a BERT model is biased towards more common entity names. Nonetheless, the models are able to learn to abstract away from the biases of the connotations and lexical semantics of entity names and find the entities with the right semantic roles even

for rare entities, given that the model has been exposed to both rare and common entities in the same contexts during training.

Interestingly, in the case of BERT the largest boosts in the SURP class are produced by the *negation#* and *srl#* models, showing that focusing on compositional information such as semantic roles and negation helps the model to be less biased towards very prominent lexical information of stereotypical entities as discussed in Section 5. In contrast, in the case of RoBERTa, *surprisal#* does not yield an improvement on the SURP class. RoBERTa requires all of the enhanced models to be combined into an *ensemble* in order to get rid of the biases, suggesting that its focus on (biased) lexical representations is even stronger than BERT or DistilBERT’s.

Finally, the *ensemble* models perform better on virtually all classes and provide a better overall score. This is to be expected as the enhanced models, while performing at a similar level, make different errors and complement each other with their respective specializations.

8 Conclusion

By and large, this paper provides additional evidence that models like BERT lack linguistic abstraction abilities, often relying on superficial features such as entity name biases or word order to answer questions in the conversational question answering task. More precisely, we show that BERT, RoBERTa and DistilBERT models mostly fail on questions that require inference over the compositional aspects of language, such as semantic roles and negation. Furthermore, we find that while RoBERTa improves over BERT’s performance, a large portion of its gain comes from better lexical representations and it appears to fall short of solving the compositional semantic problems. What is more, we provide the first evaluation and analysis of DistilBERT on CoQA, showing that DistilBERT, more so than BERT, relies on lexical information most and lacks capacity to learn compositional representations. Finally, we show that all the BERT-like models tested can be enhanced to a varying extent by feeding them linguistic knowledge through a multitask approach. Even a small amount of training data for linguistic information such as negation can provide a very large boost to the model performance on the QA classes which rely on that information.

References

- Cristian Bucilă, Rich Caruana, and Alexandru Niculescu-Mizil. 2006. [Model compression](#). In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541, Philadelphia, Pennsylvania, USA.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). *arXiv preprint arXiv:1810.04805*.
- Allyson Ettinger. 2020. [What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. Bradford Books.
- H Paul Grice. 1989. *Studies in the Way of Words*. Harvard University Press.
- John Hale. 2001. [A probabilistic Earley parser as a psycholinguistic model](#). In *Proceedings of Second Meeting of the North American Chapter of the Association for Computational Linguistics*, pages 1–8, Pittsburgh, Pennsylvania, USA.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. [Distilling the knowledge in a neural network](#). *arXiv preprint arXiv:1503.02531*.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- Nanjing Jiang and Marie-Catherine de Marneffe. 2019. [Evaluating BERT for natural language inference: A case study on the CommitmentBank](#). 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 6088–6093, Hong Kong, China.
- Ying Ju, Fubang Zhao, Shijie Chen, Bowen Zheng, Xuefeng Yang, and Yunfeng Liu. 2019. [Technical report on Conversational Question Answering](#). *arXiv preprint arXiv:1909.10772*.
- Aditya Khandelwal and Suraj Sawant. 2020. [NegBERT: A transfer learning approach for negation detection and scope resolution](#). In *Proceedings of The 12th Language Resources and Evaluation Conference (LREC)*, pages 5739–5748, Marseille, France.
- Roger Levy. 2008. [Expectation-based syntactic comprehension](#). *Cognition*, 106(3):1126–1177.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#). *arXiv preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2017. [Decoupled weight decay regularization](#). *arXiv preprint arXiv:1711.05101*.
- Harish Tayyar Madabushi and Mark Lee. 2016. [High Accuracy Rule-based Question Classification using Question Syntax and Semantics](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1220–1230, Osaka, Japan.
- Roser Morante and Eduardo Blanco. 2012. [*SEM 2012 Shared Task: Resolving the Scope and Focus of Negation](#). In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 265–274, Montréal, Canada.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. [CoQA: A conversational question answering challenge](#). *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Kyle Richardson, Hai Hu, Lawrence S Moss, and Ashish Sabharwal. 2019. [Probing natural language inference models through semantic fragments](#). *arXiv preprint arXiv:1909.07521*.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A Primer in BERTology: What we know about how BERT works](#). *arXiv preprint arXiv:2002.12327*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter](#). *arXiv preprint arXiv:1910.01108*.
- Peng Shi and Jimmy Lin. 2019. [Simple BERT Models for Relation Extraction and Semantic Role Labeling](#). *arXiv preprint arXiv:1904.05255*.
- Nathaniel J. Smith and Roger Levy. 2008. [Optimal Processing Times in Reading: a Formal Model and Empirical Investigation](#). In *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci 2008)*, volume 30, pages 595–600, Washington, DC, USA.
- Alon Talmor, Yanai Elazar, Yoav Goldberg, and Jonathan Berant. 2019. [oLMpics—On what Language Model Pre-training Captures](#). *arXiv preprint arXiv:1912.13283*.
- Alex Warstadt, Yu Cao, Ioana Grosu, Wei Peng, Hagen Blix, Yining Nie, Anna Alsop, Shikha Bordia, Haokun Liu, Alicia Parrish, et al. 2019. [Investigating BERT’s Knowledge of Language: Five Analysis Methods with NPIs](#). 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 2877–2887, Hong Kong, China.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface's transformers: State-of-the-art natural language processing](#). *ArXiv*, abs/1910.03771.

Jindou Wu, Yunlun Yang, Chao Deng, Hongyi Tang, Bingning Wang, Haoze Sun, Ting Yao, and Qi Zhang. 2019. [Sogou Machine Reading Comprehension Toolkit](#). *arXiv e-prints*, page arXiv:1903.11848.