

Attending to Entities for Better Text Understanding

Pengxiang Cheng

Department of Computer Science
The University of Texas at Austin
pxcheng@utexas.edu

Katrin Erk

Department of Linguistics
The University of Texas at Austin
katrin.erk@utexas.edu

Abstract

Recent progress in NLP witnessed the development of large-scale pre-trained language models (GPT, BERT, XLNet, etc.) based on Transformer (Vaswani et al. 2017), and in a range of end tasks, such models have achieved state-of-the-art results, approaching human performance. This clearly demonstrates the power of the stacked self-attention architecture when paired with a sufficient number of layers and a large amount of pre-training data. However, on tasks that require complex and long-distance reasoning where surface-level cues are not enough, there is still a large gap between the pre-trained models and human performance. Strubell et al. (2018) recently showed that it is possible to inject knowledge of syntactic structure into a model through supervised self-attention. We conjecture that a similar injection of semantic knowledge, in particular, coreference information, into an existing model would improve performance on such complex problems. On the LAMBADA (Paperno et al. 2016) task, we show that a model trained from scratch with coreference as auxiliary supervision for self-attention outperforms the largest GPT-2 model, setting the new state-of-the-art, while only containing a tiny fraction of parameters compared to GPT-2. We also conduct a thorough analysis of different variants of model architectures and supervision configurations, suggesting future directions on applying similar techniques to other problems.

1 Introduction

When does it help performance on NLP tasks to explicitly take linguistic structure into account? Large-scale pre-trained language models such as ELMo (Peters et al. 2018), GPT (Radford et al. 2018), BERT (Devlin et al. 2019), and XLNet (Yang et al. 2019) have recently achieved state-of-the-art results on a wide range of tasks. These models, mostly built on a stacked self-attention architecture as in the Transformer (Vaswani et al. 2017), are not explicitly trained to take linguistic structure into account, but they have been shown to encode some linguistic knowledge anyway, in particular knowledge related to syntactic structure (Tenney et al. 2019; Jawahar, Sagot, and Seddah 2019). But on some tasks that require complex and long-distance reasoning, there is still a large gap between the pre-trained models

and human performance, including tasks that require knowledge of broader discourse (Paperno et al. 2016) or coreference (Dasigi et al. 2019), or that require identifying valid reasoning (Niven and Kao 2019). For such tasks in particular, it is interesting to test whether explicit information about linguistic structure can be helpful, how such structure should be injected, and whether it can be helpful in addition to recent language models. Strubell et al. (2018) investigated how to inject syntactic knowledge into a Transformer-style model by applying supervision on self-attention weights. Intriguingly, their model, which they apply to semantic role labeling, benefits from the use of ELMo embeddings, and improves over a benchmark that uses those embeddings but no supervised self-attention.

In this paper, we consider a task that was explicitly designed to require long-distance knowledge, the LAMBADA task (Paperno et al. 2016). LAMBADA is a language modeling task on narrative text passages, where test set data points are chosen to be easily solvable for humans given a larger preceding context of several sentences, but impossible to solve for humans given only a single sentence (see Figure 1 for an example). In the paper that originally introduced LAMBADA, Paperno et al. (2016) report model accuracies of only 7.3%. Since then, more recent models (GPT-2) have improved performance to 63.24% (Radford et al. 2019), while human performance is above 80%. So the LAMBADA dataset clearly has the characteristic described above, with a large gap between pre-trained models and human performance.

We test whether an injection of linguistic knowledge in a similar manner to Strubell et al. (2018), adding supervised self-attention to an existing, non-Transformer model, will improve performance on this complex task. As LAMBADA focuses on narrative texts, we hypothesize that semantic knowledge about the entities mentioned in the passage will be particularly useful for solving the task (i.e., in Figure 1, it is important to know that “you” refers to “Jon” and “he” refers to “Tony” in the first two sentences to make the right prediction). We phrase entity knowledge as knowledge about the coreference chains in the passage. We find that a BiDAF-based (Seo et al. 2017) model trained with coreference as auxiliary supervision achieves state-of-the-

art performance with only a tiny fraction of the parameters of the previous best model, the largest GPT-2 model.

We further analyze the results in more detail, finding evidence that the auxiliary supervision enables the model to better capture coreference information in the text. To provide some insights on how to apply similar techniques to other problems, we experiment with different model variants to test where best to insert the supervision into the system, and we also test different types of linguistic knowledge as supervision signals. Our code is available at <https://github.com/pxch/att-ent>.

2 Background and Related Work

Pre-trained Language Models There has been rapid progress on large-scale pre-trained language models that provide contextualized embeddings for downstream tasks, setting up new state-of-the-art over traditional fixed vector word embeddings like word2vec (Mikolov et al. 2013) or GloVe (Pennington, Socher, and Manning 2014). Peters et al. (2018) first introduced ELMo, a bidirectional-LSTM language model pre-trained on the 1B Word Benchmark, achieving the best results at the time on a broad range of tasks, including reading comprehension, semantic role labeling, coreference resolution, and many others.

With the gaining popularity of the Transformer architecture (Vaswani et al. 2017) in NLP, later efforts shifted focus to pre-trained Transformer models. Radford et al. (2018) introduced GPT, by pre-training a 12-layer Transformer model as a generative language model on the BooksCorpus (Zhu et al. 2015). GPT outperformed previous state-of-the-art on 9 out of the 12 tasks studied. BERT (Devlin et al. 2019) enhanced the GPT model by allowing bidirectional self-attention with a new “masked language model” pre-training objective, and achieved even better results than GPT. GPT-2 (Radford et al. 2019) directly extended GPT with 10x larger pre-training corpus and much more complex model architectures (up to 1.5 billion parameters), and set new state-of-the-art on a number of language modeling tasks. More recent models, including XLNet (Yang et al. 2019) and RoBERTa (Liu et al. 2019), further improved over BERT by introducing new pre-training and optimization strategies.

Linguistic Structure in Pre-trained Models The aforementioned pre-trained models do not explicitly take any linguistic structure into account, as the pre-training objective is to predict the next word, a randomly masked word, or the next sentence. While these pre-trained models achieved state-of-the-art results on many tasks, it is still largely unknown to what extent implicit knowledge of linguistic structures, such as syntactic structure or coreference, contributes to the improvement. Tenney et al. (2019) designed a list of probing tasks to test how well the contextualized representations learned from ELMo / GPT / BERT do on some core NLP pipeline tasks, and found out that contextualized embeddings improve largely on syntactic tasks (like part-of-speech tagging and parsing) but not so much on semantic tasks (like coreference).

Strubell et al. (2018) recently achieved state-of-the-art

Context: "By the way, Elizabeth asked if I'd seen you," Tony lied. He wanted Jon to leave so he could talk with Ezekiel alone. There was something that aunt Casey, Patella and Gabriella had said about Tom that had bothered him ever since meeting Ezekiel earlier that afternoon.

Target sentence: "I'm sure she'll find me," Jon remarked curtly, trying to cut short the conversation with ___ .

Target word: Tony

Figure 1: An example from the LAMBADA dataset.

performance on semantic role labeling by injecting syntactic knowledge into a Transformer-style model. In their LISA model, one self-attention head is guided to learn dependency parsing via an auxiliary supervision signal that encourages each token to only attend to its syntactic parent. They also showed that such syntactically-informed self-attention can be combined with ELMo embeddings to further improve performance over a baseline with only ELMo and self-attention but no auxiliary supervision. In this paper, we want to investigate whether linguistic knowledge of semantic structures can be injected in a similar manner.

The LAMBADA Task Paperno et al. (2016) introduced the LAMBADA dataset, a specially designed language modeling task where each data point is a passage composed of a context (on average 4 to 5 sentences) and a target sentence, and the task is to guess the last word of the target sentence. The data comes from the BooksCorpus (Zhu et al. 2015), and is filtered by human subjects such that it is easy for humans to guess the target word when provided with the whole passage, but impossible to guess given only the target sentence. An example is shown in Figure 1.

Paperno et al. (2016) also reported the results of some standard language models on the task, which are extremely low as none of them reached an accuracy of 1%, while a baseline of selecting a random capitalized word from passage gave an accuracy of 7.3%, indicating the difficulty of the task.

Since then, Chu et al. (2017) proposed to view LAMBADA as reading comprehension, with the context sentences as the context and the target sentence without the last word as the query. The model is then asked to select a word from the context as the answer. Despite the fact that models under this setup will decidedly fail on 19% of the test cases where the target word is not in the context, doing so still greatly improved performance to 49%. Dhingra et al. (2018) improved the number to 55.69% by combining the Gated-Attention Reader (Dhingra et al. 2017) with a “Coref-GRU” layers (where both the previous token and the co-referent antecedent serve as the input to the current GRU cell). Hoang, Wiseman, and Rush (2018) combined the Attention-Sum Reader (Kadlec et al. 2016) with a multi-task objective to track entities in the context, further improving performance to 59.23%. Both these experiments proved the effectiveness of coreference knowledge in the task.

There have also been some efforts in applying

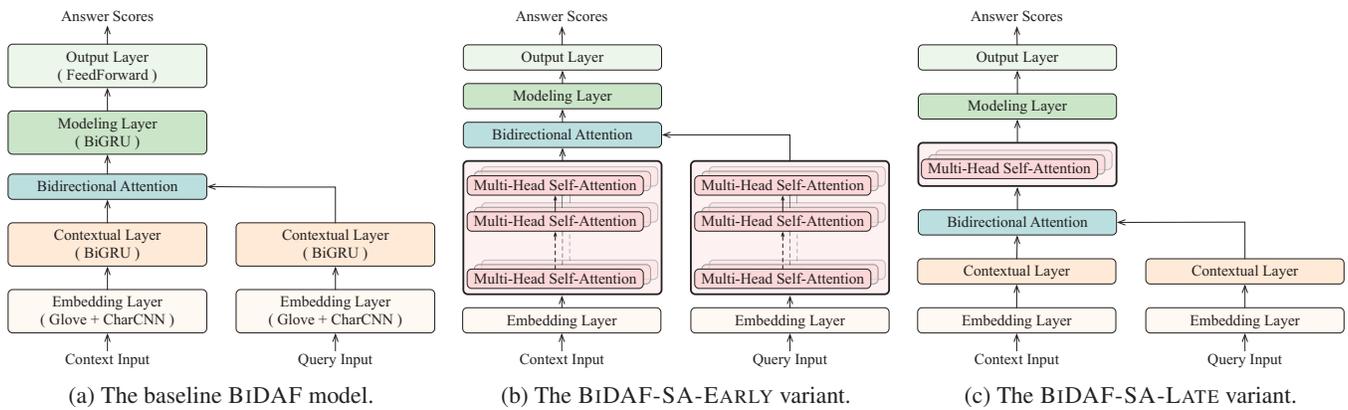


Figure 2: The original BiDAF model (Seo et al. 2017) that we use as a baseline in our experiment, and two variants with a self-attention encoder (in red) being added either as the contextual layer or after the bidirectional attention layer.

Transformer-style models to the task. Dehghani et al. (2019) achieved 56.25% with Universal Transformers. Most recently, Radford et al. (2019) reported 63.24% with the largest GPT-2 model (1.5 billion parameters), setting the current state-of-the-art. Nonetheless, it is still far from the human performance of 86% estimated by Chu et al. (2017).

3 Method

3.1 Task

We adopt the same setting as most of the previous work on LAMBADA, that is, to view the task as reading comprehension. Formally, we concatenate all tokens in the context sentences to get the context input $\mathbf{x} = \{x_1 \dots x_n\}$. We represent all but the last word from the target sentence as the query input $\mathbf{q} = \{q_1 \dots q_m\}$, and the last word of the target sentence as the answer a .

The model computes a probability of being the correct answer for each word in the context $P(x_i | \mathbf{x}, \mathbf{q})$. Because the answer a might occur multiple times in the context, at training time, we sum the probabilities of all correct tokens, and compute the loss as the negative log-likelihood of the summed probability:

$$\mathcal{L}_0 = -\log P(a | \mathbf{x}, \mathbf{q}) = -\log \sum_{i: x_i = a} P(x_i | \mathbf{x}, \mathbf{q}) \quad (1)$$

At test time, a pointer sum mechanism (Kadlec et al. 2016) is used to predict the word type with the highest summed probability among all distinct word types in the context.

3.2 Model

The aim of this paper is to test whether linguistic knowledge of semantic structures can be injected into an existing model via supervised self-attention, and whether the performance of such a model on the LAMBADA task can be matched with the large-scale pre-trained language models (i.e., GPT-2).¹

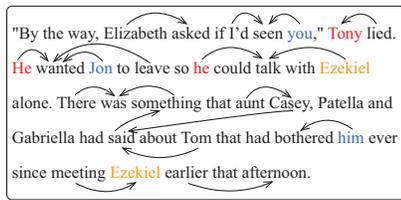
¹We take this route, rather than adding supervision to an existing GPT-2 model, because of the high computational cost of training such a large Transformer model from scratch.

As discussed in Section 2, a range of different reading comprehension models (i.e., Gated-Attention Reader, Attention-Sum Reader) have been tested in the previous work, and they all showed reasonably strong performance on the task (Dhingra et al. 2018; Hoang, Wiseman, and Rush 2018). Therefore, we decide to start with a conventional reading comprehension model, and fuse into it a simpler and shallower stacked self-attention architecture (with fewer layers, fewer attention heads, and smaller hidden size compared to GPT-2). We choose another widely-used reading comprehension model, the BiDAF model (Seo et al. 2017), as our starting point, because BiDAF has consistently outperformed the aforementioned models in many reading comprehension benchmarks.

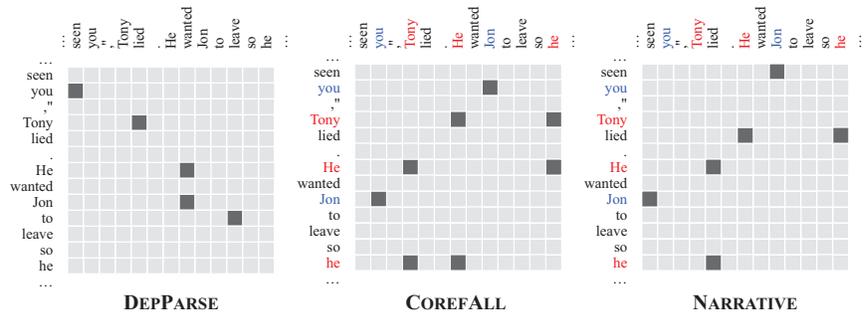
BiDAF Baseline The original BiDAF model, as illustrated in Figure 2a, mainly consists of the following components:

1. **Embedding Layer:** represents each token in the context and the query by a concatenation of Glove embeddings and Character-CNN embeddings.
2. **Contextual Layer:** encodes the context sequence and the query sequence with a bidirectional-LSTM encoder.
3. **Bidirectional Attention Layer:** computes both context-to-query and query-to-context attentions, which are then used to merge the query representations and the context representations to get query-aware vector representations for each context word.
4. **Modeling Layer:** encodes the query-aware context representation with another bidirectional-LSTM encoder to capture the interaction among context words conditioned on the query.
5. **Output Layer:** predicts the probability for each context word being the correct answer with a feed-forward layer followed by a softmax layer.

Our baseline model is mostly the same as the original BiDAF model, except for a few small changes: we substitute the LSTMs for GRUs; we add a Layer Normalization



(a) The dependency parses (arrows) and coreference chains (color-coded) of a context input (same as in Figure 1), which are used to construct different auxiliary supervision signals, as shown on the right.



(b) Examples of 3 different supervision types for self-attention (only showing part of the full matrices due to space limit). Light gray represents 0, and dark gray stands for 1.

Figure 3: An example showing how we construct different types of supervision signals from pre-processed text input.

(Ba, Kiros, and Hinton 2016) after the bidirectional attention layer and after the modeling layer to improve stability (see Section 4).

BiDAF with Self-Attention In order to inject semantic knowledge into the model via supervised attention, we need to fuse a stacked multi-head self-attention encoder into the BiDAF model. Intuitively, there are two options on where the self-attention encoder fits in:

- a) Use the encoder to replace the **Contextual Layer**, as shown in Figure 2b.

This is inspired by the trend of using self-attention encoders to replace traditional RNN-based encoders in many NLP problems. Also, a common practice in using BERT is to first encode raw input with BERT and then pass the output to higher-level task-specific layers, which is similar to what we do here. We name this variant BiDAF-SA-EARLY.

- b) Add the encoder after the **Bidirectional Attention Layer**, as shown in Figure 2c.

This is inspired by the BiDAF++ model (Clark and Gardner 2018), where a standard self-attention layer is added after the bidirectional attention layer to help reason over multiple paragraphs. Here we instead use multi-head self-attention, since applying auxiliary supervision on an attention layer with just one attention head leads to inferior performance in our preliminary experiments. We name this variant BiDAF-SA-LATE.

We also explore another variant that combines the two options, called the BiDAF-SA-BOTH model.

3.3 Auxiliary Supervision for Self Attention

Similar to Strubell et al. (2018), we want to apply auxiliary supervision on a self-attention encoder to guide the model to learn some specific linguistic structure. Our model receives context input, namely the passage, as well as query input, which is the target sentence minus the last word. We focus on investigating auxiliary supervision on the context input,

because the context input, with 4 to 5 sentences on average, should exhibit much richer linguistic structures than the query input, which is a single sentence. To examine what kind of linguistic structures are beneficial to the problem, we experiment with 3 types of supervision signals:

Syntax Supervision: Given the dependency parses for each sentence in the context, we construct the target self-attention weights by putting a weight of 1 from each token to its syntactic head token, and 0 otherwise, as shown in the left column of Figure 3b. We name this type of supervision **DEPPARSE**.

This is similar to the auxiliary supervision used in Strubell et al. (2018), except that we have multiple sentences rather than just one sentence. If an attention head is trained with this syntax supervision, we constrain the self-attention window by sentence boundaries, that is, each token can only attend to other tokens in the same sentence, to make it easier for the model to approach the target self-attention weights.

Coreference Supervision: Given a list of coreference chains from the context (each coreference chain contains a set of mentions that refer to the same entity), we construct the target self-attention weights by putting a weight of 1 between each pair of mention heads in the same coreference chain, and 0 otherwise, as shown in the middle column of Figure 3b. We name it **COREFALL**.

We also test other variants of coreference supervision, namely, guiding the head of each mention to only attend to the head of the most recent previous mention, or to the head of the immediately following mention. We refer to these two variants as **COREFPREV** and **COREFNEXT** respectively.

Narrative Supervision: Since the LAMBADA dataset is built from a corpus of novels, we hypothesize that narrative structures, that is, the sequence of events and their participants, could also be important for predicting the missing word. The interaction between the predicate and arguments of a single event is largely captured by the syntax supervision described above. Therefore, we combine the depen-

	TRAIN	DEV	TEST
Size	709,568	4,869	5,153
% Answer-in-context	100%	82.4%	81.7%
Filtered by human subjects	No	Yes	Yes

Table 1: A brief summary of the LAMBADA dataset.

dependency parses and coreference chains to construct another type of self-attention targets that reflect higher-level narrative knowledge, as shown in the right column of Figure 3b: For each event argument a , we put a weight of 1 between a and all predicates that have an argument that co-refers with a . We name this supervision **NARRATIVE**.

Note that while we require some extra information (i.e., dependency parses and coreference chains, as shown in Figure 3a) to construct the auxiliary supervision signals, we do NOT rely on any gold annotations on either the training set or the test set. All the information can be obtained automatically from running existing NLP tools. We discuss the pre-processing steps more in Section 4.

If any auxiliary supervision is applied to a self-attention head, we compute a loss from the supervision matrix $S \in \mathcal{R}^{n \times n}$ and the attention weights $A \in \mathcal{R}^{n \times n}$ as:

$$\mathcal{L}_s = \frac{1}{k} \sum_{i=1}^n \left[-\log \left(\sum_{j=1}^n A_{ij} * S_{ij} \right) * \sum_{j=1}^n S_{ij} \right] \quad (2)$$

where k is the number of rows in S such that there is at least one non-zero element in the row. To explain, for each token with at least one supervision target, we compute the negative log-likelihood loss against all of its supervision targets, and then get the mean value.

The model is trained in an end-to-end fashion, no matter whether any auxiliary supervision is applied. For a model combined with auxiliary supervisions s_1, s_2, \dots , the total loss being optimized is:

$$\mathcal{L} = \mathcal{L}_0 + \lambda * \sum_i \mathcal{L}_{s_i} \quad (3)$$

where λ is a hyper-parameter to balance between the final prediction loss and the auxiliary losses.

4 Experimental Results

Dataset & Pre-processing When introducing the LAMBADA dataset, Paperno et al. (2016) divided the BooksCorpus randomly into 2 partitions, and only applied the human subjects filtering process to the second half to create the development / test set, while leaving the first half raw data untouched to be the training set. With the reading comprehension setup introduced by Chu et al. (2017), they also constructed a new training set of ~ 1.6 M instances out of the original raw data, by requiring that the last word in the target sentence must exist in the context. Follow-up work (Dhingra et al. 2018; Hoang, Wiseman, and Rush 2018) further filtered the new training set by removing all instances where the target word is a stopword, leaving ~ 700 k instances. We follow the same setup here. A summary is shown in Table 1.

As discussed in Section 3, we also need to get the dependency trees and coreference chains from the data in order to construct the target attention weights for auxiliary supervisions. We use the neural dependency parser and the statistical coreference system from Stanford CoreNLP toolkit (Manning et al. 2014) to pre-process the whole dataset. Further discussion on the choice of pre-processing alternatives will be in Section 5.

Implementation Details We build our models and run all the experiments with AllenNLP (Gardner et al. 2017). For the baseline BiDAF model, we mostly follow the hyper-parameters of the original model: due to space limits, we provide a detailed description of hyper-parameter choices in the supplemental material. We train the model for 10 epochs, and perform early-stopping when the validation accuracy does not increase for two consecutive epochs. We use validation accuracy to select the best epoch, from which then weights are then used for test set evaluation.

For the multi-head self-attention encoders in the BiDAF-SA-* variants, we always use 4 attention heads per layer. For BiDAF-SA-EARLY, we include 4 layers in the stacked self-attention encoder, as preliminary studies show that using only 1 or 2 layers damages performance heavily, while using more than 4 layers gives no significant improvement. For BiDAF-SA-LATE, we only add 1 multi-head self-attention layer, because again, preliminary results show no further gain of using 2 or more layers. In some experiments, we also try replacing the embedding layer with the pre-trained ELMo embeddings (Peters et al. 2018).

We find that performance is very sensitive to the initial random state, possibly due to the fact that there is a large statistical discrepancy between the training set and the development / test sets (because the training set is not filtered by human subjects). We observed a similar effect when we reran existing models from the literature (Dhingra et al. 2018; Hoang, Wiseman, and Rush 2018). Therefore, for each model variant, we train 4 different runs with different random seeds, and report the average and maximum performance (in parentheses in the following tables) across the 4 runs.

Results on LAMBADA We compare our method with the two best previous approaches that did not use large-scale pre-training language models (Dhingra et al. 2018; Hoang, Wiseman, and Rush 2018), and GPT-2 (Radford et al. 2018). Note that we do not compare to other major pre-trained LMs, because BERT (Devlin et al. 2019) and its follow-ups like XLNet (Yang et al. 2019) and RoBERTa (Liu et al. 2019) all used the BooksCorpus as part of the pre-training data. As the LAMBADA task is constructed from the BooksCorpus, BERT and other models would gain an unfair advantage on the task because all the test instances have been accessed by these models during pre-training.

We present our main results in Table 2. Here we first focus on the BiDAF-SA-EARLY model and the COREFALL supervision, because intuitively, knowledge about coreference chains in the passage is likely to be the most beneficial

Models	Accuracy (%)
Dhingra et al. (2018)	55.69
Hoang, Wiseman, and Rush (2018)	59.23
GPT-2 (Radford et al. 2019)	63.24
BiDAF	59.12 (59.54)
BiDAF-SA-EARLY	60.54 (60.88)
BiDAF-SA-EARLY + COREFALL	61.51 (61.94)
BiDAF-SA-EARLY + ELMo	61.38 (61.87)
BiDAF-SA-EARLY + ELMo + COREFALL	63.71 (64.62)

Table 2: Accuracy on the LAMBADA TEST set: average accuracy across 4 runs, with max accuracy in parentheses.

factor for solving the task. Results from other variants are discussed in Section 5.

We see that our BiDAF baseline already performs similarly to the previous best results before GPT-2. Adding the COREFALL auxiliary supervision consistently improves accuracy, with or without ELMo embeddings, but we see larger improvement from COREFALL with ELMo embeddings (~2.3 points) compared to that without ELMo embeddings (~1 point). This confirms our hypothesis that the injection of semantic knowledge via supervised self-attention can be helpful in addition to recent pre-trained language models. Also, the fact that ELMo itself only brings less than 1 point of improvement without COREFALL emphasizes the difficulty of the task. With both ELMo embeddings and COREFALL supervision, we achieve an average accuracy of 63.71% (and the best run achieves 64.62%), outperforming the largest GPT-2 model. This is quite surprising, considering that our model only contains 2.6 million tunable parameters², significantly smaller than the number of parameters in GPT-2 (1.5 billion).

5 Analysis

In this section, we aim to understand why the coreference supervision helps, what is the best possible way to apply auxiliary supervision, and how different types of supervision signals compare.

Does pre-processing quality affect performance? The statistical coreference system from Stanford CoreNLP, from which we construct the supervision signals, is not the current best coreference model in terms of benchmark metrics. We also experimented with a more recent end-to-end neural coreference model (Lee et al. 2017)³, with much higher benchmark scores, as supervision signal for our BiDAF-SA-EARLY + COREFALL model. Surprisingly, this yields inferior performance (61.13% on average, compared to 61.51% with the Stanford coreference results).

We manually examined the output of both coreference systems on some data points, and found that the neural coreference system often produces highly erroneous output, possibly because it is highly optimized towards its news-centric training data, the OntoNotes dataset (Hovy et al.

²Even if we take the frozen ELMo parameters into account, the total number of parameters is only 96 million.

³We use a re-implementation from AllenNLP.

Context: "By the way, Elizabeth asked if I'd seen you," Tony lied. He wanted Jon to leave so he could talk with Ezekiel alone. There was something that aunt Casey, Patella and Gabriella had said about Tom that had bothered him ever since meeting Ezekiel earlier that afternoon.
Target sentence: "I'm sure she'll find me," Jon remarked curtly, trying to cut short the conversation with ____.
Target word: Tony Prediction (Wrong): Ezekiel

Figure 4: Does pre-processing quality affect performance? An example where a wrong coreference chain (color-coded) from a neural coreference system (which we do not use for our experiments) leads to the wrong prediction. A better coreference output from the Stanford coreference system on this example is shown in Figure 3a.

DEV subset	#	no supervision	with supervision
Require coref	21	48.5 ± 4.1	62.0 ± 3.5
Noun	2,006	58.42 ± 0.32	59.48 ± 0.23
Pronoun	2,138	72.31 ± 0.15	76.72 ± 0.53
Not NE	2,848	54.15 ± 0.21	54.90 ± 0.33
PERSON	1,646	72.71 ± 0.08	77.95 ± 0.66

Table 3: Does COREFALL learn coreference? Accuracy on some DEV subsets, BiDAF-SA-EARLY + ELMo model with and without COREFALL supervision. We report average and standard deviation across 4 runs.

2006), while LAMBADA consists of narrative texts. Figure 4 shows an example where a wrong coreference chain from the neural system leads to a wrong prediction. This is the same example as in Figure 3a, which shows the coreference output from the Stanford system. In this example, it is hard to predict the right answer without knowing that “you” refers to “Jon” and “he” refers to “Tony”, both of which are predicted correctly by the Stanford system.

This indicates that a better coreference signal could lead to even better results on the task. We leave it to future work, given some very recent work that further improved coreference performance (Joshi et al. 2019).

Does COREFALL really learn coreference knowledge?

We want to know whether the improvement from COREFALL supervision is because the supervision will actually allow the model to better learn coreference structures, or due to some unknown confounding factors. Chu et al. (2017) manually analyzed 100 random instances from the LAMBADA DEV set to identify the type of reasoning needed for humans to make the right prediction, and found that 21 out of the 100 instances require coreference resolution. We test our models on these 21 instances. To obtain a larger set of instances, we also compare the cases in the DEV set where the target word is a noun to cases where the target is a pronoun, and we compare the cases where the target word is a PERSON to the cases where it is not a named entity (the most common named entity type is PERSON, with all other types occurring very rarely, so we focus on PERSON).

Models	Accuracy (%)
BiDAF-SA-EARLY	60.54 (60.88)
BiDAF-SA-EARLY + COREFALL	61.51 (61.94)
BiDAF-SA-LATE	59.48 (59.58)
BiDAF-SA-LATE + COREFALL	61.19 (61.54)
BiDAF-SA-BOTH	60.88 (61.27)
BiDAF-SA-BOTH + COREFALL (early)	61.72 (62.35)
BiDAF-SA-BOTH + COREFALL (late)	61.54 (61.67)

Table 4: Location of supervision: Accuracy (average of 4 runs, with max in parentheses) on the LAMBADA TEST set for different locations to fuse in the self-attention encoder and to apply auxiliary supervision.

The results in Table 3 show that not only does COREFALL supervision improve accuracy on the 21 instances manually classified as requiring coreference, it also strongly boosts performance on the ‘‘Pronoun’’ and ‘‘PERSON’’ subsets, in comparison to their ‘‘Noun’’ and ‘‘Not NE’’ counterparts. Though not a direct proof, this intuitively supports the claim that the auxiliary supervision does enable the model to better capture coreference information, which is likely to help to reason particularly over pronouns and named entities.

Where should the supervision be applied? Is it more beneficial to apply auxiliary supervision at an earlier stage of the model, i.e., at the contextual layer as in the BiDAF-SA-EARLY model, or at a later stage, i.e., after the bidirectional attention layer as in the BiDAF-SA-LATE model? We compare performance using COREFALL supervision. Also, to disentangle the effect of architectural change, we experiment with the BiDAF-SA-BOTH model with supervision being added at different stages.

Table 4 shows that without supervision, BiDAF-SA-EARLY offers much better results than BiDAF-SA-LATE. Although adding supervision to BiDAF-SA-LATE leads to a larger relative improvement, applying supervision at an earlier stage still leads to a better absolute performance than doing so at a later stage, which is also confirmed by the numbers on BiDAF-SA-BOTH. This is not surprising, as intuitively coreference information about the context input should be beneficial to getting better query-aware context representations, rather than the other way around.⁴

Are other types of supervision also useful? We have so far focused on the COREFALL supervision. In Table 5 we show the results of applying other types of auxiliary supervision.

All other types of auxiliary supervision, except for NARRATIVE, show inferior performance compared to COREFALL. This is as expected. As the LAMBADA task is specifically designed to require broader discourse context, intra-sentence syntactic structure (DEPPARSE) should not play an important role. The COREFPREV and COREFNEXT variants

⁴In all our experiments with BiDAF-SA-EARLY, we apply auxiliary supervision on the 3rd layer (out of all 4 layers). We also test applying supervision on other layers, and find the 3rd layer generally works the best, though the difference is not significant.

Models	Accuracy (%)
BiDAF-SA-EARLY	60.54 (60.88)
BiDAF-SA-EARLY + COREFALL	61.51 (61.94)
BiDAF-SA-EARLY + DEPPARSE	61.06 (61.34)
BiDAF-SA-EARLY + COREFPREV	60.94 (61.71)
BiDAF-SA-EARLY + COREFNEXT	61.27 (61.63)
BiDAF-SA-EARLY + NARRATIVE	61.86 (62.39)

Table 5: Supervision types: Accuracy on the LAMBADA TEST set (average of 4 runs, with max in parentheses) with different types of auxiliary supervision.

of coreference information only provide guidance towards the immediately preceding or following mention in the same coreference chain. Such knowledge will fall short when the reasoning over a long coreference chain is crucial in making the prediction.

The NARRATIVE supervision provides slightly better performance than COREFALL. This is also not surprising, as the NARRATIVE signal is derived from both dependency parses and coreference chains. Theoretically, this type of supervision should capture useful linguistic structures from both COREFALL, which makes the main contribution to the performance improvement, and DEPPARSE, which might offer some additional boost.⁵ We further verify the hypothesis by computing the agreement between the predictions of two models on the DEV set, and find that on average, a run with COREFALL and a run with NARRATIVE agree on 89.3% of all DEV instances, confirming that it is largely the coreference signal that leads to the performance improvement observed with the NARRATIVE supervision.

6 Conclusion

In this paper we have investigated whether the injection of semantic knowledge into an existing model (BiDAF) via supervised self-attention can lead to better performance on tasks requiring complex and long-distance reasoning. On the LAMBADA dataset, where the current best result from GPT-2 (Radford et al. 2019) is still far below human performance, we show that a BiDAF model trained with coreference as auxiliary supervision achieves the new state-of-the-art, while requiring only a tiny fraction of the parameters in GPT-2. We further tested model variants to test where in a BiDAF model it is most useful to add a self-attention layer with supervision, and how different types of linguistic knowledge compared as supervision signal.

This paper takes a first step in explicitly using structural semantic knowledge to inform self-attention, which leads to many interesting future directions. First, we want to test other types of linguistic knowledge, for example, semantic role labeling or AMR parsing (O’Gorman et al. 2018). We also want to see how our current approach can be applied to other tasks, for example, the new QUOREF dataset (Dasigi et al. 2019) that requires resolving coreference among entities

⁵We also briefly test combining multiple supervision signals, and find that COREFALL + NARRATIVE leads to slightly better performance than using NARRATIVE alone, but the difference is not significant as well.

to answer a question. In this paper, we extract supervision signals from existing NLP pipeline tools, and the signals are actually quite noisy (especially on coreference). It would be interesting to see whether such semantic structures can be learned jointly when pre-training language models, with some distant supervision (for example, in a Wikipedia document, tokens with links pointing to the same Wikipedia page should be considered as co-referential mentions).

Acknowledgments

This research was supported by the DARPA AIDA program under AFRL grant FA8750-18-2-0017. We acknowledge the Texas Advanced Computing Center for providing grid resources that contributed to these results, and some results presented in this paper were obtained using the Chameleon testbed supported by the National Science Foundation. We would like to thank the anonymous reviewers for their valuable feedback.

References

- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Chu, Z.; Wang, H.; Gimpel, K.; and McAllester, D. 2017. Broad context language modeling as reading comprehension. In *Proceedings of EACL*, 52–57.
- Clark, C., and Gardner, M. 2018. Simple and effective multi-paragraph reading comprehension. In *Proceedings of ACL*, 845–855.
- Dasigi, P.; Liu, N. F.; Marasovic, A.; Smith, N. A.; and Gardner, M. 2019. Quoref: A reading comprehension dataset with questions requiring coreferential reasoning. In *Proceedings of EMNLP*.
- Dehghani, M.; Gouws, S.; Vinyals, O.; Uszkoreit, J.; and Kaiser, L. 2019. Universal transformers. In *International Conference on Learning Representations*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL*, 4171–4186.
- Dhingra, B.; Liu, H.; Yang, Z.; Cohen, W.; and Salakhutdinov, R. 2017. Gated-attention readers for text comprehension. In *Proceedings of ACL*, 1832–1846.
- Dhingra, B.; Jin, Q.; Yang, Z.; Cohen, W.; and Salakhutdinov, R. 2018. Neural models for reasoning over multiple mentions using coreference. In *Proceedings of NAACL*, 42–48.
- Gardner, M.; Grus, J.; Neumann, M.; Tafjord, O.; Dasigi, P.; Liu, N. F.; Peters, M.; Schmitz, M.; and Zettlemoyer, L. S. 2017. AllenNLP: A deep semantic natural language processing platform.
- Hoang, L.; Wiseman, S.; and Rush, A. 2018. Entity tracking improves cloze-style reading comprehension. In *Proceedings of EMNLP*, 1049–1055.
- Hovy, E.; Marcus, M.; Palmer, M.; Ramshaw, L.; and Weischedel, R. 2006. OntoNotes: The 90% solution. In *Proceedings of NAACL: Short Papers*, 57–60.
- Jawahar, G.; Sagot, B.; and Seddah, D. 2019. What does BERT learn about the structure of language? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3651–3657.
- Joshi, M.; Chen, D.; Liu, Y.; Weld, D. S.; Zettlemoyer, L.; and Levy, O. 2019. SpanBERT: Improving pre-training by representing and predicting spans. *arXiv preprint arXiv:1907.10529*.
- Kadlec, R.; Schmid, M.; Bajgar, O.; and Kleindienst, J. 2016. Text understanding with the attention sum reader network. In *Proceedings of ACL*, 908–918.
- Lee, K.; He, L.; Lewis, M.; and Zettlemoyer, L. 2017. End-to-end neural coreference resolution. In *Proceedings of ACL*, 188–197.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Manning, C.; Surdeanu, M.; Bauer, J.; Finkel, J.; Bethard, S.; and McClosky, D. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of ACL: System Demonstrations*, 55–60.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, 3111–3119.
- Niven, T., and Kao, H.-Y. 2019. Probing neural network comprehension of natural language arguments. In *Proceedings of ACL*, 4658–4664.
- O’Gorman, T.; Regan, M.; Griffith, K.; Hermjakob, U.; Knight, K.; and Palmer, M. 2018. AMR beyond the sentence: the multi-sentence AMR corpus. In *Proceedings of the 27th International Conference on Computational Linguistics*, 3693–3702.
- Paperno, D.; Kruszewski, G.; Lazaridou, A.; Pham, N. Q.; Bernardi, R.; Pezzelle, S.; Baroni, M.; Boleda, G.; and Fernández, R. 2016. The LAMBADA dataset: Word prediction requiring a broad discourse context. In *Proceedings of ACL*, 1525–1534.
- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *Proceedings of EMNLP*, 1532–1543.
- Peters, M.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep contextualized word representations. In *Proceedings of NAACL*, 2227–2237.
- Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. 2018. Improving language understanding with unsupervised learning.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language models are unsupervised multitask learners.
- Seo, M.; Kembhavi, A.; Farhadi, A.; and Hajishirzi, H. 2017. Bidirectional attention flow for machine comprehension. In *International Conference on Learning Representations*.
- Strubell, E.; Verga, P.; Andor, D.; Weiss, D.; and McCallum, A. 2018. Linguistically-informed self-attention for semantic role labeling. In *Proceedings of EMNLP*, 5027–5038.
- Tenney, I.; Xia, P.; Chen, B.; Wang, A.; Poliak, A.; McCoy, R. T.; Kim, N.; Durme, B. V.; Bowman, S.; Das, D.; and Pavlick, E. 2019. What do you learn from context? Probing for sentence structure in contextualized word representations. In *International Conference on Learning Representations*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, 5998–6008.
- Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J.; Salakhutdinov, R.; and Le, Q. V. 2019. XLNet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*.
- Zhu, Y.; Kiros, R.; Zemel, R.; Salakhutdinov, R.; Urtasun, R.; Torralba, A.; and Fidler, S. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of ICCV*, 19–27.