

Structural Embedding of Syntactic Trees for Machine Comprehension

Rui Liu*, Junjie Hu*, Wei Wei*, Zi Yang*, Eric Nyberg

School of Computer Science

Carnegie Mellon University

5000 Forbes Ave, Pittsburgh PA 15213, USA

{ruil, junjieh, weiwei, ziy, ehn}@cs.cmu.edu

Abstract

Deep neural networks for machine comprehension typically utilize only word or character embeddings without explicitly taking advantage of structured linguistic information such as constituency trees and dependency trees. In this paper, we propose *structural embedding of syntactic trees* (SEST), an algorithm framework to utilize structured information and encode them into vector representations that can boost the performance of algorithms for machine comprehension. We evaluate our approach using a state-of-the-art neural attention model on the SQuAD dataset. Experimental results demonstrate that our model can accurately identify the syntactic boundaries of the sentences and extract answers that are syntactically coherent over the baseline methods.

1 Introduction

Reading comprehension such as SQuAD (Rajpurkar et al., 2016) or NewsQA (Trischler et al., 2016) requires identifying a span from a given context, which is an extension to the traditional question answering task, aiming at responding to questions posed by human with natural language (Nyberg et al., 2002; Ferrucci et al., 2010; Liu, 2017; Yang, 2017). Many works have been proposed to leverage deep neural networks for such question answering tasks, most of which involve learning the query-aware context representations (Dhingra et al., 2016; Seo et al., 2017; Wang and Jiang, 2016; Xiong et al., 2017). Although deep learning based methods demonstrated great potentials for question answering, none of them take syntactic information of the sentences such as con-

stituency tree and dependency tree into consideration. Such techniques have been proven to be useful in many natural language understanding tasks in the past and illustrated noticeable improvements such as the work by (Rajpurkar et al., 2016). In this paper, we adopt similar ideas but apply them to a neural attention model for question answering.

The constituency tree (Manning et al., 1999) of a sentence defines internal nodes and terminal nodes to represent phrase structure grammars and the actual words. Figure 1 illustrates the constituency tree of the sentence “the architect or engineer acts as the project coordinator”. Here, “the architect or engineer” and “the project coordinator” are labeled as noun phrases (“NP”), which is critical for answering the question below. Here, the question asks for the name of certain occupation that can be best answered using a noun phrase. Utilizing the knowledge of a constituency relations, we can reduce the size of the candidate space and help the algorithm to identify the correct answer.

Whose role is to design the works, prepare the specifications and produce construction drawings, administer the contract, tender the works, and manage the works from inception to completion?

On the other hand, a dependency tree (Manning et al., 1999) is constructed based on the dependency structure of a sentence. Figure 2 displays the dependency tree for sentence

The Annual Conference, roughly the equivalent of a diocese in the Anglican Communion and the Roman Catholic Church or a synod in some Lutheran denominations such as the Evangelical Lutheran Church in America, is the basic unit of organization within the UMC.

*Authors contributed equally to this work.




Figure 1: The constituency tree of context “the architect or engineer acts as the project coordinator”

“The Annual Conference” being the subject of “the basic unit of organization within the UMC” provides a critical clue for the model to skip over a large chunk of the text when answering the question “What is the basic unit of organization within the UMC”. As we show in the analysis section, adding dependency information dramatically helps identify dependency structures within the sentence, which is otherwise difficult to learn.

In this paper, we propose Structural Embedding of Syntactic Trees (SEST) that encode syntactic information structured by constituency tree and dependency tree into neural attention models for the question answering task. Experimental results on SQuAD dataset illustrates that the syntactic information helps the model to choose the answers that are both succinct and grammatically coherent, which boosted the performance on both qualitative studies and numerical results. Our focus is to show adding structural embedding can improve baseline models, rather than directly compare to published SQuAD results. Although the methods proposed in the paper are demonstrated using syntactic trees, we note that similar approaches can be used to encode other types of tree structured information such as knowledge graphs and ontology relations.

2 Methodology

The general framework of our model is illustrated in Figure 3. Here the input of the model is the embedding of the context and question while the output is two indices *begin* and *end* which indicate the begin and end indices of the answer in the context space.

The input of the model contains two parts: the word/character model and the syntactic model.

The shaded portion of our model in Figure 3 represents the encoded syntactic information of both context and question that are fed into the model. To gain an insight of how the encoding works, consider a sentence which syntactic tree consists of four nodes (o_1, o_2, o_3, o_4). A specific word is represented to be a sequence of nodes from its leave all the way to the root. We cover how this process work in detail in Section 3.1.1 and 3.1.2. Another input that will be fed into deep learning model is the embedding information for words and characters respectively. There are many ways to convert words in a sentence into a high-dimensional embedding. We choose GloVe (Pennington et al., 2014b) to obtain a pre-trained and fixed vector for each word. Instead of using a fixed embedding, we use Convolutional Neural Networks (CNN) to model character level embedding, which values can be changed during training (Kim, 2014). To integrate both embeddings into the deep neural model, we feed the concatenation of them for the question and the context to be the input of the model.

The inputs are processed in the embedding layer to form more abstract representations. Here we choose a multi-layer bi-directional Long Short Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997) to obtain more abstract representations for words in the contexts and questions.

After that, we employ an attention layer to fuse information from both the contexts and the questions. Various matching mechanisms using attentions have been extensively studied for machine comprehension tasks (Xiong et al., 2017; Seo et al., 2017; Wang et al., 2016; Wang and Jiang, 2016). We use the Bi-directional Attention flow model (Seo et al., 2017) which performs context-to-question and question-to-context attentions in both directions. The context-to-question attention signifies which question words are most relevant to each context word. For each context word, the attention weight is first computed by a softmax function with question words, and the attention vector of each context word is then computed by a weighted sum of the question words’ embeddings obtained from the embedding layer. The question-to-context attention summarizes a context vector by performing a soft attention with context words given the question. We then represent each context word as the concatenation of the embedding vector obtained from the embedding layer, the atten-




Figure 2: Partial dependency parse tree of an example context “The Annual Conference, roughly the equivalent of a diocese in the Anglican Communion and the Roman Catholic Church or a synod in some Lutheran denominations such as the Evangelical Lutheran Church in America, is the basic unit of organization within the UMC.”

tion vector obtained from the context-to-question attention and the context vector obtained from the question-to-context attention. We then feed the concatenated vectors to a stacked bi-directional LSTM with two layers to obtain the final representations of context words. We note that our proposed structural embedding of syntactic trees can be easily applied to any attention approaches mentioned above.

For the machine comprehension task in this paper, the answer to the question is a phrase in the context. In the output layer, we use two softmax functions over the output of the attention layer to predict the begin and end indices of the phrase in the context.

3 Structural Embedding of Syntactic Tree

We detail the procedures of two alternative implementations of our methods: the Structural Embedding of Constituency Trees model (SECT) and the Structural Embedding of Dependency Trees model (SEDT). We assume that the syntactic information has already been generated in the pre-processing step using tools such as the Stanford CoreNLP (Manning et al., 2014).

3.1 Syntactic Sequence Extraction

We first extract a *syntactic collection* $\mathcal{C}(p)$ for each word p , which consists of a set of nodes $\{o_1, o_2, \dots, o_{d-1}, o_d\}$ in the syntactic parse tree \mathcal{T} . Each node o_i can be a word, a grammatical category (e.g., part-of-speech tagging), or a dependency link label, depending on the type of syntactic tree we use. To construct syntactic embeddings, the first thing we need to do is to define a specific *processing order* \mathcal{A} over the syntactic collection $\mathcal{C}(p)$, in which way we can extract a *syntactic sequence* $\mathcal{S}(p)$ for the word p .

3.1.1 Structural Embedding of Constituency Trees (SECT)

The constituency tree is a syntactic parse tree constructed by phrase structure grammars (Manning et al., 1999), which defines the way to hierarchically construct a sentence from words in a bottom-up manner based on constituency relations. Words in the contexts or the questions are represented by leaf nodes in the constituency tree while the non-terminal nodes are labeled by categories of the grammar. Non-terminal nodes summarize the grammatical function of the sub-tree. Figure 1 shows an example of the constituency tree with “the architect or engineer” being annotated as a noun phrase (NP).

A path originating from the leaf node to the root node captures the syntactic information in the constituency tree in a hierarchical way. The higher the node is, the longer span of words the sub-tree of this node covers. Hence, to extract the syntactic sequence $\mathcal{S}(p)$ for a leaf node p , it is reasonable to define the processing order $\mathcal{A}(p)$ from the leaf p all the way to its root. For example, the syntactic sequence for the phrase “the project coordinator” in Figure 1 is detected as (NP, PP, VP, S). In practice, we usually take the last hidden units of Bi-directional encoding mechanisms such as Bi-directional LSTM to represent the sequence state, as is indicated in Figure 4 (a). We set a window size to limit the amount of information that is used in our models. For example, if we choose the window size as 2, then the syntactic sequence becomes (NP, PP). This process is introduced for both performance and memory utilization consideration, which is discussed in detail in Section 4.5.

In addition, a non-terminal node at a particular position in the syntactic sequence defines the begin and end indices of a phrase in the context. By measuring the similarity between syntactic se-




Figure 3: Model Framework. The neural network for training and testing is built by components with solid lines, which includes the embedding layer, attention layer, and output layer. The shaded area highlights the part of the framework that involves syntactic information. Components with dashed lines is an example to illustrate how syntactic information is decoded. Here a sentence is decomposed into a syntactic tree with four nodes and the syntactic information for a specific word is recorded as the path from its position in the syntactic tree to the root, i.e. (o_1, o_2, o_3) in this case.

quences $\mathcal{S}(p)$ extracted for each word p of both the question and the context, we are able to locate the boundaries of the answer span. This is done in the attention layer shown in Figure 3.

3.1.2 Structural Embedding of Dependency Trees (SED_T)

The dependency tree is a syntactic tree constructed by dependency grammars (Manning et al., 1999), which defines the way to connect words by directed links that represent dependencies. A dependency link is able to capture both long and short distance dependencies of words. Relations on links vary in their functions and are labeled with different categories. For example, in the dependency tree plotted in Figure 2, the link from “unit” to “Conference” indicates that the target node is a nominal subject (i.e. NSUBJ) of the source node.

The syntactic collection $\mathcal{C}(p)$ for dependency tree is defined as p ’s children, each represented by its word embedding concatenated with a vec-

tor that uniquely identifies the dependency label. The processing order $\mathcal{A}(p)$ for dependency tree is then defined to be the dependent’s original order in the sentence.

Take the word “unit” as an example, we encode the dependency sub-tree using a Bi-directional LSTM, as indicated in Figure 4 (b). In such as a sub-tree, since children are directly linked to the root, they are position according to the original sequence in the sentence. Similar to the syntactic encoding of C-Tree, we take the last hidden states as its embedding.

Similar to SECT, we use a window of size l to limit the amount of syntactic information for the learning models by choosing only the l -nearest dependents, which is again reported in Section 4.5.

3.2 Syntactic Sequence Encoding

Similar to previous work (Cho et al., 2014; Kim, 2014), we use a neural network to encode a variable-length syntactic sequence into a fixed-length vector representation. The encoder can be a Recurrent Neural Network (RNN) or a Convolutional Neural Network (CNN) that learns a structural embedding for each node such that embedding of nodes under similar syntactic trees are close in their embedding space.

We can use a Bi-directional LSTM as our RNN encoder, where the hidden state \mathbf{v}_t^p is updated according to Eq. 1. Here \mathbf{x}_t^p is the t^{th} node in the syntactic sequence of the word p , which is a vector that uniquely identifies each syntactic node. We obtain the structural embedding of the given word p , $\mathbf{v}_{\text{Bi-LSTM}}^p = \mathbf{v}_T^p$ to be the final hidden state.

$$\mathbf{v}_t^p = \text{Bi-LSTM}(\mathbf{v}_{t-1}^p, \mathbf{x}_t^p) \quad (1)$$

Alternatively, we can also use CNN to obtain embeddings from a sequence of syntactic nodes. We denote l as the length of the filter of the CNN encoder. We define $\mathbf{x}_{i:i+l}^p$ as the concatenation of the vectors from \mathbf{x}_i^p to \mathbf{x}_{i+l-1}^p within the filter. The i^{th} element in the j^{th} feature map can be obtained in Eq. 2. Finally we obtain the structural embedding of the given word p by $\mathbf{v}_{\text{CNN}}^p$ in Eq. 3.

$$\mathbf{c}_{i,j}^p = f(\mathbf{w}_j \cdot \mathbf{x}_{i:i+l-1}^p + b_j) \quad (2)$$

$$\mathbf{v}_{\text{CNN}}^p = \max_{\text{row}}(\mathbf{c}^p) \quad (3)$$

where \mathbf{w}_j and b_j are the weight and bias of the j^{th} filter respectively, f is a non-linear activation function and $\max_{\text{row}}(\cdot)$ takes the maximum value along rows in a matrix.




Figure 4: Two examples are used to illustrate how the syntactic information is encoded for SECT and SEDT respectively. Take Bi-directional LSTM as examples, where \mathbf{x} is a vector such as word embedding, \mathbf{v} and \mathbf{u} are the outputs of the forward and backward LSTMs respectively. For SECT, we encode the syntactic sequence (NP, PP, VP) for the word “coordinator” in Figure 1. We use fixed vectors for syntactic tags (e.g., NP, PP and VP), initialized with multivariate normal distribution. The final representation for the target word “coordinator” can be represented as the concatenation $[\mathbf{E}\mathbf{w}; \mathbf{u}_0; \mathbf{v}_4]$, where $\mathbf{E}\mathbf{w}$ is the word embedding for “coordinator” that is 100 dimensions in our experiments and each of the encoded vector \mathbf{u} and \mathbf{v} can be 30 dimensional. For SEDT, we encode the word “unit” in Figure 2 with its dependent nodes including “Conference”, “is”, “the”, “basic”, “organization”, ordered by their positions in the original sentence. Each word is represented with its word embedding. Similar to SECT, the final representation is the concatenation $[\mathbf{E}\mathbf{w}; \mathbf{u}_0; \mathbf{v}_6]$, which will be sent to the input layer of a neural network.

4 Experiments

We conducted systematic experiments on the SQuAD dataset (Rajpurkar et al., 2016). We compared our methods against Bi-Directional Attention Flow (BiDAF), as well as the SECT models described in Section 3.

4.1 Preprocessing

A couple of preprocessing steps is performed to ensure that the deep neural models get the correct input. We segmented context and questions into sentences by using NLTK’s Punkt sentence segmenter¹. Words in the sentences were then converted into symbols by using PTB Tokenizer². Syntactic information including POS tags and syntactic trees were acquired by Stanford CoreNLP utilities (Manning et al., 2014). For the parser, we collected constituent relations and dependency relations for each word by using tree annotation and enhanced dependencies annotation respectively. To generate syntactic sequence, we removed sequences whose first node is a punctuation (“\$”, “:”, “#”, “:”, “ ””, “ “ ””, “ “ ””, “,”). To use dependency labels, we removed all the subcategories (e.g., “nmod:poss” \Rightarrow “nmod”).

¹<http://www.nltk.org/api/nltk.tokenize.html>

²<http://nlp.stanford.edu/software/tokenizer.shtml>

4.2 Experiment Setting

We run our experiments on a machine that contains a single GTX 1080 GPU with 8GB VRAM. All of the models being compared have the same settings on character embedding and word embedding. As introduced in Section 2, we use a variable character embedding with a fixed pre-trained word embedding to serve as part of the input into the model. The character embedding is implemented using CNN with a one-dimensional layer consists of 100 units with a channel size of 5. It has an input depth of 8. The max length of SQuAD is 16 which means there are a maximum 16 words in a sentence. The fixed word embedding has a dimension of 100, which is provided by the GloVe data set (Pennington et al., 2014a). The settings for syntactic embedding are slightly different for each model. The BiDAF model does not deal with syntactic information. The POS model contains syntactic information with 39 different POS tags that serve as both input and output. For SECT and SEDT the input of the model has a size of 8 with 30 units to be output. Both of them has a maximum length size that is set to be 10 and 20 respectively, which values will be further discussed in Section 4.5. They also have two different ways to encode the syntactic information as indi-

cated in Section 3: LSTM and CNN. We apply the same sets of parameters when we experiment them with the two models. We report the results on the SQuAD development set and the blind test set.

4.3 Predictive Performance

We first compared the performance of single models between the baseline approach BiDAF and the proposed SEST approaches, including SEPOS, SECT-LSTM, SECT-CNN, SEDT-LSTM, and SEDT-CNN, on the development dataset of SQuAD. For each model, we conducted 5 different single experiments and evaluated them using two metrics: “Exact match” (EM), which calculates the ratio of questions that are answered correctly by strict string comparison, and the F1 score, which calculates the harmonic mean of the precision and recall between predicted answers and ground true answers at the character level. As shown in Table 1, we reported the maximum, the mean, and the standard deviation of EM and F1 scores across all single runs for each approach, and highlighted the best model using bold font. SECT-LSTM is the second best method, which confirms the predictive powers of different types of syntactic information. We could see that SEDT-LSTM model outperforms the baseline method and other proposed methods in terms of both EM and F1. Another observation is that our proposed models achieve higher relative improvements in EM scores than F1 scores over the baseline methods, providing the evidence that syntactic information can accurately locate the boundaries of the answer.

Moreover, we found that both SECT-LSTM and SEDT-LSTM have better performance than their CNN counterparts, which suggests that LSTM can more effectively preserve the syntactic information. As a result, we conducted further analysis of only SECT-LSTM and SEDT-LSTM models in the subsequent subsections and drop the suffix “-LSTM” for abbreviation. We built an ensemble model from the 5 single models for the baseline method BiDAF and our proposed methods SEPOS, SECT-LSTM, and SEDT-LSTM. The ensemble model choose the answer with the highest sum of confidence scores among the 5 single models for each question. We compared these models on both the development set and official test set and reported the results in Table 2. We found that the models have higher performance on

the test set than the development set, which coincides with the previous results on the same data set (Seo et al., 2017; Xiong et al., 2017).

4.4 Contribution of Syntactic Sequence

To take a closer look at how syntactic sequences affect the performance, we removed the character/word embedding from our model seen in Figure 3 and conducted experiments based on the syntactic input alone. In particular, we are interested in two aspects related to syntactic sequences: First, the ability to predict answers of questions of syntactic sequences compared to complete random sequences. Second, the amount of impacts brought by our proposed ordering introduced in Section 3.1.1 and Section 3.1.2 compared to random ordering.

We compared the performance of the models using syntactic information along in their original order (i.e. SECT-Only and SEDT-Only) against their counterparts with the same syntactic tree nodes but with randomly shuffled order (i.e. SECT-Random-Order and SEDT-Random-Order) as well as the baselines with randomly generated tree nodes (i.e. SECT-Random and SEDT-Random). Here we choose the length of window size to be 10. The predictive results in terms of EM and F1 metrics are reported in Table 3. From the table we see that both the ordering and the contents of the syntactic tree are important for the models to work properly: constituency and dependency trees achieved over 20% boost on performance compared to the randomly generated ones and our proposed ordering also out-performed the random ordering. It also worth mentioning that the ordering of dependency trees seems to have less impact on the performance compared to that of the constituency trees. This is because sequences extracted from constituency trees contain hierarchical information, which ordering will affect the output of the model significantly. However, sequences extracted from dependency trees are all children nodes, which are often interchangeable and don’t seem to be affected by ordering much.

4.5 Window Size Analysis

As we have mentioned in the earlier sections, limiting the window size is an important technique to prevent excessive usage on VRAM. In practice, we found that limiting the window size also benefits the performance of our models. In Table 4 we compared the predictive performance of SECT

Method	Single				Ensemble	
	EM		F1		EM	F1
	Max	Mean (\pm SD)	Max	Mean (\pm SD)		
BiDAF	67.10	66.92 (\pm 0.23)	76.92	76.79 (\pm 0.08)	70.97	79.53
SEPOS	67.65	66.05 (\pm 2.94)	77.25	75.80 (\pm 2.65)	71.46	79.70
SECT-LSTM	67.91	67.65 (\pm 0.31)•	77.47	77.19 (\pm 0.21) •	71.76	80.09
SEDT-LSTM	68.13	67.89 (\pm0.10)•	77.58	77.42 (\pm0.19) •	72.03	80.28
SECT-CNN	67.29	64.04 (\pm 4.28)	76.91	73.99 (\pm 3.89)	69.70	78.49
SEDT-CNN	67.88	66.53 (\pm 1.91)	77.27	76.21 (\pm 1.67)	71.58	79.80

Table 1: Performance comparison on the development set. Each setting contains five runs trained consecutively. Standard deviations across five runs are shown in the parenthesis for single models. Dots indicate the level of significance.

Method	Single		Ensemble	
	EM	F1	EM	F1
BiDAF	67.69	77.07	72.33	80.33
SECT-LSTM	68.12	77.21	72.83	80.58
SEDT-LSTM	68.48	77.97	73.02	80.84

Table 2: Performance comparison on the official blind test set. Ensemble models are trained over the five single runs with the identical network and hyper-parameters.

Method	EM	F1
SECT-Random	5.64	12.85
SECT-Random-Order	30.04	39.98
SECT-Only	34.21	44.53
SEDT-Random	0.92	8.82
SEDT-Random-Order	31.82	43.65
SEDT-Only	32.96	44.37

Table 3: Performance comparisons of models with only syntactic information against their counterparts with randomly shuffled node sequences and randomly generated tree nodes using the SQuAD Dev set

and SEDT models by varying the length of their window sizes from 1 to maximum on the development set. In general the results illustrate that performances of the models increase with the length of the window. However, we found that for SECT model, its mean performance reached the peak while standard deviations narrowed when window size reaches 10. We also observed that larger window size does not generate predictive results that is as good as the one with window size set to 10. This suggests that there exists an optimal window size for the constituency tree. One possible ex-

Method	Len	EM	F1
SECT	1	65.58 (\pm 2.58)	75.31 (\pm 2.39)
	5	65.74 (\pm 3.77)	75.48 (\pm 3.39)
	10	67.51 (\pm 0.34)	77.14 (\pm 0.39)
	Max	67.48 (\pm 0.33)	77.09 (\pm 0.45)
SEDT	1	66.23 (\pm 2.5)	73.85 (\pm 2.22)
	10	67.39 (\pm 0.09)	76.93 (\pm 0.21)

Table 4: Performance means and standard deviations of different window sizes on the development set.

planation is increasing the window size leads to the increase in the number of syntactic nodes in the extracted syntactic sequence. Although subtrees might be similar between context and question, it is very unlikely that the complete trees are the same. Because of that, allowing the syntactic sequence to extend beyond the certain heights will introduce unnecessary noise into the learned representation, which will compromise the performance of the models. Similar conclusion holds for the SEDT model, which has an improved performance and decreased variance with the window size is set to 10. We did not perform experiments with window size beyond 10 for SEDT since it will consume VRAM that exceeds the capacity of our computing device.

4.6 Overlapping Analysis

To further understand the performance benefits of incorporating syntactic information into the question answering problem, we can take a look at the questions on which models disagree. Figure 5 is the Venn Diagram on the questions that have been corrected identified by SECT, SEDT and the baseline BiDAF model. Here we see that the vast

Question	BiDAF	SECT
Whose role is to design the works, prepare the specifications and produce construction drawings, administer the contract, tender the works, and manage the works from inception to completion?	[the architect or engineer] _{NP} [acts as [the project coordinator]] _{NP}] _{VP}	[the architect or engineer] _{NP}
What did Luther think the study of law meant?	[represented [uncertainty] _{NP}] _{VP}	[uncertainty] _{NP}
What caused the dynamos to be burnt out?	[the powerful high frequency currents] _{NP} [set up [in [them] _{NP}] _{PP}] _{VP}	[powerful high frequency currents] _{NP}

Table 5: Questions that are correctly answered by SECT but not BiDAF

Question	Context	BiDAF	SEDT
In the layered model of the Earth, the mantle has two layers below it. What are they?	These advances led to the development of a layered model of the Earth, with a crust and lithosphere on top, the mantle below (separated within itself by seismic discontinuities at 410 and 660 kilometers), and the outer core and inner core below that.	seismic discontinuities at 410 and 660 kilometers), and the outer core and inner core	the outer core and inner core
What percentage of farmland grows wheat?	More than 50% of this area is sown for wheat, 33% for barley and 7% for oats.	33%	50%
What is the basic unit of organization within the UMC?	The Annual Conference, roughly the equivalent of a diocese in the Anglican Communion and the Roman Catholic Church or a synod in some Lutheran denominations such as the Evangelical Lutheran Church in America, is the basic unit of organization within the UMC.	Evangelical Lutheran Church in America	The Annual Conference

Table 6: Questions that are correctly answered by SEDT but not BiDAF




Figure 5: Venn Diagram on the number of correct answers predicted by BiDAF, SECT and SEDT

majority of the correctly answered questions are shared across all three models. The rest of them indicates questions that models disagree and are distributed fairly evenly.

To understand the types of the questions that

syntactic models can do better, we extracted three questions that were correctly answered by SECT and SEDT but not the baseline model. In Table 5, all of the three questions are “Wh-questions” and expect the answer of a noun phrase (NP). Without knowing the syntactic information, BiDAF answered questions with unnecessary structures such as verb phrases (vp) (e.g. “acts as ...”, “represented ...”) or prepositional phrases (pp) (e.g. “in ...”) in addition to NPs (e.g. “the architect engineer”, “uncertainty” and “powerful high frequency currents”) that normal human would answer. For that reason, answers provided by BiDAF failed the exact match although its answers are semantically equivalent to the ones provided by SECT. Having incorporated constituency information provided an huge advantage in inferring the answers that are most natural for a human.

The advantages of using the dependency tree in our model can be illustrated using the questions in Table 6. Here again we listed the ones that are correctly identified by SEDT but not BiDAF. As we can see that the answer provided by BiDAF for first question broke the parenthesis incorrectly, this problem that can be easily solved by utilizing dependency information. In the second example, BiDAF failed to identify the dependency structures between “50%” and the keyword being asked “wheat”, which resulted in an incorrect answer that has nothing to do with the question. SEDT, on the other hand, answered the question correctly. In the third question, the key to the answer is to correctly identify the subject of question phrase “is the basic unit of organization”. Using the dependency tree as illustrated in Figure 2, SEDT is able to identify the subject phrase correctly, namely “The Annual Conference”. However, BiDAF failed to answer the question correctly and selected a noun phrase as the answer.

5 Related Work

Reading Comprehension. Reading comprehension is a challenging task in NLP research. Since the release of the SQuAD data set, many works have been done to construct models on this massive question answering data set. Rajpurkar et. al. are among the first authors to explore the SQuAD. They used logistic regression with pos tagging information (Rajpurkar et al., 2016) and provided a strong baseline for all subsequent models. A steep improvement was given by the RaSoR model (Lee et al., 2016) which utilized recurrent neural networks to consider all possible subphrases of the context and evaluated them one by one. To avoid comparing all possible candidates and to improve the performance, Match-LSTM (Wang and Jiang, 2016) was proposed by using a pointer network (Vinyals et al., 2015) to extract the answer span from the context. The same idea was taken to the BiDAF (Seo et al., 2017) model by introducing a bi-directional attention mechanism. Despite the above-mentioned strong models for the machine comprehension task, none of them considers syntactic information into their prediction models.

Representations of Texts and Words. One of the main issues in reading comprehension is to identify the latent representations of texts and words (Cui et al., 2016; Lee et al., 2016; Wang et al., 2016; Xiong et al., 2017; Yu et al., 2016).

Many pre-trained libraries such as word2vec (Mikolov et al., 2013) and Glove (Pennington et al., 2014a) have been widely used to map words into a high dimensional embedding space. Another approach is to generate embeddings by using neural networks models such as Character Embedding (Kim, 2014) and Tree-LSTM (Tai et al., 2015). One thing that worth mentioning is that although Tree-LSTM does utilize syntactic information, it targets at the phrases or sentences level embedding other than the word level embedding we have discussed in this paper. Many machine comprehension models include both pre-trained embeddings and variable embeddings that can be changed through a training stage (Seo et al., 2017).

6 Conclusion

In this paper, we proposed methods to embed syntactic information into the deep neural models to improve the accuracy of our model in the machine comprehension task. We formally defined our SEST framework and proposed two instances to it: the structural embedding of constituency trees (SECT) and the structural embedding of dependency trees (SEDT). Experimental results on SQuAD data set showed that our proposed approaches outperform the state-of-the-art BiDAF model, proving that the proposed embeddings play a significant part in correctly identifying answers for the machine comprehension task. In particular, we found that our model can perform especially well on exact match metrics, which requires syntactic information to accurately locate the boundaries of the answers. Similar approaches can be used to encode other tree structures such as knowledge graphs and ontology relations.

This work opened several potential new lines of research: 1) In the experiments of our paper we utilized the BiDAF model to retrieve answers from the context. Since there are no structures in the BiDAF models to specifically optimize for syntactic information, an attention mechanism that is designed for to utilize syntactic information should be studied. 2) Another direction of research is to incorporate SEST with deeper neural networks such as VD-CNN (Conneau et al., 2017) to improve learning capacity for syntactic embedding. 3) Tree structured information such as knowledge graphs and ontology structure should be studied and improve question answering tasks using similar techniques to the ones proposed in the paper.

References

- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP*. pages 1724–1734.
- Alexis Conneau, Holger Schwenk, Loïc Barrault, and Yann Lecun. 2017. Very deep convolutional networks for text classification. In *EACL*.
- Yiming Cui, Zhipeng Chen, Si Wei, Shijin Wang, Ting Liu, and Guoping Hu. 2016. Attention-over-attention neural networks for reading comprehension. *arXiv preprint arXiv:1607.04423*.
- Bhuwan Dhingra, Hanxiao Liu, William W Cohen, and Ruslan Salakhutdinov. 2016. Gated-attention readers for text comprehension. *arXiv preprint arXiv:1606.01549*.
- David Ferrucci, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya A Kalyanpur, Adam Lally, J William Murdock, Eric Nyberg, John Prager, et al. 2010. Building watson: An overview of the deepqa project. *AI magazine* 31(3):59–79.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *EMNLP*. pages 1746–1751.
- Kenton Lee, Tom Kwiatkowski, Ankur Parikh, and Dipanjan Das. 2016. Learning recurrent span representations for extractive question answering. *arXiv preprint arXiv:1611.01436*.
- Rui Liu. 2017. *A Phased Ranking Model for Information Systems*. Ph.D. thesis, Carnegie Mellon University.
- Christopher D Manning, Hinrich Schütze, et al. 1999. *Foundations of statistical natural language processing*, volume 999. MIT Press.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](http://www.aclweb.org/anthology/P/P14/P14-5010). In *Association for Computational Linguistics (ACL) System Demonstrations*. pages 55–60. <http://www.aclweb.org/anthology/P/P14/P14-5010>.
- 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.
- Eric Nyberg, Teruko Mitamura, Jaime G Carbonell, Jamie Callan, Kevyn Collins-Thompson, Krzysztof Czuba, Michael Duggan, Laurie Hiyakumoto, N Hu, Yifen Huang, et al. 2002. The javelin question-answering system at trec 2002. In *TREC*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014a. Glove: Global vectors for word representation. In *EMNLP*. Association for Computational Linguistics, pages 1532–1543.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014b. [Glove: Global vectors for word representation](http://www.aclweb.org/anthology/D14-1162). In *Empirical Methods in Natural Language Processing (EMNLP)*. pages 1532–1543. <http://www.aclweb.org/anthology/D14-1162>.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *EMNLP*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *ICLR*.
- Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. Improved semantic representations from tree-structured long short-term memory networks. In *ACL*.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2016. Newsqa: A machine comprehension dataset. *arXiv preprint arXiv:1611.09830*.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *NIPS*. pages 2692–2700.
- Shuohang Wang and Jing Jiang. 2016. Machine comprehension using match- lstm and answer pointer. *arXiv preprint arXiv:1608.07905*.
- Zhiguo Wang, Haitao Mi, Wael Hamza, and Radu Florian. 2016. Multi-perspective context matching for machine comprehension. *arXiv preprint arXiv:1612.04211*.
- Caiming Xiong, Victor Zhong, and Richard Socher. 2017. Dynamic coattention networks for question answering. In *ICLR*.
- Zi Yang. 2017. *Analytics Meta Learning*. Ph.D. thesis, Carnegie Mellon University.
- Yang Yu, Wei Zhang, Kazi Hasan, Mo Yu, Bing Xiang, and Bowen Zhou. 2016. End-to-end answer chunk extraction and ranking for reading comprehension. *arXiv preprint arXiv:1610.09996*.