

On Incorporating Structural Information to improve Dialogue Response Generation

Nikita Moghe[‡], Priyesh Vijayan², Balaraman Ravindran^{3,4}, and Mitesh M. Khapra^{3,4}

¹School of Informatics, University of Edinburgh

²School of Computer Science, McGill University and Mila

³Indian Institute of Technology Madras

⁴Robert Bosch Centre for Data Science and Artificial Intelligence (RBC-DSAI),
Indian Institute of Technology Madras

nikita.moghe@ed.ac.uk priyesh.vijayan@mail.mcgill.ca

Abstract

We consider the task of generating dialogue responses from background knowledge comprising of domain specific resources. Specifically, given a conversation around a movie, the task is to generate the next response based on background knowledge about the movie such as the plot, review, Reddit comments *etc.* This requires capturing structural, sequential, and semantic information from the conversation context and background resources. We propose a new architecture that uses the ability of BERT to capture deep contextualized representations in conjunction with explicit structure and sequence information. More specifically, we use (i) Graph Convolutional Networks (GCNs) to capture structural information, (ii) LSTMs to capture sequential information, and (iii) BERT for the deep contextualized representations that capture semantic information. We analyze the proposed architecture extensively. To this end, we propose a plug-and-play *Semantics-Sequences-Structures* (SSS) framework which allows us to effectively combine such linguistic information. Through a series of experiments, we make some interesting observations. First, we observe that the popular adaptation of the GCN model for NLP tasks where structural information (GCNs) was added on top of sequential information (LSTMs) performs poorly on our task. This leads us to explore interesting ways of combining semantic and structural information to improve performance. Second, we observe that while BERT already outperforms other deep contextualized representations such as ELMo, it still benefits from the additional structural information explicitly added using GCNs. This is a bit surprising given the recent claims that BERT already captures structural information. Lastly, the proposed SSS framework gives an improvement of 7.95% BLEU score over the baseline.

[‡]The work was done by Nikita and Priyesh at Indian Institute of Technology Madras.

1 Introduction

Neural conversation systems that treat dialogue response generation as a sequence generation task (Vinyals and Le, 2015) often produce generic and incoherent responses (Shao et al., 2017). The primary reason for this is that, unlike humans, such systems do not have any access to background knowledge about the topic of conversation. For example, while chatting about movies, we use our background knowledge about the movie in the form of plot details, reviews, and comments that we might have read. To enrich such neural conversation systems, some recent works (Moghe et al., 2018; Dinan et al., 2019; Zhou et al., 2018) incorporate external knowledge in the form of documents which are relevant to the current conversation. For example, Moghe et al. (2018) released a dataset containing conversations about movies where every alternate utterance is extracted from a background document about the movie. This background document contains plot details, reviews, and Reddit comments about the movie. The focus thus shifts from sequence generation to identifying relevant snippets from the background document and modifying them suitably to form an appropriate response given the current conversational context.

Intuitively, any model for this task should exploit semantic, structural and sequential information from the conversation context and the background document. For illustration, consider the chat shown in Figure 1 from the Holl-E movie conversations dataset (Moghe et al., 2018). In this example, Speaker 1 nudges Speaker 2 to talk about how James’s wife was irritated because of his career. The right response to this conversation comes from the line beginning at “*His wife Mae ...*”. However, to generate this response, it is essential to understand that (i) *His* refers to James from the previous sentence; (ii) *quit boxing* is a contigu-

Source Doc: ... At this point James Brad-dock (Russel Crowe) was a light heavyweight boxer, who was forced to retired from the ring after breaking his hand in his last fight. **His wife Mae had prayed for years that he would quit boxing, before becoming permanently injured.** ...

Conversation:

Speaker 1(N): Yes very true, this is a real rags to riches story. Russell Crowe was excellent as usual.

Speaker 2(R): Russell Crowe owns the character of James Bradock, the unlikely hero who makes the most of his second chance. He's a good fighter turned hack.

Speaker 1(N): Totally! Oh by the way do you remember his wife ... how she wished he would stop

Speaker 2(P): His wife Mae had prayed for years that he would quit boxing, before becoming permanently injured.

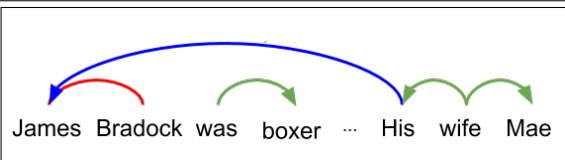


Figure 1: Sample conversation from the Holl-E Dataset. The text in bold in the first block is the background document which is used to generate the last utterance in this conversation. N, P, and R correspond to the type of background knowledge used: None, Plot, and Review as per the dataset definitions. For simplicity, we show only a few of the edges for the background knowledge at the bottom. The edge in blue corresponds to the co-reference edge, the edges in green are dependency edges and the edge in red is the entity edge.

ous phrase, and (iii) *quit* and *he would stop* mean the same. We need to exploit (i) **structural** information, such as, the co-reference edge between *His-James* (ii) the **sequential** information in *quit boxing* and (iii) the **semantic** similarity (or synonymy relation) between *quit* and *he would stop*.

To capture such multi-faceted information from the document and the conversation context we propose a new architecture that combines BERT with explicit sequence and structure information. We start with the deep contextualized word representations learnt by BERT which capture distributional semantics. We then enrich these representations with sequential information by allowing the words to interact with each other by passing

them through a bidirectional LSTM as is the standard practice in many NLP tasks. Lastly, we add explicit structural information in the form of dependency graphs, co-reference graphs, and entity co-occurrence graphs. To allow interactions between words related through such structures, we use GCNs which essentially aggregate information from the neighborhood of a word in the graph.

Of course, combining BERT with LSTMs in itself is not new and has been tried in the original work (Devlin et al., 2019) for the task of Named Entity Recognition. Similarly, Bastings et al. (2017) combine LSTMs with GCNs for the task of machine translation. To the best of our knowledge, this is the first work that combines BERT with explicit structural information. We investigate several interesting questions in the context of dialogue response generation. For example,

1. Are BERT-based models best suited for this task?
2. Should BERT representations be enriched with sequential information first or structural information?
3. Are dependency graph structures more important for this task or entity co-occurrence graphs?
4. Given the recent claims that BERT captures syntactic information, does it help to explicitly enrich it with syntactic information using GCNs?

To systematically investigate such questions we propose a simple plug-and-play *Semantics-Sequences-Structures* (SSS) framework which allows us to combine different semantic representations (GloVe (Pennington et al., 2014), BERT(Devlin et al., 2018), ELMo (Peters et al., 2018a)) with different structural priors (dependency graphs, co-reference graphs, etc.). It also allows us to use different ways of combining structural and sequential information, e.g., LSTM first followed by GCN or vice versa, or both in parallel. Using this framework we perform a series of experiments on the *Holl-E* dataset and make some interesting observations. First, we observe that the conventional adaptation of GCNs for NLP tasks, where contextualized embeddings obtained through LSTMs are fed as input to a GCN, exhibits poor performance. To overcome this, we propose some simple alternatives and show that they lead to

better performance. Second, we observe that while BERT performs better than GloVe and ELMo, it still benefits from explicit structural information captured by GCNs. We find this interesting because some recent works (Tenney et al., 2019; Jawahar et al., 2019; Hewitt and Manning, 2019) suggest that BERT captures syntactic information, but our results suggest that there is still more information to be captured by adding explicit structural priors. Third, we observe that certain graph structures are more useful for this task than others. Lastly, our best model which uses a specific combination of semantic, sequential, and structural information improves over the baseline by 7.95% on the BLEU score.

2 Related work

There is an active interest in using external knowledge to improve the informativeness of responses for goal-oriented as well as chit-chat conversations (Lowe et al., 2015; Ghazvininejad et al., 2018; Moghe et al., 2018; Dinan et al., 2019). Even the teams participating in the annual Alexa Prize competition (Ram et al., 2017) have benefited by using several knowledge resources. This external knowledge can be in the form of knowledge graphs or unstructured texts such as documents.

Many NLP systems including conversation systems use RNNs as their basic building block which typically captures n -gram or sequential information. Adding structural information through tree-based structures (Tai et al., 2015) or graph-based structures (Marcheggiani and Titov, 2017) on top of this has shown improved results on several tasks. For example, GCNs have been used to improve neural machine translation (Marcheggiani et al., 2018) by exploiting the semantic structure of the source sentence. Similarly, GCNs have been used with dependency graphs to incorporate structural information for semantic role labelling (Marcheggiani and Titov, 2017), neural machine translation (Bastings et al., 2017) and entity relation information in question answering (De Cao et al., 2019) and temporal information for neural dating of documents (Vashishth et al., 2018).

There have been advances in learning deep contextualized word representations (Peters et al., 2018b; Devlin et al., 2019) with a hope that such representations will implicitly learn structural and relational information with the interaction between words at multiple layers (Jawahar et al., 2019; Pe-

ters et al., 2018c). These recent developments have led to many interesting questions about the best way of exploiting rich information from sentences and documents. We try to answer some of these questions in the context of background aware dialogue response generation.

3 Background

In this section, we provide a background on how GCNs have been leveraged in NLP to incorporate different linguistic structures.

The Syntactic-GCN proposed in (Marcheggiani and Titov, 2017) is a GCN (Kipf and Welling, 2017) variant which can model multiple edge types and edge directions. It can also dynamically determine the importance of an edge. They only work with one graph structure at a time with the most popular structure being the dependency graph of a sentence. For convenience, we refer to Syntactic-GCNs as GCNs from here on.

Let G denote a graph defined on a text sequence (sentence, passage or document) with nodes as words and edges representing a directed relation between words. Let \mathcal{N} denote a dictionary of list of neighbors with $\mathcal{N}(v)$ referring to the neighbors of a specific node v , including itself (self-loop). Let $dir(u, v) \in \{in, out, self\}$ denote the direction of the edge, (u, v) . Let \mathcal{L} be the set of different edge types and let $L(u, v) \in \mathcal{L}$ denote the label of the edge, (u, v) . The $(k + 1)$ -hop representation of a node v is computed as

$$h_v^{(k+1)} = \sigma \left(\sum_{u \in \mathcal{N}(v)} g_{(u,v)}^{(k)} (W_{dir(u,v)}^{(k)} h_u^{(k)} + b_{L(u,v)}^{(k)}) \right) \quad (1)$$

where σ is the activation function, $g_{(u,v)} \in \mathbb{R}$ is the predicted importance of the edge (u, v) and $h_v \in \mathcal{R}^m$ is node, v 's embedding. $W_{dir(u,v)} \in \{W_{in}, W_{out}, W_{self}\}$ depending on the direction $dir(u, v)$ and W_{in}, W_{self} and $W_{out} \in \mathcal{R}^{m \times m}$. The importance of an edge $g_{(u,v)}$ is determined by an edge gating mechanism w.r.t. the node of interest, u as given below:

$$g_{(u,v)} = sigmoid(h_u \cdot W_{dir(u,v)} + b_{L(u,v)}) \quad (2)$$

In summary, a GCN computes new representation of a node u by aggregating information from its neighborhood $\mathcal{N}(v)$. When $k=0$, the aggregation happens only from immediate neighbors, i.e., 1 hop neighbors. As the value of k increases the aggregation implicitly happens from a larger neighborhood.

4 Proposed Model

Given a document D and a conversational context Q the task is to generate the response $\mathbf{y} = y_1, y_2, \dots, y_m$. This can be modeled as the problem of finding a \mathbf{y} that maximizes the probability $P(\mathbf{y}|D, Q)$ which can be further decomposed as

$$\mathbf{y} = \arg \max_{\mathbf{y}} \prod_{t=1}^m P(y_t | y_1, \dots, y_{t-1}, Q, D)$$

As has become a standard practice in most NLG tasks, we model the above probability using a neural network comprising of an encoder, a decoder, an attention mechanism, and a copy mechanism. The copy mechanism essentially helps to directly copy words from the document D instead of predicting them from the vocabulary. Our main contribution is in improving the document encoder where we use a plug-and-play framework to combine semantic, structural, and sequential information from different sources. This enriched document encoder could be coupled with any existing model. In this work, we couple it with the popular Get To The Point (GTTP) model (See et al., 2017) as used by the authors of the *Holl-E* dataset. In other words, we use the same attention mechanism, decoder, and copy mechanism as GTTP but augment it with an enriched document encoder. Below, we first describe the document encoder and then very briefly describe the other components of the model. We also refer the reader to the supplementary material for more details.

4.1 Encoder

Our encoder contains a semantics layer, a sequential layer and a structural layer to compute a representation for the document words which is a sequence of words w_1, w_2, \dots, w_m . We refer to this as a plug-and-play document encoder simply because it allows us to plug in different semantic representations, different graph structures, and different simple but effective mechanisms for combining structural and semantic information.

Semantics Layer: Similar to almost all NLP models, we capture semantic information using word embeddings. In particular, we utilize the ability of BERT to capture deep contextualized representations and later combine it with explicit structural information. This allows us to evaluate (i) whether BERT is better suited for this task as compared to other embeddings such as ELMo and GloVe and

(ii) whether BERT already captures syntactic information completely (as claimed by recent works) or can it benefit from additional syntactic information as described below.

Structure Layer: To capture structural information we propose multi-graph GCN, *M-GCN*, a simple extension of GCN to extract relevant multi-hop multi-relational dependencies from multiple structures/graphs efficiently. In particular, we generalize G to denote a labelled multi-graph, *i.e.*, a graph which can contain multiple (parallel) labelled edges between the same pair of nodes. Let \mathcal{R} denote the set of different graphs (structures) considered and let $G = \{\mathcal{N}_1, \mathcal{N}_2 \dots \mathcal{N}_{|\mathcal{R}|}\}$ be a set of dictionary of neighbors from the $|\mathcal{R}|$ graphs. We extend the Syntactic GCN defined in Eqn: 1 to multiple graphs by having $|\mathcal{R}|$ graph convolutions at each layer as given in Eqn: 3. Here, $g_conv(\mathcal{N})$ is the graph convolution defined in Eqn: 1 with σ as the identity function. Further, we remove the individual node (or word) i from the neighbourhood list $\mathcal{N}(i)$ and model the node information separately using the parameter W_{self} .

$$h_i^{(k+1)} = ReLU((h_i^{(k)} W_{self}^{(k)} + \sum_{\mathcal{N} \in G} g_conv(\mathcal{N})) \quad (3)$$

This formulation is advantageous over having $|\mathcal{R}|$ different GCNs as it can extract information from multi-hop pathways and can use information across different graphs with every GCN layer (hop). Note that h_i^0 is the embedding obtained for word v from the semantic layer. For ease of notation, we use the following functional form to represent the final representation computed by M-GCN after k -hops starting from the initial representation h_i^0 , given G .

$$h_i = M-GCN(h_i^0, G, k)$$

Sequence Layer: The purpose of this layer is to capture sequential information. Once again, following standard practice, we pass the word representations computed by the previous layer through a bidirectional LSTM to compute a sequence contextualized representation for each word. As described in the next subsection, depending upon the manner in which we combine these layers, the previous layer could either be the structure layer or the semantics layer.

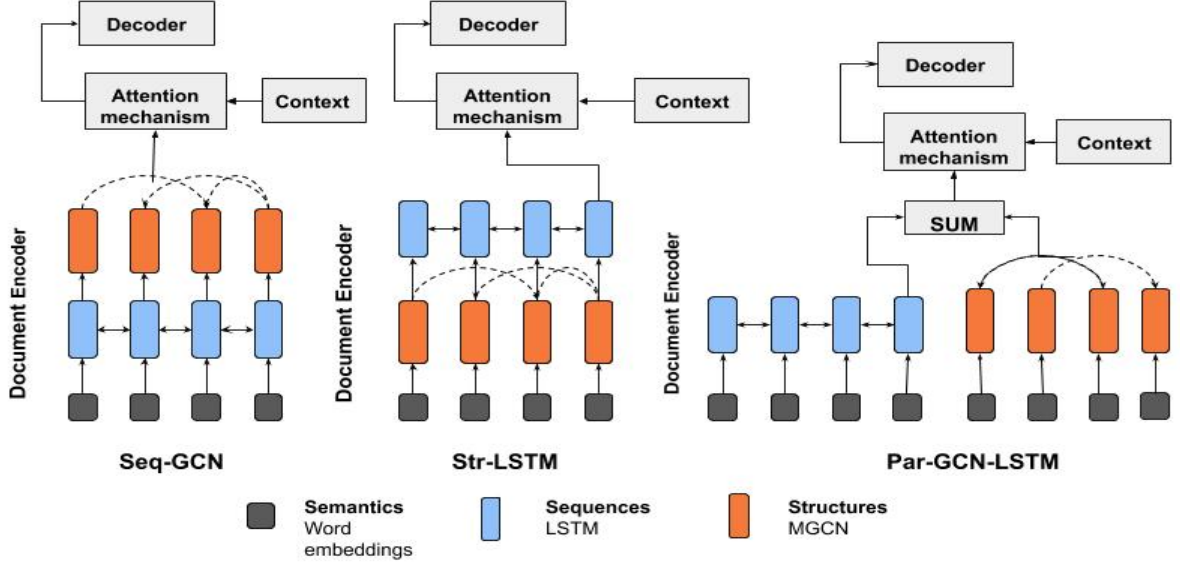


Figure 2: The SSS framework. The Word Embeddings include GloVe, ELMo, and BERT. *Seq-GCN* considers obtaining an LSTM representation first which this then passed through an M-GCN module. In *Str-LSTM*, we compute the M-GCN representation first which is then passed to an LSTM layer while *Par-GCN-LSTM* computes both LSTM and M-GCN representations independently which are then combined into a final representation.

4.2 Combining structural and sequential information

As mentioned earlier, for a given document D containing words $w_1, w_2, w_3, \dots, w_m$, we first obtain word representations $x_1, x_2, x_3, \dots, x_m$ using BERT (or ELMo or GloVe). At this point, we have three different choices for enriching the representations using structural and sequential information: (i) structure first followed by sequence (ii) sequence first followed by structure or (iii) structure and sequence in parallel. We depict these three choices pictorially in Figure 2 and describe them below with appropriate names for future reference. Please note that the choice of “Seq” denotes the sequential nature of LSTMs while “Str” denotes the structural nature of GCNs. Though we use a specific variant of GCN, described as M-GCN in the previous section, any other variant of GCN can be replaced in the “Str” layer.

4.2.1 Sequence contextualized GCN (*Seq-GCN*)

Seq-GCN is similar to the model proposed in (Bastings et al., 2017; Marcheggiani and Titov, 2017) where the word representations $x_1, x_2, x_3, \dots, x_m$ are first fed through a BiLSTM to obtain sequence contextualized representations as shown below.

$$h_i^{seq} = BiLSTM(h_{i-1}^{seq}, x_i)$$

These representations $h_1, h_2, h_3, \dots, h_m$ are

then fed to the M-GCN along with the graph G to compute a k -hop aggregated representation as shown below:

$$h_i^{str} = M-GCN(h_i^{seq}, G, k)$$

This final representation $h_i^{final} = h_i^{str}$ for the i -th word thus combines semantics, sequential and structural information in that order. This is a popular way of combining GCNs with LSTMs but our experiments suggest that this does not work well for our task. We thus explore two other variants as explained below.

4.2.2 Structure contextualized LSTM (*Str-LSTM*)

Here, we first feed the word representations $x_1, x_2, x_3, \dots, x_m$ to M-GCN to obtain structure aware representations as shown below.

$$h_i^{str} = M-GCN(x_i, G, k)$$

These structure aware representations are then passed through a BiLSTM to capture sequence information as shown below:

$$h_i^{seq} = BiLSTM(h_{i-1}^{seq}, h_i^{str})$$

This final representation $h_i^{final} = h_i^{seq}$ for the i -th word thus combines semantics, structural and sequential information in that order.

4.2.3 Parallel GCN-LSTM (*Par-GCN-LSTM*)

Here, both M-GCN and BiLSTMs are fed with word embeddings x_i as input to aggregate structural and sequential information independently as shown below:

$$\begin{aligned}h_i^{str} &= M\text{-GCN}(x_i, G, k) \\h_i^{seq} &= BiLSTM(h_{i-1}^{seq}, x_i)\end{aligned}$$

The final representation, h_i^{final} , for each word is computed as $h_i^{final} = h_i^{str} + h_i^{seq}$ and combines structural and sequential information in parallel as opposed to a serial combination in the previous two variants.

4.3 Decoder, Attention, and Copy Mechanism

Once the final representation for each word is computed, an attention weighted aggregation, c_t , of these representations is fed to the decoder at each time step t . The decoder itself is a LSTM which computes a new state vector s_t at every timestep t as

$$s_t = LSTM(s_{t-1}, c_t)$$

The decoder then uses this s_t to compute a distribution over the vocabulary where the probability of the i -th word in the vocabulary is given by $p_i = \text{softmax}(Vs_t + Wc_t + b)_i$. In addition, the decoder also has a copy mechanism wherein, at every timestep t , it could either choose the word with the highest probability p_i or copy that word from the input which was assigned the highest attention weight at timestep t . Such copying mechanism is useful in tasks such as ours where many words in the output are copied from the document D . We refer the reader to the GTTP paper for more details of the standard copy mechanism.

5 Experimental setup

In this section, we briefly describe the dataset and task setup followed by the pre-processing steps we carried to obtain different linguistic graph structures on this dataset. We then describe the different baseline models. Our code is available at: https://github.com/nikitacs16/horovod_gcn_pointer_generator

5.1 Dataset description

We evaluate our models using Holl-E, an English language movie conversation dataset (Moghe et al.,

2018) which contains $\sim 9k$ movie chats and $\sim 90k$ utterances. Every chat in this dataset is associated with a specific background knowledge resource from among the plot of the movie, the review of the movie, comments about the movie, and occasionally a fact table. Every even utterance in the chat is generated by copying and or modifying sentences from this unstructured background knowledge. The task here is to generate/retrieve a response using conversation history and appropriate background resources. Here, we focus only on the *oracle* setup where the correct resource from which the response was created is provided explicitly. We use the same train, test, and validation splits as provided by the authors of the paper.

5.2 Construction of linguistic graphs

We consider leveraging three different graph-based structures for this task. Specifically, we evaluate the popular syntactic word dependency graph (*Dep-G*), entity co-reference graph (*Coref-G*) and entity co-occurrence graph (*Ent-G*). Unlike the word dependency graph, the two entity-level graphs can capture dependencies that may span across sentences in a document. We use the dependency parser provided by SpaCy (<https://spacy.io/>) to obtain the dependency graph (*Dep-G*) for every sentence. For the construction of the co-reference graph (*Coref-G*), we use the NeuralCoref model (<https://github.com/huggingface/neuralcoref>) integrated with SpaCy. For the construction of the entity graph (*Ent-G*), we first perform named-entity recognition using SpaCy and connect all the entities that lie in a window of $k = 20$.

5.3 Baselines

We categorize our baseline methods as follows:

Without Background knowledge: We consider the simple Sequence-to-Sequence (S2S) (Vinyals and Le, 2015) architecture that conditions the response generation only on the *previous* utterance and completely ignores the other utterances as well as the background document. We also consider HRED (Serban et al., 2016), a hierarchical variant of the S2S architecture which conditions the response generation on the entire conversation history in addition to the last utterance. Of course, we do not expect these models to perform well as they completely ignore the background knowledge but we include them for the sake of completeness.

With Background Knowledge: To the S2S architecture we add an LSTM encoder to encode the

document. The output is now conditioned on this representation in addition to the previous utterance. We refer to this architecture as S2S-D. Next, we use GTTP (See et al., 2017) which is a variant of the S2S-D architecture with a copy-or-generate decoder; at every time-step, the decoder decides to copy from the background knowledge or generate from the fixed vocabulary. We also report the performance of the BiRNN + GCN architecture that uses the dependency graph only as discussed in (Marcheggiani and Titov, 2017). Finally, we note that in our task many words in the output need to be copied sequentially from the input background document which makes it very similar to the task of span prediction as used in Question Answering. We thus also evaluate BiDAF (Seo et al., 2017), a popular question-answering architecture, that extracts a span from the background knowledge as a response using complex attention mechanisms. For a fair comparison, we evaluate the spans retrieved by the model against the ground truth responses.

We use BLEU-4 and ROUGE (1/2/L) as the evaluation metrics as suggested in the dataset paper. Using automatic metrics is more reliable in this setting than the open domain conversational setting as the variability in responses is limited to the information in the background document. We provide implementation details in Appendix A.

6 Results and Discussion

In Table 1, we compare our architecture against the baselines as discussed above. SSS(BERT) is our proposed architecture in terms of the SSS framework. We report best results within SSS chosen across 108 configurations comprising of four different graph combinations, three different contextual and structural infusion methods, three M-GCN layers, and, three embeddings. The best model was chosen based on performance of the validation set. From Table 1, it is clear that our improvements in incorporating structural and sequential information with BERT in the SSS encoder framework significantly outperforms all other models.

6.1 Qualitative Evaluation

We conducted human evaluation for the SSS models from Table 1 against the generated responses of GTTP. We presented 100 randomly sampled outputs to three different annotators. The annotators were asked to pick from four options: A, B, both, and none. The annotators were told these were con-

Model	BLEU	ROUGE		
		1	2	L
S2S	4.63	26.91	9.34	21.58
HRED	5.23	24.55	7.61	18.87
S2S-D	11.71	26.36	13.36	21.96
GTTP	13.97	36.17	24.84	31.07
BiRNN+GCN	14.70	36.24	24.60	31.29
BiDAF	16.79	26.73	18.82	23.58
SSS(GloVe)	18.96	38.61	26.92	33.77
SSS(ELMo)	19.32	39.65	27.37	34.86
SSS(BERT)	22.78	40.09	27.83	35.20

Table 1: Results of automatic evaluation. The architectures within the SSS framework outperform the baseline methods with our proposed architecture SSS(BERT) performing the best.

versations between friends. Tallying the majority vote, we obtain win/loss/both/none for SSS(BERT) as 29/25/29/17, SSS(GloVe) as 24/17/47/12 and SSS(ELMo) as 22/23/41/14. This suggests qualitative improvement using the SSS framework. We also provide some generated examples in Appendix B1. We found that the SSS framework had less confusion in generating the opening responses than the GTTP baseline. These “conversation starters” have a unique template for every opening scenario and thus have different syntactic structures respectively. We hypothesize that the presence of dependency graphs over these respective sentences helps to alleviate the confusion as seen in Example 1. The second example illustrates why incorporating structural information is important for this task. We also observed that the SSS encoder framework does not improve on the aspects of human creativity such as diversity, initiating a context-switch, and common-sense reasoning as seen in Example 3.

6.2 Ablation studies on the SSS framework

We report the component-wise results for the SSS framework in Table 2. The *Sem* models condition the response generation directly on the word embeddings. As expected, we observe that ELMo and BERT perform much better than GloVe embeddings.

The *Sem+Seq* models condition the decoder on the representation obtained after passing the word embeddings through the LSTM layer. These models outperform their respective *Sem* models. The gain with ELMo is not significant because the underlying architecture already has two BiLSTM layers which are already being fine-tuned for the task.

Emb	Paradigm	BLEU	ROUGE		
			1	2	L
GloVe	Sem	4.4	29.72	11.72	22.99
	Sem+Seq	14.83	36.17	24.84	31.07
	SSS	18.96	38.61	26.92	33.77
ELMo	Sem	14.36	32.04	18.75	26.71
	Sem+Seq	14.61	35.54	24.58	30.71
	SSS	19.32	39.65	27.37	34.86
BERT	Sem	11.26	33.86	16.73	26.44
	Sem+Seq	18.49	37.85	25.32	32.58
	SSS	22.78	40.09	27.83	35.2

Table 2: Performance of components within the SSS framework. BERT based models outperform both ELMo and GloVe based architectures in the respective paradigms. Notably, adding all the three levels of information: semantic, sequential, and structural is useful.

Hence the addition of one more LSTM layer may not contribute to learning any new sequential word information. It is clear from Table 2 that the SSS models, that use structure information as well, obtain a significant boost in performance, validating the need for incorporating all three types of information in the architecture.

6.3 Combining structural and sequential information

The response generation task of our dataset is a span based generation task where phrases of text are expected to be copied or generated as they are. The sequential information is thus crucial to reproduce these long phrases from background knowledge. This is strongly reflected in Table 3 where *Str-LSTM* which has the LSTM layer on top of GCN layers performs the best across the hybrid architectures discussed in Figure 2. The *Str-LSTM* model can better capture sequential information with structurally and syntactically rich representations obtained through the initial GCN layer. The *Par-GCN-LSTM* model performs second best. However, in the parallel model, the LSTM cannot leverage the structural information directly and relies only on the word embeddings. *Seq-GCN* model performs the worst among all the three as the GCN layer at the top is likely to modify the sequence information from the LSTMs.

6.4 Understanding the effect of structural priors

While a combination of intra-sentence and inter-sentence graphs is helpful across all the models, the best performing model with BERT embeddings

relies only on the dependency graph. In the case of GloVe based experiments, the entity and co-reference relations were not independently useful with the *Str-LSTM* and *Par-GCN-LSTM* models, but when used together gave a significant performance boost, especially for *Str-LSTM*. However, most of the BERT based and ELMo based models achieved competitive performance with the individual entity and co-reference graphs. There is no clear trend across the models. Hence, probing these embedding models is essential to identify which structural information is captured implicitly by the embeddings and which structural information needs to be added explicitly. For the quantitative results, please refer to Appendix B2.

6.5 Structural information in deep contextualised representations

Earlier work has suggested that deep contextualized representations capture syntax and co-reference relations (Peters et al., 2018c; Jawahar et al., 2019; Tenney et al., 2019; Hewitt and Manning, 2019). We revisit Table 2 and consider the *Sem+Seq* models with ELMo and BERT embeddings as two architectures that *implicitly* capture structural information. We observe that the SSS model using the simpler GloVe embedding outperforms the ELMo *Sem+Seq* model and performs slightly better than the BERT *Sem+Seq* model.

Given that the SSS models outperform the corresponding *Sem+Seq* model, the extent to which the deep contextualized word representations learn the syntax and other linguistic properties implicitly are questionable. Also, this calls for better loss functions for learning deep contextualized representations that can incorporate structural information explicitly.

More importantly, all the configurations of SSS (GloVe) have a lesser memory footprint in comparison to both ELMo and BERT based models. Validation and training of GloVe models require one-half, sometimes even one-fourth of computing resources. Thus, the simple addition of structural information through the GCN layer to the established Sequence-to-Sequence framework that can perform comparably to stand-alone expensive models is an important step towards Green AI (Schwartz et al., 2019).

Emb	Seq-GCN				Str-LSTM				Par-GCN-LSTM			
	BLEU	ROUGE			BLEU	ROUGE			BLEU	ROUGE		
		1	2	L		1	2	L		1	2	L
GloVe	15.61	36.6	24.54	31.68	18.96	38.61	26.92	33.77	17.1	37.04	25.70	32.2
ELMo	18.44	37.92	26.62	33.05	19.32	39.65	27.37	34.86	16.35	37.28	25.67	32.12
BERT	20.43	40.04	26.94	34.85	22.78	40.09	27.83	35.20	21.32	39.9	27.60	34.87

Table 3: Performance of different hybrid architectures to combine structural information with sequence information. We observe that using structural information followed by sequential information, Str-LSTM, provides the best results.

7 Conclusion and Future Work

We demonstrated the usefulness of incorporating structural information for the task of background aware dialogue response generation. We infused the structural information explicitly in the standard semantic+sequential model and observed a performance boost. We studied different structural linguistic priors and different ways to combine sequential and structural information. We also observe that explicit incorporation of structural information helps the richer deep contextualized representation based architectures. The framework provided in this work is generic and can be applied to other background aware dialogue datasets and several tasks such as summarization and question answering. We believe that the analysis presented in this work would serve as a blueprint for analyzing future work on GCNs ensuring that the gains reported are robust and evaluated across different configurations.

Acknowledgements

We would like to thank the anonymous reviewers for their useful comments and suggestions. We would like to thank the Department of Computer Science and Engineering, and Robert Bosch Center for Data Sciences and Artificial Intelligence, IIT Madras (RBC-DSAI) for providing us with adequate resources. Lastly, we extend our gratitude to the volunteers who participated in the human evaluation experiments.

References

- Joost Bastings, Ivan Titov, Wilker Aziz, Diego Marcheggiani, and Khalil Sima'an. 2017. [Graph convolutional encoders for syntax-aware neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1957–1967, Copenhagen, Denmark. Association for Computational Linguistics.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2019. [Question answering by reasoning across documents with graph convolutional networks](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2306–2317, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. [Wizard of wikipedia: Knowledge-powered conversational agents](#). *International Conference on Learning Representations*.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. [A knowledge-grounded neural conversation model](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 5110–5117.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.

- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What does BERT learn about the structure of language?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Thomas N. Kipf and Max Welling. 2017. [Semi-supervised classification with graph convolutional networks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Ryan Lowe, Nissan Pow, Iulian Serban, Laurent Charlin, and Joelle Pineau. 2015. [Incorporating unstructured textual knowledge sources into neural dialogue systems](#). In *Neural Information Processing Systems Workshop on Machine Learning for Spoken Language Understanding*.
- Diego Marcheggiani, Joost Bastings, and Ivan Titov. 2018. [Exploiting semantics in neural machine translation with graph convolutional networks](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 486–492, New Orleans, Louisiana. Association for Computational Linguistics.
- Diego Marcheggiani and Ivan Titov. 2017. [Encoding sentences with graph convolutional networks for semantic role labeling](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1506–1515, Copenhagen, Denmark. Association for Computational Linguistics.
- Nikita Moghe, Siddhartha Arora, Suman Banerjee, and Mitesh M. Khapra. 2018. [Towards exploiting background knowledge for building conversation systems](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2322–2332, Brussels, Belgium. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018a. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018b. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Luke Zettlemoyer, and Wen-tau Yih. 2018c. [Dissecting contextual word embeddings: Architecture and representation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1499–1509, Brussels, Belgium. Association for Computational Linguistics.
- Matthew E. Peters, Sebastian Ruder, and Noah A. Smith. 2019. [To tune or not to tune? adapting pre-trained representations to diverse tasks](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepLanLP-2019)*, pages 7–14, Florence, Italy. Association for Computational Linguistics.
- Ashwin Ram, Rohit Prasad, Chandra Khatri, Anu Venkatesh, Raefer Gabriel, Qing Liu, Jeff Nunn, Behnam Hedayatnia, Ming Cheng, Ashish Nagar, Eric King, Kate Bland, Amanda Wartick, Yi Pan, Han Song, Sk Jayadevan, Gene Hwang, and Art Pettigrew. 2017. [Conversational AI: the science behind the alexa prize](#). *Alexa Prize Proceedings*.
- Roy Schwartz, Jesse Dodge, Noah A Smith, and Oren Etzioni. 2019. [Green AI](#).
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. [Bidirectional attention flow for machine comprehension](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016. [Building end-to-end dialogue systems using generative hierarchical neural network models](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 3776–3784.
- Yuanlong Shao, Stephan Gouws, Denny Britz, Anna Goldie, Brian Strope, and Ray Kurzweil. 2017. [Generating high-quality and informative conversation responses with sequence-to-sequence models](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2210–2219, Copenhagen, Denmark. Association for Computational Linguistics.

Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. [Improved semantic representations from tree-structured long short-term memory networks](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1556–1566, Beijing, China. Association for Computational Linguistics.

Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.

Shikhar Vashishth, Shib Sankar Dasgupta, Swayambhu Nath Ray, and Partha Talukdar. 2018. [Dating documents using graph convolution networks](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1605–1615, Melbourne, Australia. Association for Computational Linguistics.

Oriol Vinyals and Quoc V. Le. 2015. [A neural conversational model](#). In *ICML Deep Learning Workshop*.

Kangyan Zhou, Shrimai Prabhumoye, and Alan W Black. 2018. [A dataset for document grounded conversations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 708–713, Brussels, Belgium. Association for Computational Linguistics.

A Implementation Details

A.1 Base Model

The authors of [Moghe et al. \(2018\)](#) adapted the architecture of Get to the Point ([See et al., 2017](#)) for background aware dialogue response generation task. In the summarization task, the input is a *document* and the output is a *summary* whereas in our case the input is a $\{\text{resource/document, context}\}$ pair and the output is a *response*. Note that the context includes the previous two utterances (dialog history) and the current utterance. Since, in both the tasks, the output is a sequence (*summary* v/s *response*) we don’t need to change the decoder (*i.e.*, we can use the decoder from the original model as it is). However, we need to change the input fed to the decoder. We use an RNN to compute a representation of the conversation history. Specifically, we consider the previous k utterances as a single sequence of words and feed these to an RNN. Let M be the total length of the context (*i.e.*, all the k utterances taken together) then the RNN computes representations $h_1^d, h_2^d, \dots, h_M^d$ for all the words in

the context. The final representation of the context is then the attention weighted sum of these word representations:

$$\begin{aligned} f_i^t &= v^T \tanh(W_c h_i^d + V s_t + b_d) \\ m^t &= \text{softmax}(f^t) \\ d_t &= \sum_i m_i^t h_i^d \end{aligned} \quad (4)$$

Similar to the original model, we use an RNN to compute the representation of the document. Let N be the length of the document then the RNN computes representations $h_1^r, h_2^r, \dots, h_N^r$ for all the words in the resource (we use the superscript r to indicate resource). We then compute the query aware resource representation as follows.

$$\begin{aligned} e_i^t &= v^T \tanh(W_r h_i^r + U s_t + V d_t + b_r) \\ a^t &= \text{softmax}(e^t) \\ c_t &= \sum_i a_i^t h_i^r \end{aligned} \quad (5)$$

where c_t is the attended context representation. Thus, at every decoder time-step, the attention on the document words is also based on the currently attended context representation.

The decoder then uses r_t (document representation) and s_t (decoder’s internal state) to compute a probability distribution over the vocabulary P_{vocab} . In addition, the model also computes p_{gen} which indicates that there is a probability p_{gen} that the next word will be *generated* and a probability $(1 - p_{gen})$ that the next word will be *copied*. We use the following modified equation to compute p_{gen}

$$p_{gen} = \sigma(w_r^T r_t + w_s^T s_t + w_x^T x_t + b_g) \quad (6)$$

where x_t is the previous word predicted by the decoder and fed as input to the decoder at the current time step. Similarly, s_t is the current state of the decoder computed using this input x_t . The final probability of a word w is then computed using a combination of two distributions, *viz.*, (P_{vocab}) as described above and the attention weights assigned to the document words as shown below

$$P(w) = p_{gen} P_{vocab}(w) + (1 - p_{gen}) \sum_{i:w_i=w} a_i^t \quad (7)$$

where a_i^t are the attention weights assigned to every word in the document as computed in Equation 5. Thus, effectively, the model could learn to copy a word i if p_{gen} is low and a_i^t is high. This is

the baseline with respect to the LSTM architecture (Sem + Seq). For, GCN based encoders, the h_i^r is the final outcome after the desired GCN/LSTM configuration.

A.2 Hyperparameters

We selected the hyper-parameters using the validation set. We used Adam optimizer with a learning rate of 0.0004 and a batch size of 64. We used GloVe embeddings of size 100. For the RNN-based encoders and decoders, we used LSTMs with a hidden state of size 256. We used gradient clipping with a maximum gradient norm of 2. We used a hidden state of size 512 for *Seq-GCN* and 128 for the remaining GCN-based encoders. We ran all the experiments for 15 epochs and we used the checkpoint with the least validation loss for testing. For models using ELMo embeddings, a learning rate of 0.004 was most effective. For the BERT-based models, a learning rate of 0.0004 was suitable. Rest of the hyper-parameters and other setup details remain the same for experiments with BERT and ELMo. Our work follows a task specific architecture as described in the previous section. Following the definitions in (Peters et al., 2019), we use the “feature extraction” setup for both ELMo and BERT based models.

B Extended Results

B.1 Qualitative examples

We illustrate different scenarios from the dataset to identify the strengths and weaknesses of our models under the *SSS* framework in Table 4. We compare the outputs from the best performing model on the three different embeddings and use GTTP as our baseline. The best performing combination of sequential and structural information for all the three models in the *SSS* framework is *Str-LSTM*. The best performing *SSS*(GloVe) and *SSS*(ELMo) architectures use all the three graphs while *SSS*(BERT) uses only the dependency graph.

We find that the *SSS* framework improves over the baseline for the cases of opening statements (see Example 1). The baseline had confusion in picking opening statements and often mixed the responses for “Which is your favorite character?”, “Which is your favorite scene” and “What do you think about the movie?”. The responses to these questions have different syntactic structures - “My favorite character is XYZ”, “I liked the one in which XYZ”, and “ I think this movie

is XYZ” where XYZ was the respective crowd-sourced phrase. The presence of dependency graphs over the respective sentences may help to alleviate the confusion.

Now consider the example under *Hannibal* in Table 4. We find that the presence of a co-reference graph between “Anthony Hopkins” in the first sentence and “he” in the second sentence can help in continuing the conversation on the actor “Anthony Hopkins”. Moreover, connecting tokens in “Anthony Hopkins” to refer to “he” in the second sentence is possible because of the explicit entity-entity connection between the two tokens. However, this is applicable only to *SSS*(GloVe) and *SSS*(ELMo) as their best performing versions use these graphs along with the dependency graph while the best performing *SSS*(BERT) only uses dependency graph and may have learnt the inter-sentence relations implicitly.

There is a limited diversity of responses generated by the *SSS* framework as it often resorts to the patterns seen during training while it is not copying from the background knowledge. We also identify that *SSS* framework cannot handle the cases where Speaker2 initiates a context switch, *i.e.*; when Speaker2 introduces a topic that has not been discussed in the conversation so far. In the chat on *The Road Warrior* in Table 4, we find that *Mad Max: Fury Road* has been used to initiate a discussion that compares the themes of both the movies. All the models produce irrelevant responses.

B.2 Quantitative results

We explore the effect of using different graphs in Table 5.

Movie	Rocky V (Example 1)	Hannibal (Example 2)	The Road Warrior (Example 3)
Resource	Rocky V, terribly underrated. I liked the one in which Rocky My favorite character was Adrian because she was ...I think it was enjoyable, though slightly less than th.	Anthony Hopkins gave However, the material he was given to work almost seemed like an excuse to present us	Box Office: \$ 9,003,011 Similar Movies: Lethal Weapon, Mad Max: Fury Road ...
Chat	Speaker 1: Which is your favourite character in this?	Speaker 1: What is your opinion about the movie? Speaker 2: I think it was a fair sequel. Speaker 1: Anthony Hopkins gave an impeccable performance.	Speaker 1: Which is your favourite character in this? Speaker 2: My favorite character was Dog Speaker 1: I liked the feral kid even after he's found in the car and Max evicts him without ceremony he doesn't give up.
Ground Truth	my favorite character was adrian because she was strong and did her best to keep her family together through tough times .	however , the material he was given to work with was not as good as silence of the lambs .	Have you seen the new Mad Max: Fury Road?
Baseline	i liked the one in which rocky and gunn fight in the street .	in fairness , perhaps there was no way it could be .	lethal weapon .
SSS(GloVe)	my favorite character was adrian because she was strong and did her best to keep her together through tough times .	however , the material he was given to work with was not as good as silence of the lambs .	it made \$ 9,003,011
SSS(ELMo)	my favorite character was adrian because she was strong and did her best to keep her family together through tough times .	the material he was given to work with was not as good as silence of the lambs .	[UNK] only one man can make the difference in the future .
SSS(BERT)	my favorite character was adrian because she was strong and did her best to keep her family together through tough times .	the material he was given to work with was not as good as silence of the lambs .	yes .[UNK] only one man can make the difference in the future .

Table 4: Sample outputs from the SSS framework compared with baseline and ground truth responses.

Graph	GloVe				ELMo				BERT			
	BLEU	ROUGE			BLEU	ROUGE			BLEU	ROUGE		
		1	2	L		1	2	L		1	2	L
Dep	16.79	37.77	25.89	32.88	17.00	37.56	26.14	32.77	22.78	40.09	27.83	35.2
Dep+Ent	14.44	35.14	24.61	30.43	18.34	39.55	28.00	34.76	19.33	39.37	27.52	34.33
Dep+Coref	16.58	37.60	25.72	32.63	18.56	40.08	28.42	35.06	20.99	40.10	28.66	35.11
Dep+Ent +Coref	18.96	38.61	26.92	33.77	19.32	39.65	27.37	34.86	20.37	39.11	27.2	34.19

Table 5: Comparing performance of different structural priors across different semantic information on the *Str-LSTM* architecture.