

Discourse-Aware Semantic Self-Attention for Narrative Reading Comprehension

Todor Mihaylov and Anette Frank

Research Training Group AIPHES

Department of Computational Linguistics, Heidelberg University

Heidelberg, Germany

{mihaylov, frank}@cl.uni-heidelberg.de

Abstract

In this work, we propose to use linguistic annotations as a basis for a *Discourse-Aware Semantic Self-Attention* encoder that we employ for reading comprehension on narrative texts. We extract relations between discourse units, events and their arguments as well as coreferring mentions, using available annotation tools. Our empirical evaluation shows that the investigated structures improve the overall performance (up to +3.4 *Rouge-L*), especially intra-sentential and cross-sentential discourse relations, sentence-internal semantic role relations, and long-distance coreference relations. We show that dedicating self-attention heads to intra-sentential relations and relations connecting neighboring sentences is beneficial for finding answers to questions in longer contexts. Our findings encourage the use of discourse-semantic annotations to enhance the generalization capacity of self-attention models for reading comprehension.

1 Introduction

Transformer-based self-attention models (Vaswani et al., 2017) have been shown to work well on many natural language tasks that require large-scale training data, such as Machine Translation (Vaswani et al., 2017; Dai et al., 2019), Language Modeling (Radford et al., 2018a; Devlin et al., 2019; Dai et al., 2019; Radford et al., 2019) or Reading Comprehension (Yu et al., 2018), and can even be trained to perform surprisingly well in several multi-modal tasks (Kaiser et al., 2017b).

Recent work (Strubell et al., 2018) has shown that for downstream semantic tasks with much smaller datasets, such as Semantic Role Labeling (SRL) (Palmer et al., 2005), self-attention models greatly benefit from the use of linguistic information such as dependency parsing annotations. Motivated by this work, we examine to what extent

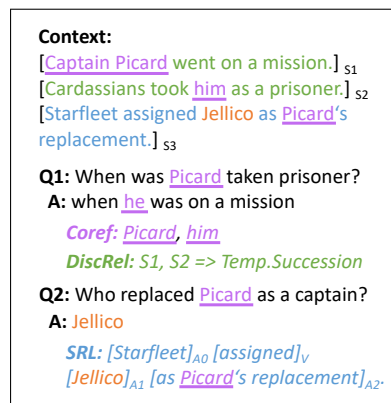


Figure 1: Motivational example: context and questions with required discourse and semantic annotations.

we can use discourse and semantic information to extend self-attention-based neural models for a higher-level task such as Reading Comprehension.

Reading Comprehension is a task that requires a model to answer natural language questions, given a text as context: a paragraph or even full documents. Many datasets have been proposed for the task, starting with a small multi-choice dataset (Richardson et al., 2013), large-scale automatically created cloze-style datasets (Hermann et al., 2015; Hill et al., 2016) and big manually annotated datasets such as Onishi et al. (2016); Rajpurkar et al. (2016); Joshi et al. (2017); Kocisky et al. (2018). Previous research has shown that some datasets are not challenging enough, as simple heuristics work well with them (Chen et al., 2016; Weissenborn et al., 2017b; Chen et al., 2017). In this work we focus on the recent NarrativeQA (Kocisky et al., 2018) dataset that was designed not to be easy to answer and that requires a model to read narrative stories and answer questions about them.

In terms of model architecture, previous work in reading comprehension and question answer-

DR (NE)	S-ARG1	S-ARG1	S-ARG1	S-ARG1	S-ARG1	S-ARG1	S-ARG1	.	S-ARG2	S-ARG2	S-ARG2	S-ARG2	S-ARG2	S-ARG2	.
SRL	B-ARG0	I-ARG0	B-V	B-ARG1	B-ARG2	I-ARG2	I-ARG2	.	B-ARG0	B-V	B-ARG1	B-ARG2	I-ARG2	I-ARG2	.
Coref	O	C5	O	C6	O	O	C6	.	C7	O	C8	O	C6	O	.
Tokens	The Cardassians took him as a prisoner . Starleet assigned Jellico as Picard's replacement .														

Figure 2: Example on different discourse-semantic annotations: DiscRel (Discourse Relations) (NE - Non-Explicit), SRL (Semantic Role Labeling), Coref (Co-reference resolution). The distinct horizontal lines show the interaction between the tokens: Coref - full context, SRL - single sentence, Non-Explicit DR - two neighbouring sentences.

ing has focused on integrating external knowledge (linguistic and/or knowledge-based) into recurrent neural network models using Graph Neural Networks (Song et al., 2018), Graph Convolutional Networks (Sun et al., 2018; De Cao et al., 2019), attention (Das et al., 2017; Mihaylov and Frank, 2018; Bauer et al., 2018) or pointers to coreferent mentions (Dhingra et al., 2017).

In contrast, in this work we examine the impact of *discourse-semantic annotations* (Figure 1) in a *self-attention architecture*. We build on the QANet (Yu et al., 2018) model by modifying the encoder of its self-attention modeling layer. In particular, we *specialize self-attention heads to focus on specific discourse-semantic annotations*, such as, e.g., an ARG1 relation in SRL, a CAUSATION relation holding between clauses in shallow discourse parsing, or coreference relations holding between entity mentions.

Our contributions are the following:

- To our knowledge we are the first to explicitly introduce discourse information into a neural model for reading comprehension.
- We design a *Discourse-Aware Semantic Self-Attention* mechanism, an extension to the standard self-attention models – without significant increase of computation complexity.
- We analyze the impact of different discourse and semantic annotations for narrative reading comprehension and report improvements of up to 3.4 *Rouge-L* over the base model.
- We perform empirical fine-grained evaluation of the discourse-semantic annotations on specific question types and context size.

Code and data will be available at <https://github.com/Heidelberg-NLP/discourse-aware-semantic-self-attention>.

2 Discourse-aware Semantic Annotations

Understanding narrative stories requires the ability to identify events and their participants and to identify how these events are related in discourse (e.g., by *causation*, *contrast*, or *temporal sequence*) (Mani, 2012). Our aim is to extract structured knowledge about these phenomena from long texts and to integrate this information in a neural self-attention model, in order to examine to what extent such knowledge can enhance the efficiency of a strong reading comprehension model applied to NarrativeQA.

Specifically, we enhance self-attention with knowledge about entity coreference (Coref), their participation in events (SRL) and the relation between events in narrative discourse (Shallow Discourse Parsing (Xue et al., 2016), DR).

All these linguistic information types are *relational* in nature. For integrating relational knowledge into the self-attention mechanism, we follow a two-step approach: i) we extract such relations from a multi-sentence paragraph and *project them down to the token level*, specifically to the tokens of the text fragments that they involve; ii) we design a neural self-attention model that *uses the interaction information between these tokens in a multi-head self-attention module*.

To be able to map the extracted linguistic knowledge to paragraph tokens, we need annotations that are easy to map to token level (see Figure 2). This can be achieved with tools for annotation of span-based Semantic Role Labeling, Coreference Resolution, and Shallow Discourse Parsing.

Events and Their Participants Relations between characters in a story are expressed in text through their participation in states or actions in which they fill a particular event argument with a specific semantic role (see Figure 2). For annotation of events and their participants we use the state-of-the-art SRL system of He et al. (2017) as

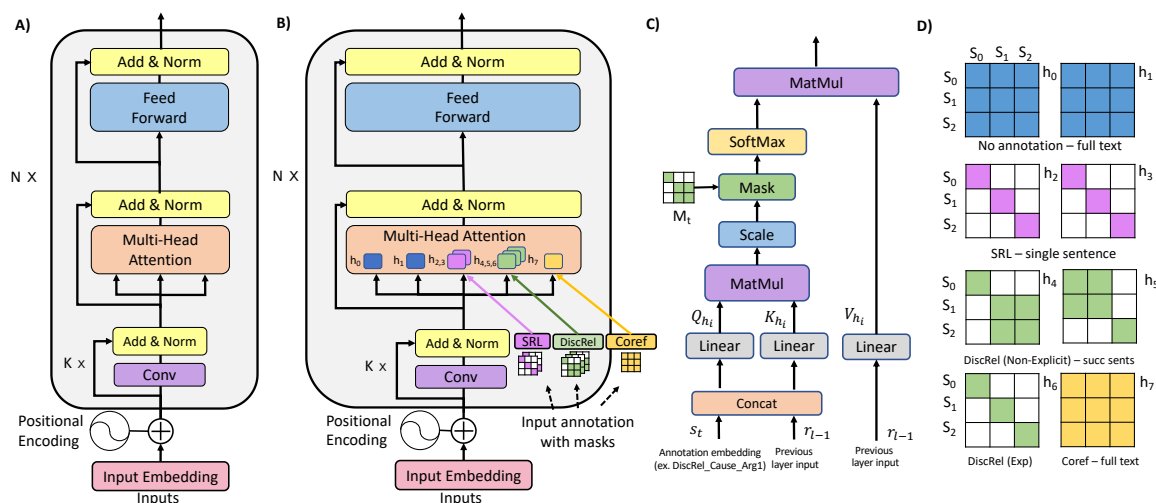


Figure 3: A) Base Multi-Head Self-Attention Encoder Block, B) Discourse-Aware Semantic Self-Attention (DASSA) Encoder Block, C) Single Attention Head with discourse/semantic Information, D) Example of attention scope masks for different attention heads and different information.

implemented in AllenNLP (Gardner et al., 2018). The system splits paragraphs into sentences and tokens, performs POS (part of speech tagging) and for each verb token V it predicts semantic tags such as ARG0, ARG1 (Argument Role 0, 1 of verb V), etc. When several argument-taking predicates are realized in a sentence, we obtain more than single semantic argument structure, and each token in the sentence can be involved in the argument structure of more than one verb. We refer to these annotations as different semantic views (Khashabi et al., 2018a), e.g., ‘semantic view for verb 1’. Different self-attention heads will be able to attend to individual semantic views.

Coreference Resolution Narrative texts abound of entity mentions that refer to the same entity in the discourse. We hypothesize that by directing the self-attention to this specific coreference information, we can encourage the model to focus on tokens that refer to the same entity mention. Although token-based self-attention models are able to attend over wide-ranged context spans, we hypothesize that it will be beneficial to allow the model to focus directly on the parts of the text that refer to the same entity. For coreference annotation we use the *medium* size model from the neuralcoref spaCy extension available at <https://github.com/huggingface/neuralcoref>. For each token we give as information the label of the corresponding coreference cluster (see Figure 2) that it belongs to. Therefore, tokens from the same coreference cluster get the same label as input.

Discourse Relations In narrative texts, events are connected by discourse relations such as *causation*, *temporal succession*, etc. (Mani, 2012). In this work we adopt the 15 fine-grained discourse relation sense types from the annotation scheme of the Penn Discourse Tree Bank (PDTB) (Prasad et al., 2008). For producing discourse relation annotations we use the discourse relation sense disambiguation system from Mihaylov and Frank (2016) which is trained on the data provided by the CoNLL Shared Task on Shallow Discourse Parsing (Xue et al., 2016). In this annotation scheme discourse relations are divided into two main types: *Explicit* and *Non-Explicit*. *Explicit* relations are usually connected with an explicit *discourse connective*, such as *because*, *but*, *if*. *Non-Explicit*¹ relations are not explicitly marked with a discourse connective and the arguments are usually contained in two consecutive sentences (see Figure 2). To extract explicit discourse relations we take into account only arguments that are in the same sentence. We consider as separate arguments (ARG1 and ARG2) text sequences that are on the left and right of an explicit discourse connective (CONN): ex. ‘[Jeff went home]_{ARG1.CCR} [because]_{CONN} [he was hungry.]_{ARG2.CCR}, where CCR is Contingency.Cause.Reason’. To provide *Non-Explicit* discourse relation sense annotations, we annotate every consecutive pair of sentences with a predicted discourse relation sense type.

¹*Non-Explicit* relations include *Implicit*, *AltLex* and *EnRel* relation from PDTB. See Xue et al. (2016) for details.

3 A Discourse-Aware Semantic Self-Attention Neural Model

3.1 QANet

As a base reading comprehension model we use QANet (Yu et al., 2018). QANet is a standard token-based self-attention model with the following components, which are common across many recent models: 1. *Input Embedding Layer* uses pre-trained word embeddings and convolutional character embeddings; 2. *Encoder Layer* consists of stacked *Encoder Blocks* (see Figure 3, A) based on *Multi-Head Self-Attention* (Vaswani et al., 2017) and *depth-wise separable convolution* (Chollet, 2016; Kaiser et al., 2017a); 3. *Context-to-Query Attention Layer* is a standard layer, that builds a token-wise attention-weighted question-aware context representation; 4. *Modeling Layer* has the same structure as 2. above but uses as input the output of layer 3.; 5. *Output layer* is used for prediction of *start* and *end* answer pointers. For detailed information about these layers, please refer to Yu et al. (2018). In this work we replace the standard *Multi-Head Self-Attention* with *Discourse-Aware Semantic Self-Attention*, using several different semantic and discourse annotation types. We describe this below and explain the differences to the standard *Encoder Block*.

3.2 Discourse-Aware Semantic Self-Attention

In Figure 3 we show the difference between the Base Multi-Head Self-Attention Encoder Block A) and the Discourse-Aware Semantic Self-Attention Encoder Block B). Both consist of positional-encoding + convolutional-layer $\times K$ + multi-head-self-attention + feed-forward layer. The difference is that B is provided additional inputs that are used by multi-head self-attention. The multi-head self-attention is a concatenation of outputs from multiple single self-attention heads h_i followed by a linear layer. A single head of the extended multi-head self-attention is shown in Figure 3C and is formally defined as

$$a_{h_i} = \text{mask_softmax} \left(\frac{Q_{h_i} K_{h_i}^T}{\sqrt{d_h}}, M_t \right) V_{h_i} \quad (1)$$

$$Q_{h_i} = W_{h_i}^Q [r_{l-1}; s_t] \in \mathbb{R}^{n \times d_h} \quad (2)$$

$$K_{h_i} = W_{h_i}^K [r_{l-1}; s_t] \in \mathbb{R}^{n \times d_h} \quad (3)$$

$$V_{h_i} = W_{h_i}^V r_{l-1} \in \mathbb{R}^{n \times d_h} \quad (4)$$

, where Q_{h_i} , K_{h_i} , V_{h_i} are components of the query-key-value attention and $\sqrt{d_h}$ is used for

weight scaling as originally proposed in Vaswani et al. (2017). $W_{h_i}^Q$, $W_{h_i}^K$, $W_{h_i}^V$ are weights, specific for head h_i , $i \in 1..H$ ², r_{l-1} is the input from the previous encoder block, s_t is an embedding vector for the linguistic annotation type t ('SRL_Arg1', 'DiscRel_Cause.Reason_Arg2', etc.), a_{h_i} is the output of head h_i . M_t is a sentence-wise attention mask as shown in Figure 3D. s_t and M_t are the main difference compared to the standard self-attention (Figure 3C).

In principle, representing edges of a graph (e.g., the V-ARG1 role from SRL) requires memory of $n^2 d_h H$, where n is the length of the context, which would be a bottleneck for computation on a GPU with limited memory (8-16GB). Instead, we adopt a strategy where the *relation* is represented as a *source and target node* and an *attention scope* (one sentence for SRL; two sentences for DR (Non-Exp); full context for Coref). The latter is controlled using the attention mask. The combination of flat token labels and mask reduces the maximum memory required for representing the information in the knowledge-enhanced head to $2nd_h H$. The attention masks, which we use for reducing the attention scope of the different semantic and discourse annotations, are shown in Figure 3D. These masks ensure that the corresponding attention heads will only attend to tokens from the corresponding scope (SRL: single sentence; DR (NonE): two sentences, etc.). The attention masks are symmetric to the matrix diagonal. Therefore, they can easily be computed 'on-the-fly' given only the sentence boundaries (corresponding to the horizontal lines in Figure 2).

To reduce the model memory further and still benefit from the full-context self-attention, we use the *Discourse-Aware Semantic Self-Attention* encoder (Figure 3B) only for blocks [1,3,5] of the *Modeling Layer* that consists of 7 stacked encoder blocks (indexed 0 to 6). Blocks [0,2,6] are set as the base encoders that look at the entire context (Figure 3A).

4 Data and Task Description

NarrativeQA We perform experiments with the NarrativeQA (Kocisky et al., 2018) reading comprehension dataset. This dataset requires understanding of narrative stories (English) in order to provide answers for a given question. It offers two

²Number of heads $H=8$ as in original QANet specification if not specified otherwise.

sub-tasks: (i) answering questions about a long narrative summary (up to 1150 tokens) of a book or movie, or (ii) answering questions about entire books or movie scripts of lengths up to 110k tokens. We are focusing on the summary setting (i) and refer to the summary as *document* or *context*. The dataset contains 1572 documents in total, divided into *Train* (1102 docs, 32.7k questions), *Dev* (115 documents, 3.5k questions) and *Test* (355 documents, 10.5k questions) sets.

Generative QA as Span Prediction An interesting aspect of the NarrativeQA dataset is that in contrast to most other RC datasets, the two answers provided for each question are written by human annotators. Therefore, answers typically differ in form from the context passages that license them. To map the human-generated answers to answer candidate spans from the context, we use *Rouge-L* (Lin, 2004) to calculate a similarity score between token n-grams from the provided answer and token n-grams from candidate answers selected from the context (we select candidate spans of the same length as the given answer). If two answer candidates have the same *Rouge-L* score, we calculate the score between the candidates' surrounding tokens (window size: 15 tokens to the left and right) and the question tokens, and choose the candidate with the higher score. We retrieve the best candidate answer span for each answer and use the candidate with the higher *Rouge-L* score as supervision for training. We refer to this method for answer retrieval as *Oracle (Ours)*.

5 Related Work

Reading Comprehension with Knowledge Recent work has proposed different approaches for integrating external knowledge into neural models for the high-level downstream tasks reading comprehension (RC) and question answering (QA). One line of work leverages external knowledge from knowledge bases for RC (Xu et al., 2016; Weissenborn et al., 2017a; Ostermann et al., 2018; Mihaylov and Frank, 2018; Bauer et al., 2018; Wang et al., 2018b) and QA (Das et al., 2017; Sun et al., 2018; Tandon et al., 2018). These approaches make use of implicit (Weissenborn et al., 2017a) or explicit (Mihaylov and Frank, 2018; Sun et al., 2018; Bauer et al., 2018) attention-based knowledge aggregation or leverage features from knowledge base relations (Wang et al., 2018b).

Another line of work builds on linguistic knowledge from downstream tasks, such as coreference resolution (Dhingra et al., 2017) or notions of co-occurring candidate mentions (De Cao et al., 2019) and OpenIE triples (Khot et al., 2017) into RNN-based encoders. Recently, several pre-trained language models (Peters et al., 2018; Radford et al., 2018b; Devlin et al., 2019) have been shown to incrementally boost the performance of well-performing models for several short paragraph reading comprehension tasks (Peters et al., 2018; Devlin et al., 2019) and question answering (Sun et al., 2019), as well as many tasks from the GLUE benchmark (Wang et al., 2018a). Approaches based on BERT (Devlin et al., 2019) usually perform best when the weights are fine-tuned for the specific training task. Earlier, many papers that do not use self-attention models or even neural methods have also tried to use semantic parse labels (Yih et al., 2016), or annotations from upstream tasks (Khashabi et al., 2018b).

Self-Attention Models in NLP Vanilla self-attention models (Vaswani et al., 2017) use positional encoding, sometimes combined with local convolutions (Yu et al., 2018) to model the token order in text. Although they are scalable due to their recurrence-free nature, most self-attention models do not well work when trained with fixed-length context, due to the fact that they often learn global token positions observed during training, rather than relative. To address this issue, Shaw et al. (2018) proposes relative position encoding to model the distance between tokens in the context. Dai et al. (2019) address the problem of moving beyond fixed-length context by adding recurrence to the self-attention model. Dai et al. (2019) argue that the fixed-length segments used for language modeling hurt the performance due to the fact that they do not respect sentence or any other semantic boundaries. In this work we also support the claim that the lack of semantic, and also discourse *boundaries* is an issue, and therefore we aim to introduce structured linguistic information into the self-attention model. We hypothesize that the lack of local discourse context is a problem for answering narrative questions, where the answer is contained inside the same sentence, or neighbouring sentences and therefore, by offering discourse-level semantic structure to the attention heads, offer ways to restrict, or focus the model to wider or narrower structures, depending on what is needed

Model	B-1	R-L
(Kocisky et al., 2018)		
Human	44.43	57.02
Oracle (original)	54.14	59.92
Seq2seq (no context) †	15.89	13.15
ASR †	23.30	22.26
BiDAF	33.72	36.30
Previous work		
BiAtt + MRU (Tay et al., 2018a)	36.55	41.44
DecaProp (Tay et al., 2018b)	44.35	44.69
MHPGM + NOIC (Bauer et al., 2018) †	43.63	44.16
RMR (Hu et al., 2018b)	48.40	51.50
RMR (Ens) (Hu et al., 2018b)	50.10	53.90
RMR + A2D (Hu et al., 2018b)	50.40	53.30
This work		
Oracle (ours)	70.71	70.82
BiDAF	47.19	49.63
QANet	46.37	48.66
+ DR (Exp)	50.12	52.14
+ DR (Exp) EMA	51.16	53.26

Table 1: Results on the NarrativeQA *Test* set. Models with † are generative, while the rest use span prediction.

for specific questions.

Self-attention architectures can be seen as graph architectures (imagine the token (node) interactions as adjacency matrix) and are applied to graph problems (Veličković et al., 2018; Li et al., 2019). Therefore, in very recent work Koncel-Kedziorski et al. (2019) have used a self-attention encoder as a graph encoder for text generation, in a dual encoder model. A dual-encoder model similar to Koncel-Kedziorski et al. (2019) is suitable for a setting where the input is knowledge from a graph knowledge-base. For a text-based setting like ours, where word order is important and the tokens are part of semantic arguments, an approach that tries to encode linguistic information in the same architecture (Strubell et al., 2018) is more appropriate. Therefore our method is most related to LISA (Strubell et al., 2018), which uses joint multi-task learning of POS and Dependency Parsing to inject syntactic information for Semantic Role Labeling. In contrast, we do not use multi-task learning, but directly encode semantic information extracted by pre-processing with existing tools.

NarrativeQA The summary setting of the NarrativeQA dataset (Kocisky et al., 2018) has in the past been addressed with attention mechanisms by the following models: *BiAtt + MRU* (Tay et al., 2018a) is similar to *BiDAF* (Seo et al., 2017). It is bi-attentive (attends from context-to-query and vice versa) but enhanced with a MRU (Multi-Range Reasoning Units). MRU is a compositional encoder that splits the context tokens

Config	DR-E	DR-NE	SRL	Coref	No
QANet (baseline)	-	-	-	-	8
DR (All)	2	2	-	-	4
DR (Exp)	2	-	-	-	6
DR (NonE)	-	2	-	-	6
Coref	-	-	-	3	5
SRL	-	-	3	-	5
SRL+ DR (Exp)	2	-	3	-	3
SRL + DR (NonE)	-	2	3	-	3
SRL + DR (All)	2	2	3	-	1
SRL + DR (Exp) + Coref	2	-	3	1	2
SRL + DR (All) + Coref	2	2	3	1	4

Table 2: The number of attention heads by discourse-semantic type. ‘No’ means that no linguistic annotation types are provided (attends to all tokens).

into ranges (n-grams) of different sizes and combines them in summed n-gram representations and fully-connected layers. *DecaProp* (Tay et al., 2018b) is a neural architecture for reading comprehension, that densely connects all pairwise layers, modeling relationships between passage and query across all hierarchical levels. Bauer et al. (2018) observed that some of the questions require external commonsense knowledge and developed *MHPGM-NOIC* - a *seq2seq* generative model with a copy mechanism that also uses commonsense knowledge and ELMo (Peters et al., 2018) contextual representations. Hu et al. (2018b) used an implementation of Reinforced Mnemonic Reader (*RMR*) (Hu et al., 2018a). They also proposed *RMR + A2D*, a novel teacher-student attention distillation method to train a model to mirror the behavior of the ensemble model *RMR (Ens)*.

6 Experiments and Results

In this section we describe the experiments and results of our proposed model in different configurations. We compare the results of different models using overall results (Table 1) on the dataset, but also the performance for different question types (Figure 4) and context sizes (Figure 5).

6.1 Overall Results

Table 1 compares our baselines and proposed model to prior work. We report results for *Bleu-1*, and *Rouge-L* scores. The first section lists results on the NarrativeQA dataset as reported in Kocisky et al. (2018). *Oracle (original)* uses the gold answers as queries to match a token sequence (with the answer length) in the context that

	Absolute		Improvement over QANet baseline														
all	70.82	48.66	0.97	1.21	1.34	3.43	0.98	0.56	1.34	1.61	1.53	1.84	1.40	1.31	1.13	1.21	
how	53.48	33.00	2.26	0.71	1.00	1.32	1.86	1.11	-0.17	0.53	1.00	0.85	0.44	1.34	0.08	1.96	
how far*	65.76	48.62	2.93	8.33	-26.40	10.45	10.45	2.86	10.45	16.70	2.93	10.45	10.45	-22.14	13.67	10.45	
how long*	69.92	62.91	2.31	3.68	6.25	3.33	-1.34	-2.25	-0.37	-0.61	-2.32	2.50	4.04	-1.22	4.52	2.01	
how many	54.57	46.75	5.36	12.36	2.45	2.54	5.50	2.50	8.63	6.99	3.26	5.42	3.00	5.58	10.15	7.05	
how much*	67.59	61.87	5.55	4.48	9.44	12.23	1.91	-4.11	0.09	5.75	3.91	11.52	4.73	7.28	9.43	4.22	
how old*	28.62	23.52	-4.13	4.22	-0.03	-9.91	1.05	-13.88	3.48	-0.65	5.22	1.87	1.14	-5.66	3.69	3.92	
other*	50.11	18.59	-2.02	7.73	-5.60	-1.53	0.88	6.14	3.44	2.17	7.77	6.74	7.91	2.62	-0.48	3.43	
what	69.73	47.88	0.69	0.77	0.40	3.11	0.28	0.71	0.62	0.68	1.24	1.36	1.61	0.97	0.54	0.55	
when	64.11	39.66	5.35	4.25	2.05	5.71	4.04	4.88	5.03	5.81	4.80	6.98	2.57	3.25	4.21	6.55	
where	80.66	62.75	-1.38	-0.22	-1.08	1.78	-0.01	0.10	0.61	-0.38	0.44	1.15	0.21	0.77	1.19	-0.40	
which	85.64	55.99	-2.78	-1.59	0.94	3.44	-2.81	-0.24	-1.20	-0.95	-1.57	3.14	0.91	-1.06	-0.72	1.28	
who	82.97	55.97	0.93	1.88	3.49	5.40	1.87	0.17	2.80	3.77	2.30	2.35	1.90	2.25	1.74	1.88	
why	49.59	33.51	2.19	0.69	1.31	1.86	1.35	0.06	0.96	1.02	1.54	0.96	-0.15	0.17	0.60	0.31	
	Oracle	QANet	BiDAF	Coref	DR (All)	DR (Exp)	DR (NonE)	DR (Exp NoSense)	DR (NonE NoSense)	Sent span 3	SRL	SRL + DR (Exp)	SRL + DR (NonE)	SRL + DR (All)	SRL + DR (Exp) + Coref	SRL + DR (All) + Coref	

Figure 4: Rouge-L performance per Question Type on the NarrativeQA *Test* set. The first two columns represent *Absolute* values. The rest are improvements over the QANet baseline model (i) by BiDAF and (ii) configurations of QANet with linguistic information. Question types with * have less than 100 instances in the *Test* set.

has the highest *Rouge-L*. In contrast, using *Oracle (Ours)*, described in Section 4, we report a +11 *Rouge-L* score improvement (Table 1: **This work**). The Oracle performance in this setting is important since the produced annotations are used for training of the span-prediction systems, and is considered upper-bound.³ *Seq2Seq (no context)* is an encoder-decoder RNN model trained only on the question. *ASR* is a version of the Attention Sum Reader (Kadlec et al., 2016) implemented as a pointer-generator that reads the question and points to words in the context that are contained in the answer. *BiDAF* is Bi-Directional Attention Flow (Seo et al., 2017) trained either with the *Oracle (original)* or *Oracle (ours)*. The models from **Previous Work** are described in Section 5. In the last section of Table 1 we present the results of our experiments (**This work**). Here, *BiDAF* and *QANet* are implementations available in the AllenNLP framework (Gardner et al., 2018). In the last two rows we give the results of QANet extended with the proposed Discourse-Aware Se-

³The previous work that uses span-prediction models do not report their *Oracle* model used for training supervision.

matic Self-Attention, using intra-sentential, *Explicit* discourse relations (*DR (Exp)*, EMA is Exponential Moving Average).

6.2 Fine-grained Evaluation

We further analyze the performance of different configurations of our model by conducting fine-grained evaluation in view of question types (Figure 4) and context length (Figure 5).

We define a range of system configurations using attention heads enhanced with different combinations of linguistic annotation types, including *Explicit* (referred to as Exp or E) and *Non-Explicit* (NonE, NE), *Discourse Relations* (DiscRel, DR), *Semantic Role Labeling* (SRL), and *Coreference* (Coref), and configurations without any such additional information (*No*). We also experiment with a setting where instead of using specific discourse relation types (such as DiscRel_Exp_Cause_Arg1), we only identify that a token is part of **any** (NoSense) discourse relation (e.g., DiscRel_Exp_Arg1) or simply a multi-sentence attention span *Sent span 3* with labels *Sent1*, *Sent2*, *Sent3* for each sentence. This is

	Absolute		Improvement over QANet baseline													
200 - 400	71.54	50.89	-0.12	0.84	1.63	3.04	0.90	1.17	0.33	2.13	1.45	2.29	0.95	0.25	0.50	0.92
400 - 600	69.69	48.73	1.79	1.28	1.35	3.39	0.26	0.96	1.40	1.11	2.28	1.68	0.98	0.78	1.95	1.33
600 - 800	70.86	48.72	0.49	1.47	0.87	2.80	0.50	-0.56	1.73	1.23	0.67	1.41	0.81	0.98	0.39	0.62
800 - 1000	70.54	46.31	1.35	1.42	1.90	4.23	1.64	1.02	1.64	1.73	1.40	1.71	2.07	2.50	1.44	1.79
1000 - 1200*	73.05	51.66	1.96	0.19	0.33	3.99	2.45	0.48	1.07	2.70	3.60	3.34	3.51	2.16	2.27	1.65

Oracle QANet BiDAF Coref DR (All) DR (Exp) DR (NonE) DR (Exp NoSense) DR (NonE NoSense) Sent span 3 SRL SRL + DR (Exp) SRL + DR (NonE) SRL + DR (All) SRL + DR (Exp) + Coref SRL + DR (All) + Coref

Figure 5: Rouge-L performance by context length on the NarrativeQA *Test* set. The first two columns represent *Absolute* values. The rest are improvements over the QANet baseline model (i) by BiDAF and (ii) different configurations of QANet with linguistic information. Rows with * have less than 100 instances in the *Test* set.

to examine whether the type of discourse relation is important or rather the attention scope (intra-sentential, cross-sentence - 2, 3 neighbouring sentences, full context).

Question Type Different question types might profit from different linguistic annotation types. We thus examine the performance of different question types, and analyze how it correlates with the presence of specific Semantic Self-Attention signals. We classify the questions into question types using a simple heuristic based on the question words as an indicator of their type (*How / Where / Why / Who / What ...*), and calculate the average *Rouge-L* for each such questions type. The resulting scores are displayed in Figure 4. In the first two columns of the figure, we report the *Oracle* score and the baseline (QANet) score. In the remaining columns we report (i) the improvement over the QANet baseline of BiDAF, and (ii) of our models with different combinations of discourse-aware semantic self-attention. In the first row we report the score for each of the models on *all* questions. We observe that best performing models on *all* questions are the ones that include *Explicit* DR, and/or SRL. In terms of hardness, *how* and *why* questions usually have the lowest score. This not surprising since *Oracle* performance is also low. For these type of questions, the RNN-based encoder (BiDAF) and self-attention with DR (Exp) or DR (NonE) perform best. Almost all models with additional linguistic information improve over the baseline on *when* questions, lead by the SRL+DR (Exp) and SRL + DR (All) + Coref. *What* questions are improved most by DR (Exp) and SRL alone or when combined. *Who* questions gain most from discourse relations and all models that contain SRL.

Context Length In Figure 5 we present the performance on documents of different lengths, in number of tokens. All presented models are trained on the examples from the *Train* set with context up to 800 tokens. Again, the models DR (Exp) and SRL+DR (Exp) show clear improvement across all context lengths. It is clear that all models show improvement over length *800-1000*. This supports our hypothesis that discourse information is required for generalizing to longer contexts. One reason is that some of the questions can be answered with a local context (one-two sentences) which are better represented given short discourse scope (one-three sentences) or long dependencies given coreference.

In the evaluation of multiple model configurations we notice that in some cases a single discourse/semantic type (e.g. DR (Exp)) performs better than in combination with others (e.g. SRL+DR (Exp)). We hypothesize that the reason is that the linguistic annotations work well in combination with free *No* attention heads (see Table 2). Currently, we place multiple annotations on the same *Encoder Block* which reduces the number of free attention heads. For instance, for SRL+DR (Exp), each knowledge-enhanced encoder block has 3 SRL + 2 DR (Exp) + 3 No heads. In future work we plan to use different annotation heads per *Encoder Block (EB)*: e.g., *EB0* has 3 SRL + 5 No; *EB1* has 2 DR (Exp) + 6 No; etc.

Success and Failure Examples In Figures 6, 7, 8 we show examples of context⁴ and questions, together with the answers from human annotators and some of the examined models.⁵ We provide

⁴The part that contains the correct answer.

⁵For easier reading, we color the **gold**, **correct**, and **wrong** answers and underline the mentions of different characters.

Context	Although <u>he</u> terrifies the fairies when <u>he</u> first arrives , <u>Peter</u> quickly gains favour with <u>them</u> . <u>He</u> amuses <u>them</u> with <u>his</u> human ways and agrees to play the panpipes at the <u>fairly</u> dances . Eventually , <u>Queen Mab</u> grants <u>him</u> the wish of <u>his</u> heart , and <u>he</u> decides to return home to <u>his</u> mother .
Question	After scaring the fairies, how does <u>Peter</u> win them over ?
Human 1:	<u>he agrees to play the panpipes at all of the fairy dances.</u> ;
Human 2:	<u>He amuses them with his human ways and plays the pipes at their dances.</u> ;
Oracle:	human ways and agrees to play the panpipes at the fairy dances ;
QANet:	gains favour ;
DR (Exp), DR (NE):	quickly gains favour with them;
Coref, SRL, SRL+DR(Exp):	He amuses them with his human ways and agrees to play the panpipes;
SRL+DR(NE):	He amuses them with his human ways and agrees to play the panpipes at the fairy dances
Rationale:	To find the correct answer we need to know that (i) ‘gains favor’ is a synonym to ‘win’ in this context (commonsense); (ii) the following (2nd) sentence is the reason for the previous (1st) (DR - the model fails in this case) (iii) ‘them’ are ‘the fairies’, ‘he’ is <u>Peter</u> (Coref)

Figure 6: Example of positive impact of SRL and Coref and negative impact from discourse relations (DR).

Context	<u>Jacob</u> frequently visits <u>Jeff</u> and <u>Kenny</u> , <u>who</u> are serving time in a juvenile hall . <u>Jacob</u> initially threatens <u>them</u> , until eventually <u>Jeff</u> commits suicide . <u>Jacob</u> befriends <u>Kenny</u> , soon learning <u>he</u> has an early release and is illegally moving to New Mexico .
Question	Why does <u>Jeff</u> committ suicide ?
Human 1:	<u>Jacob threatened them;</u>
Human 2:	<u>He is threatened by Jacob.</u> ;
Oracle:	site which he says is ;
QANet:	<u>Jeff and Kenny , who are serving time in a juvenile hall;</u>
DR (Exp), DR (NE), SRL, SRL+DR(Exp), SRL+DR(NE):	<u>Jacob initially threatens them ;</u> ;
Coref:	<u>Jacob initially threatens them , until eventually Jeff commits suicide . Jacob befriends Kenny , soon learning he has an early release and is illegally moving to New Mexico</u>
Rationale:	To find the correct answer we need to understand that ‘until eventually’ suggests that the suicide of <u>Jeff</u> is caused by <u>Jacob</u> threatening ‘them’ (DR) and that <u>Jeff</u> is part of ‘them’ (Coref).

Figure 7: Example of positive impact of SRL and Coref, and discourse relations (DR).

Context	The four orphan children of the house , <u>Edward</u> , <u>Humphrey</u> , <u>Alice</u> and <u>Edith</u> , are believed to have died in the flames . However , <u>they</u> are saved by <u>Jacob Armitage</u> , a local verderer , <u>who</u> hides <u>them</u> in <u>his</u> isolated cottage and disguises <u>them</u> as <u>his</u> grandchildren . Under <u>Armitage</u> ’s guidance , <u>the children</u> from an aristocratic lifestyle to that of simple foresters .
Question	Who rescues the children from fire at Arnwood ?
Human 1, Human 2:	<u>Jacob Armitage;</u>
Oracle:	<u>Jacob Armitage;</u>
DR (Exp), DR (NE), Coref:	<u>Jacob Armitage;</u>
QANet, SRL, SRL+DR(Exp):	<u>Pablo;</u>
SRL+DR(NE):	<u>Patience</u>
Rationale:	To find the correct answer we need to understand at least that ‘they’ are ‘the children’ (Coref) and ‘who did what to whom’ in the context (SRL).

Figure 8: Example of positive impact of Coref and DR and negative impact from SRL.

a hypothetical rationale of what we would need to answer the question. ⁶


7 Conclusion and Future Work

In this work we use linguistic annotations as a basis for a *Discourse-Aware Semantic Self-Attention* encoder that we employ for reading comprehension on narrative texts.

The provided annotations of discourse relations, events and their arguments as well as coreferring mentions, are using available annotation tools. Our empirical evaluation shows that discourse-semantic annotations combined with self-attention yields significant (+3.43 *Rouge-L*) improvement over QANet’s token-based self-attention when applied to NarrativeQA reading comprehension. We analyzed the impact of different semantic annota-

tion types on specific question types and context regions. We find, for instance, that SRL greatly improves *who* and *when* questions, and that discourse relations improve also the performance on *why* and *where* questions. While all examined annotation types contribute, particularly strong and constant gains are seen with intra-sentential DR (all context ranges), followed by SRL (short to mid-sized contexts). Coreference shows positive, but weaker impact, mostly in mid-sized contexts. A promising future direction would be to include additional external knowledge such as commonsense and world knowledge, and learn all annotations jointly with the downstream task.

Acknowledgments This work has been supported by the German Research Foundation as part of the Research Training Group Adaptive Preparation of Information from Heterogeneous Sources (AIPHES) under grant No. GRK 1994/1. We thank NVIDIA Corporation for donating GPUs used in this research.

⁶  The examples are selected from NarrativeQA *Test*, in such a way, that they depict the strength and weaknesses of the different models, corresponding to the empirical evaluation on Figure 4 and they fit in the space limit.

References

- Lisa Bauer, Yicheng Wang, and Mohit Bansal. 2018. Commonsense for generative multi-hop question answering tasks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4220–4230, Brussels, Belgium.
- Danqi Chen, Jason Bolton, and Christopher D. Manning. 2016. A thorough examination of the cnn/daily mail reading comprehension task. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 2358–2367, Berlin, Germany.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading Wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 1870–1879, Vancouver, Canada.
- François Chollet. 2016. Xception: Deep learning with depthwise separable convolutions. *CoRR*, abs/1610.02357.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy.
- Rajarshi Das, Manzil Zaheer, Siva Reddy, and Andrew McCallum. 2017. Question answering on knowledge bases and text using universal schema and memory networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 358–365.
- 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.
- 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.
- Bhuwan Dhingra, Zhilin Yang, William W. Cohen, and Ruslan Salakhutdinov. 2017. Linguistic knowledge as memory for recurrent neural networks. volume abs/1703.02620.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. AllenNLP: A deep semantic natural language processing platform. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia.
- Luheng He, Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2017. Deep semantic role labeling: What works and what’s next. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 473–483, Vancouver, Canada.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 1693–1701. Curran Associates, Inc.
- Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2016. The goldilocks principle: Reading children’s books with explicit memory representations. In *ICLR*.
- Minghao Hu, Yuxing Peng, Zhen Huang, Xipeng Qiu, Furu Wei, and Ming Zhou. 2018a. Reinforced mnemonic reader for machine reading comprehension. In *Proceedings of International Joint Conferences on Artificial Intelligence Organization*, pages 4099–4106.
- Minghao Hu, Yuxing Peng, Furu Wei, Zhen Huang, Dongsheng Li, Nan Yang, and Ming Zhou. 2018b. Attention-guided answer distillation for machine reading comprehension. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2077–2086, Brussels, Belgium.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics (ACL) 2017.
- Rudolf Kadlec, Martin Schmid, Ondřej Bajgar, and Jan Kleindienst. 2016. Text understanding with the attention sum reader network. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 908–918. Association for Computational Linguistics.
- Lukasz Kaiser, Aidan N. Gomez, and François Chollet. 2017a. Depthwise separable convolutions for neural machine translation. *CoRR*, abs/1706.03059.
- Lukasz Kaiser, Aidan N. Gomez, Noam Shazeer, Ashish Vaswani, Niki Parmar, Llion Jones, and Jakob Uszkoreit. 2017b. One model to learn them all. *CoRR*, abs/1706.05137.

- Daniel Khashabi, Tushar Khot, Ashish Sabharwal, and Dan Roth. 2018a. [Question Answering as Global Reasoning over Semantic Abstractions](#). In *AAAI*.
- Daniel Khashabi, Tushar Khot, Ashutosh Sabharwal, and Dan Roth. 2018b. Question answering as global reasoning over semantic abstractions. In *AAAI*.
- Tushar Khot, Ashish Sabharwal, and Peter Clark. 2017. [Answering complex questions using open information extraction](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 311–316.
- Diederik P. Kingma and Jimmy Lei Ba. 2015. Adam: a Method for Stochastic Optimization. *International Conference on Learning Representations 2015*, pages 1–15.
- Tomas Kocisky, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gabor Melis, and Edward Grefenstette. 2018. [The narrativeqa reading comprehension challenge](#). *Transactions of the Association for Computational Linguistics*, 6:317–328.
- Rik Koncel-Kedziorski, Dhanush Bekal, Yi Luan, Mirella Lapata, and Hannaneh Hajishirzi. 2019. [Text Generation from Knowledge Graphs with Graph Transformers](#). pages 2284–2293.
- Yuan Li, Xiaodan Liang, Zhiting Hu, Yinbo Chen, and Eric P. Xing. 2019. [Graph transformer](#).
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81, Barcelona, Spain.
- Inderjeet Mani. 2012. *Computational Modeling of Narrative*, volume 5.
- Todor Mihaylov and Anette Frank. 2016. Discourse relation sense classification using cross-argument semantic similarity based on word embeddings. In *Proceedings of the Twentieth Conference on Computational Natural Language Learning - Shared Task*.
- Todor Mihaylov and Anette Frank. 2018. [Knowledgeable reader: Enhancing cloze-style reading comprehension with external commonsense knowledge](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 821–832, Melbourne, Australia.
- Takeshi Onishi, Hai Wang, Mohit Bansal, Kevin Gimpel, and David McAllester. 2016. [Who did what: A large-scale person-centered cloze dataset](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2230–2235, Austin, Texas. Association for Computational Linguistics.
- Simon Ostermann, Ashutosh Modi, Michael Roth, Stefan Thater, and Manfred Pinkal. 2018. [Mscript: A novel dataset for assessing machine comprehension using script knowledge](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*. European Language Resource Association.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. [The proposition bank: An annotated corpus of semantic roles](#). *Computational Linguistics*, 31(1):71–106.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [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.
- Rashmi Prasad, Nikhil Dinesh, Alan Lee, Eleni Miltasakaki, Livio Robaldo, Aravind Joshi, and Bonnie Webber. 2008. The Penn Discourse treebank 2.0. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC-08)*, Marrakech, Morocco.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018a. [GPT: Improving Language Understanding by Generative Pre-Training](#). *arXiv*, pages 1–12.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018b. Improving Language Understanding by Generative Pre-Training.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Matthew Richardson, Christopher J.C. Burges, and Erin Renshaw. 2013. [MCTest: A challenge dataset for the open-domain machine comprehension of text](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 193–203, Seattle, Washington, USA. Association for Computational Linguistics.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bi-Directional Attention Flow for Machine Comprehension. In *Proceedings of International Conference of Learning Representations 2017*, pages 1–12.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). 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 464–468, New Orleans, Louisiana.

- Lin Feng Song, Zhiguo Wang, Mo Yu, Yue Zhang, Radu Florian, and Daniel Gildea. 2018. [Exploring Graph-structured Passage Representation for Multi-hop Reading Comprehension with Graph Neural Networks](#).
- Emma Strubell, Patrick Verga, Daniel Andor, David Weiss, and Andrew McCallum. 2018. [Linguistically-informed self-attention for semantic role labeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5027–5038.
- Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Kathryn Mazaitis, Ruslan Salakhutdinov, and William Cohen. 2018. [Open domain question answering using early fusion of knowledge bases and text](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4231–4242.
- Kai Sun, Dian Yu, Dong Yu, and Claire Cardie. 2019. [Improving machine reading comprehension with general reading strategies](#). 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 2633–2643, Minneapolis, Minnesota.
- Niket Tandon, Bhavana Dalvi, Joel Grus, Wen-tau Yih, Antoine Bosselut, and Peter Clark. 2018. [Reasoning about actions and state changes by injecting commonsense knowledge](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 57–66.
- Yi Tay, Anh Tuan Luu, and Siu Cheung Hui. 2018a. [Multi-granular sequence encoding via dilated compositional units for reading comprehension](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2141–2151, Brussels, Belgium. Association for Computational Linguistics.
- Yi Tay, Luu Anh Tuan, Siu Cheung Hui, and Jian Su. 2018b. [Densely connected attention propagation for reading comprehension](#). In *Proceedings of the 32Nd International Conference on Neural Information Processing Systems, NIPS'18*, pages 4911–4922.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention Is All You Need](#). In *Nips*.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph Attention Networks](#). *International Conference on Learning Representations*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018a. [Glue: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355.
- Liang Wang, Meng Sun, Wei Zhao, Kewei Shen, and Jingming Liu. 2018b. [Yuanfudao at semeval-2018 task 11: Three-way attention and relational knowledge for commonsense machine comprehension](#). In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 758–762.
- Dirk Weissenborn, Tomas Kocisky, and Chris Dyer. 2017a. [Dynamic integration of background knowledge in neural NLU systems](#). *CoRR*, abs/1706.02596.
- Dirk Weissenborn, Georg Wiese, and Laura Seiffe. 2017b. [Making neural qa as simple as possible but not simpler](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 271–280. Association for Computational Linguistics.
- Kun Xu, Siva Reddy, Yansong Feng, Songfang Huang, and Dongyan Zhao. 2016. [Question answering on Freebase via relation extraction and textual evidence](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 2326–2336, Berlin, Germany.
- Nianwen Xue, Hwee Tou Ng, Sameer Pradhan, Bonnie Webber, Attapol Rutherford, Chuan Wang, and Hongmin Wang. 2016. [The conll-2016 shared task on multilingual shallow discourse parsing](#). In *Proceedings of the Twentieth Conference on Computational Natural Language Learning - Shared Task*, Berlin, Germany. Association for Computational Linguistics.
- Wen-tau Yih, Matthew Richardson, Chris Meek, Ming-Wei Chang, and Jina Suh. 2016. [The value of semantic parse labeling for knowledge base question answering](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 201–206, Berlin, Germany.
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V. Le. 2018. [QANet: Combining Local Convolution with Global Self-Attention for Reading Comprehension](#). In *ICLR 2018*.