

Why can't memory networks read effectively?*

Simon Šuster Madhumita Sushil Walter Daelemans
Computational Linguistics & Psycholinguistics Research Center,
University of Antwerp, Belgium

{simon.suster, madhumita.sushil, walter.daelemans}@uantwerpen.be

Abstract

Memory networks have been a popular choice among neural architectures for machine reading comprehension and question answering. While recent work revealed that memory networks can't truly perform multi-hop reasoning, we show in the present paper that vanilla memory networks are ineffective even in single-hop reading comprehension. We analyze the reasons for this on two cloze-style datasets, one from the medical domain and another including children's fiction. We find that the output classification layer with entity-specific weights, and the aggregation of passage information with relatively flat attention distributions are the most important contributors to poor results. We propose network adaptations that can serve as simple remedies. We also find that the presence of unseen answers at test time can dramatically affect the reported results, so we suggest controlling for this factor during evaluation.

1 Introduction

Recent work in machine reading comprehension and question answering has put a lot of focus on proposing new datasets and developing new reading architectures. Among the latter, memory networks (Weston et al., 2014; Sukhbaatar et al., 2015; Miller et al., 2016) have played a prominent role, spanning the research on general reading comprehension, such as using children's fiction (Hill et al., 2016) and news texts (Kadlec et al., 2016), as well as the research on knowledge integration and reasoning (Kumar et al., 2016; Das et al., 2017; Kaushik and Lipton, 2018; Chen and Durrett, 2019a).¹

*The code is available at <http://bit.ly/2MOK6a4>.

¹Many architectures, often with more complex memory-organization mechanisms, have been inspired by the original memory networks, e.g. Graves et al. (2016), Weissenborn (2016), Xiong et al. (2016), Hudson and Manning (2018).

While the workings of memory networks for reading comprehension and question answering have already been intensively studied, the emphasis has so far gravitated mainly around their ability to perform *multi-hop reasoning*, where the network is expected to combine evidence from different pieces of text when deriving the correct answer. The main conclusion has been that memory networks have great difficulties learning multi-hop reasoning in an end-to-end way (Chen and Durrett, 2019a,b). Chen and Durrett have demonstrated that even after providing a strong supervision signal to the attention component, the results on WikiHop (Welbl et al., 2018) and bAbI (Weston et al., 2015) benchmarks still remained far behind other competing models. Our work complements their work and extends it by focusing on the applicability and effectiveness of memory networks in *single-hop* rather than multi-hop machine reading comprehension. Additionally, whereas Chen and Durrett's work focused primarily on memory network's use of attention and on the proposed solution of attention supervision, we investigate a feature-based approach to give the network a strong attention signal, and also discuss other factors that lead to network's ineffectiveness.

In the first part of the paper, we describe our attempt (§ 3) at establishing competitive results with memory networks on a medical reading comprehension task where other baselines and neural sequential models achieve much better results. Attracted by its transparent and extensible architecture, we had planned to extend memory networks for knowledge integration in the medical domain. Instead, we have found them to be ineffective at plain reading comprehension, so we have set out to explore the reasons behind this in more depth. We identify two major factors. First, the large parametrized output layer suffers from unseen an-

swers at test time, which can be alleviated by either including a pointing layer, which obviates the need to keep weights for each answer entity, or performing entity anonymization, which reduces the size of the output layer and thus sparsity. Second, aggregating vectors of different passage windows into a single representation is harmful. Since the attention weights are distributed rather evenly, the most relevant passage windows for the given query aren’t given enough prominence in the aggregated representation. We can reduce this problem by providing a stronger signal to the network about what passage window has the highest attention weight.

We also have a closer look (§ 4) at the scenario where memory networks have been previously shown to work well, most clearly on Children’s Book Test (CBT) (Hill et al., 2016). We find that we can indeed obtain competitive results using the same architecture and implementation. However, through experimentation we confirm the findings from the literature that the scores on CBT are inflated due to a dataset design bias, and argue that memory networks don’t effectively take into account the full evidence presented to them.

With our analysis, we are adding to the existing body of work that is (re-)examining the established machine reading and language understanding architectures, e.g. Wang et al. (2017); Kaushik and Lipton (2018); Chen and Durrett (2019a). We hope that our results can be of use also more broadly to researchers using other architectures or tasks.

2 Models and datasets

Our cloze-style reading comprehension tasks involve predicting a missing entity as the answer to a query Q , which includes a gap. A support text passage P is given in which the answer can be found.

2.1 Memory networks

We first concisely describe the generic architecture of the memory network (MemNet) (Weston et al., 2014; Sukhbaatar et al., 2015) used in our work. In a window-based representation, the passage P consists of windows around candidate answer entities, and the query Q , which is also represented as a window around the empty slot. The words in the passage windows and the query are mapped into d -dimensional vectors using an embedding matrix $E \in \mathbb{R}^{|V| \times d}$, where $|V|$ is the size of the vocab-

ulary. The query and passage window (memory slot) **encodings** are then obtained by averaging:

$$q = \frac{1}{|Q|} \sum_{i \in Q} q_i; \quad p_j = \frac{1}{|P_j|} \sum_{k \in P_j} p_k, \quad (1)$$

where q_i is the embedding of the word at position i in the query window, and p_k is the embedding of k th word in j th passage window. The words in the passage window include the candidate answer entity.

Next, an attention mechanism measures the **compatibility** between the query representation q and the window representations p_j . We formulate the attention as:

$$\alpha_j = \text{softmax}(\text{cos}(q, p_j)), \quad (2)$$

where cos is the cosine similarity.² This gives us an attention weight for every slot j , indicative of its importance with respect to the query. The attention probabilities are then used to weigh the contribution of each memory slot into an **aggregate output** vector:

$$o = \sum \alpha_j p'_j, \quad (3)$$

where p'_j are obtained using another embedding matrix $E' \in \mathbb{R}^{|V| \times d}$. This matrix is different from E as we would like to distinguish between the representation used for measuring passage-query compatibility and the representation used for answering. In the key-value variant (Miller et al., 2016; Kaushik and Lipton, 2018), p_j only represents a window encoding around the candidate, and p'_j an encoding of the candidate itself. Empirically, this doesn’t lead to improvements in our case, so we don’t use it.

Finally, in a single-hop architecture like ours, we feed the output vector o , the query vector q , and their additive and multiplicative interactions directly to the **output classifier**, which is a softmaxed linear layer:

$$\text{softmax}(W([o; q; o + q; o \odot q]) + b), \quad (4)$$

where $W \in \mathbb{R}^{C \times 4d}$, C is the number of output labels, or answer candidates, and b is the bias term. Although the original architecture feeds only $o + q$ to the output classifier, we have found our approach to be empirically superior.

²We have investigated several other possibilities, including parametrized attention (cf. Seo et al. (2017)), and finally chosen the cosine similarity as the simplest yet a well performing function.

2.2 Medical machine reading: CliCR

The dataset of English clinical case reports (CliCR) (Šuster and Daelemans, 2018) is a gap-filling reading comprehension dataset consisting of around 100,000 queries and their associated documents. The dataset was built from case reports, requiring the machine reader to answer the query with an entity which is either a medical problem, a test or a treatment. Unlike in the original paper, we first pre-process the dataset by recreating the data splits so that those instances whose answer isn't found literally in the passage text are removed. This increases the scores across the board, and, crucially, ensures that pointing-based models can also be applied to the dataset. We rerun the baselines and comparison models from the original paper, and report those in Table 1.

Baselines The baselines include choosing as answer a random entity in the test passage (**random**) and selecting the most frequent passage entity (**max-freq**). We also include a distance-based method that uses pre-trained word embeddings (**sim-entity**), where we vectorize the passage and the query, and then choose that entity from the passage whose representation has the highest cosine similarity to the query representation. These are the same baselines as reported in Šuster and Daelemans (2018).

Other neural readers For comparison, we include several other neural systems. In the Stanford Attentive reader (SARReader) (Chen et al., 2016) and the Gated-Attention reader (GARReader) (Dhingra et al., 2017), the architectures resemble that of MemNet, but the query and the passage are encoded using bidirectional GRUs. While SARReader predicts the answer with a classification layer, as in MemNet, GARReader drops it and bases the answer prediction directly on attention weights using a pointer mechanism (Vinyals et al., 2015; Kadlec et al., 2016; Wang and Jiang, 2017).

We additionally apply two models whose results haven't yet been published on CliCR. The first is BiDAF (Seo et al., 2017), which represents input words using a character-level CNN and pre-trained word embeddings, and then encodes the query and passage with an LSTM. Both query-to-passage and passage-to-query attention is computed, but it is not used to summarize the passage into a fixed-size vector. The resulting vectors from the LSTM and the attention part are then passed

to the recurrent modeling layer. We also apply QANet (Yu et al., 2018), which encodes words similarly as BiDAF, but whose modeling components are based on the idea of stacked encoding blocks, which consist of CNN, self-attention and feed-forward layers. Both BiDAF and QANet predict an answer span using a pointing mechanism.

2.3 Narrative machine reading: Children's Book Test (CBT)

To contrast the results obtained on CliCR with those from another gap-filling dataset, we use Children's Book Test (CBT) Hill et al. (2016), which consists of passages from children's literature in English. Each instance contains 21 consecutive sentences, where the last sentence represents a query with a missing slot. The goal is to fill the slot with either a named entity (NE), a common noun (CN), a verb (V) or a preposition (P), depending on the task. For each query, a list of 10 candidates is already given. We apply the same baselines as with CliCR.

3 Analysis of the failure case

In this section, we analyze the performance of MemNet on CliCR.

We have trained the network for 10 epochs and selected the best performing model based on the development set accuracy. We have tuned the learning rate of Adam (Kingma and Ba, 2014) to $\{0.01, 0.005, \underline{0.001}, 0.0005\}$, the embedding dimensionality to $\{50, \underline{100}, 200\}$, the number of hops to $\{\underline{1}, 2, 3\}$, and the window size to 2 to each side (thus in total 5 tokens including the entity). We have set the size of the memory to 300. The word embeddings were pretrained on a combination of the training set of CliCR and PubMed abstracts (Hakala et al., 2016), amounting in total to over 9 billion tokens. Other details about embedding pretraining can be found in Šuster and Daelemans (2018).

We summarize the results in Table 1 and discuss them in turn in this and the following section. The first thing to note is that although we can fit the training data well with MemNet, the generalization to unseen data remains poor as we only obtain 16.8 F1 on the test set. This is much lower than that for other neural readers, and even lower than that for the maximum frequency and embedding baselines. Simple modifications such as changing the definition of the memory from window-

based to either sentences or key-value pairs have no positive effect. We now present the factors that strongly affect this low performance.

Effect of unseen answers We first have a look at the number of training instances available for answer labels. For those answers at test time that are also observed in the training set, the median frequency is 4 (mean=12, std. dev.=39, as measured on the development set). However, as much as around 60% of all answers in the development set are never observed as answers in the training set.³ To investigate how this affects the results, we subsample a new test set (called *seen*), which satisfies the condition that all data instances must contain answers which are already observed in the training set as entities. We see a considerable improvement on this reduced test set—compared to the original performance of 16.8 F1, we now achieve an F1 of 25.0. This indicates that the low performance can be partly attributed to the unseen answers in the test set. In general, we expect this effect to be amplified in situations exactly like ours, where the output layer is large (about 567 thousand entities). Interestingly, however, the unseen answers do not present a problem to other neural readers, which can be explained by a different way that predictions are achieved in these readers. More precisely, for GAREader, BiDAF and QANet, it is irrelevant at test time whether the true answer has been observed previously or not since the prediction is a **pointer** to the passage and not an answer entity from the fixed output vocabulary. In this way, the prediction is typically achieved based on attention directly (selecting as answer the entity whose window has the highest attention weight), without the need to learn the representation for each answer label in the output space. In the case of SAREader, although it still uses an output layer with dedicated weights for each answer, the widely adopted practice is to **anonymize** the answer entities (Hermann et al., 2015; Chen et al., 2016; Wang et al., 2017). This procedure heavily reduces the output space and reduces the chance of observing a previously unseen answer at test time. It also forces the network to focus more on reading, since the anonymized entities (e.g. @entity₃) become effectively just pointers devoid of any se-

³All entities from the training set build up the possible output answer space. At test time, the correct answer may have occurred as an entity in a training passage, but not as an answer.

mantics, apart from providing co-reference information (Wang et al., 2017). As the results show, anonymization has an overwhelmingly positive effect for SAREader on CliCR.

To verify whether these alternative prediction mechanisms would have a positive effect on MemNet, we first try out a modification to its architecture where we replace the original output layer with a pointing layer, so that the prediction is a candidate from the window with the maximal attention weight. The F1 score in this case increases substantially, to 34.7. To apply anonymization, we keep the original output layer of MemNet but map all entities to numbered symbols, where the numbering is reinitiated for every data instance. For a given passage with n different entity types, the set of entities is thus $\{@entity_0, \dots, @entity_{n-1}\}$. The final number of output labels in the network will correspond to the highest n encountered in the training set. In our case, this results in an output layer with a size of only 443 labels, where almost all are observed during training. With anonymization, the F1 score only increases marginally, from 16.8 to 18.1.

The above results suggest that the output layer as included in MemNet by default may not always be the best option, and that the presence of unseen answers at test time strongly influences the results.

The role of attention Since the representations from MemNet’s compatibility component are aggregated prior to being passed to the output layer, this could lead to blurring of information from several passage windows, especially if the attention weights aren’t strongly peaked. In practice, we find that with a trained anonymized model, the average maximum attention probability is only 0.016 on the test set, with absolute variance from the mean at 0.009. This tells us that the attention distribution generally isn’t strongly peaked towards any particular passage window.

The results of MemNet-pointer and the embedding-similarity baseline suggest that contextual similarity between the query and a *single* passage window should provide a relatively reliable clue about the correct answer, so it is reasonable to expect that a stronger signal from attention could help us to train a more competitive model. We therefore explore two options. The first uses a one-hot feature to indicate the best candidate (**attention-feat**):

$$\varphi \in \{0, 1\}^C : \sum_{i=1}^C \varphi_i = 1, \quad (5)$$

$$\varphi_i = \begin{cases} 1 & \text{if } i = \text{cand}(\text{argmax}_j \alpha_j), \\ 0 & \text{otherwise.} \end{cases}$$

Here, the feature vector φ_i is one exactly at the index corresponding to the candidate answer whose window attained the maximum attention weight.⁴

This is a more direct way of biasing the model towards solutions preferred by hard selection than using attention supervision (cf. [Chen and Durrett \(2019a\)](#)). The feature vector is finally simply concatenated with other output vectors from eq. (4) to form the output representation. The output layer weights are expanded correspondingly, i.e. to $W \in \mathbb{R}^{C \times (4d+C)}$. Table 1 shows that the attention feature significantly boosts the results, from 18.1 to 33.9 F1. When we use pretrained embeddings, this advantage is even greater at 36.7 F1. Clearly, the network can use the added feature to good effect, while aggregating without the added attention signal appears to be inadequate for good performance. An obvious next question is whether the aggregated information can contribute positively to the results *at all*. If we drop the aggregated output vector and only keep the feature vector in the output layer (**attention-feat-only**), we observe a further improvement by two points, which could mean that the aggregated output vector is confusing the network, at least on these simpler instances which can be solved by using the attention feature alone.

Yet another possibility of adding a strong attention signal to the output layer, but without the use of the attention feature, is to avoid aggregation by feeding only the single-best passage window vector to the output layer (**best-window**):

$$o = p'_i, \text{ where } i = \text{argmax}_j \alpha_j. \quad (6)$$

This technique was also used by [Hill et al. \(2016\)](#), who call it *self-supervision*. We see here that we achieve results that are almost identical to those obtained with the attention feature. We can conclude that what appears to matter the most to the memory network is some indication of the best window, be it an indicator feature or the embedded

⁴We also tried using soft probabilities as values of the feature vector, but this had a less positive effect.

Model	EM	F1
MemNet	12.8	16.8
- <i>seen</i>	21.0	25.0
- ten-cands	23.9	26.3
- ten-cands (<i>seen</i>)	55.7	57.5
- pointer	27.5	34.7
- anonymization	13.3	18.1
- attention-feat	27.8	33.9
- attention-feat (+emb-pretr)	30.3	36.7
- attention-feat-only (+emb-pretr)	32.2	38.7
- best-window	30.3	36.8
- best-window (+emb-pretr)	29.0	35.4
QueryClassifier	10.7	15.1
- <i>seen</i>	16.8	21.4
random	2.7	6.9
max-freq	14.6	19.0
sim-window [†]	35.7	43.7
Stanford Attentive Reader [†]		
- anonymization	34.0	41.3
- no-anonymization	10.8	15.8
Gated-attention reader [†]		
- anonymization	42.2	50.0
- no-anonymization	38.2	45.6
QANet	39.5	46.5
BiDAF	44.7	52.8

Table 1: **CLiCR** results. Results indicated with [†] are for models from [Šuster and Daelemans \(2018\)](#). Results are reported for the test split with those questions not having an answer in the passage removed. The use of pretrained word embeddings for the model variants in the upper part of the table is indicated with *+emb-pretr*, with randomly initialized word embeddings used otherwise. In the bottom part of the table, all neural models and *sim-window* use pretrained embeddings. Finally, for QANet, the data splits are smaller in size to prune out long sequences which led to memory issues.

window itself. A simple aggregate, however, is far inferior since the query-weighted passage representation gets too diffused as a consequence of flat attention weights. This probably means that the aggregate representation ends up resembling just an average over all passage windows.

Overall, despite the fact that these modifications improve the network’s performance considerably, the memory network remains a poor competitor to other machine reading systems in Table 1. The BiDAF reader, which adds additional complexity to the attention mechanism in comparison to SAREader and GARReader, performs the best on CLiCR.⁵

⁵Although the readers that we compare to contain a larger number of parameters than MemNet, we don’t observe reliable improvements when parametrizing MemNet more heavily (increasing the sizes of embedding and output matrices, and introducing a weight matrix into the attention function).

4 Contrasting with Children’s Book Test

While our analysis from the previous section sheds some light on the workings of the memory network on CliCR, we now include in our analysis also the CBT dataset, and compare some of its characteristics to CliCR. Strong performance has been previously shown on CBT using memory networks (Hill et al., 2016). With our implementation that is identical to the one we have used on CliCR,⁶ we obtain similarly strong performance as Hill et al. (2016) and Kaushik and Lipton (2018). The results are included in Table 2.

We now discuss why the basic network without any modifications appears to work on CBT and not on CliCR.

Effect of unseen answers The first thing we note is that the performance on CBT varies a great deal across different parts of the dataset. The network does worst on the named entities part (NE). The CBT dataset statistics in Table 3 show that for named entities, only around *a half* of the test answer types are observed in the training set, whereas for other parts this percentage is much higher.

Model	NE	CN	V	P
random	10.0	10.0	10.0	10.0
max-freq	35.2	28.1	29.5	27.7
sim-window	31.4	26.0	30.7	23.2
QueryClassifier	43.1	45.3	67.2	63.5
- <i>seen</i>	53.3	46.2	68.1	63.5
MemNet	44.5	48.2	67.5	63.2
- <i>seen</i>	57.1	49.0	68.1	63.2
- Hill et al. (2016)	49.3	55.4	69.2	67.4
- Kaushik and Lipton (2018)	35.0	37.6	52.5	55.2

Table 2: Accuracy in % on **CBT** for different dataset parts of the test set, using exact match. All MemNet models use random word embeddings.

	NE	CN	V	P
test size	2500	2500	2500	2500
ans types in train	5078	4253	2322	85
ans types in test	422	696	406	43
% found in train	56	93	93	100

Table 3: CBT statistics, with a low proportion of named entities encountered as choices in the training set highlighted in bold.

⁶We have set the network’s hyperparameters to the same values as on CliCR. The size of the memory was adjusted in such a way that all passage windows are used in training for each dataset part.

Number of candidates Furthermore, an obvious difference between the two datasets is that there are always ten candidates given in CBT, whereas in CliCR all entities in the passage are possible candidates (112 on average). To see the effect of reducing the pool of candidates on CliCR, we cap the number of possible candidates to ten (as in CBT), thus including nine random candidates from the passage plus the correct answer.⁷ This results in a dramatically higher score of 57.5 F1 compared to 25 F1 on the *seen* test set of CliCR (Table 1). Of course, pre-selecting candidates makes the task less natural as we rely on an external oracle that shortlists the candidates. This practice, as performed on CBT, should generally be avoided as it is unreasonable to expect to have a strong test-time system for winnowing the candidates down without performing proper reading comprehension (cf. Chen and Durrett (2019b) for a similar observation).

Query-only model Is the network exploiting a more direct correspondence between queries and answers in CBT? Prior work has shown that a baseline predictor which excludes passage information can return strong results on some datasets (Chen and Durrett, 2019a,b; Anand et al., 2018). Our results confirm those of Kaushik and Lipton (2018), where the query-only baseline reaches or surpasses the performance of the full MemNet. This can explain why we see (seemingly) competitive results with MemNet. Despite a large proportion of unseen answers and the aggregation of window vectors, the network can still draw on the association between queries and answers, which constitutes a CBT’s construction bias.

On CliCR, the original MemNet performance is already poor from the start, so we wouldn’t expect to see a strong result when using the query input alone. Indeed, the query classifier obtains only 15.1 F1, with a gap to MemNet being more pronounced on the *seen* test set (21.4 vs. 25.0 F1). In this respect, CliCR appears to be more difficult, as the network can’t exploit the query-answer correspondences to obtain good results. Of course, the dataset could still contain other artefacts we are presently unaware of (Feng et al., 2019).

To return to Kaushik and Lipton (2018)’s study, their main finding in relation to the CBT dataset

⁷It would also be possible to carry out an experiment where we adversarially add more false candidates to CBT and then stress-test on it.

has been that it contains a bias where the query and, to a somewhat lesser extent, the passage can be used in isolation to predict answers with an unexpectedly high level of accuracy. We suspect that their results reveal not only the dataset bias, but also MemNet’s poor reading capability, especially pertaining to measuring query-passage compatibility, and reworking this information successfully into an answer. Specifically, we see that other readers in [Kaushik and Lipton \(2018\)](#) obtain much higher results than the memory network, despite the bias. For example, GAREader records a forty-point advantage on the NE part, and 33 points on the CN part. This advantage might be explained by different text encoders of the two networks, which is a GRU in GAREader and an embedding layer in the memory network. Also, the difference between the query-only model and the full model is 24.3 points (NE) for GAREader, but as little as 5.9 points for the memory network. We believe this points to memory network’s modest capability of accounting for the passage in light of the query.

5 Conclusion

Our analysis reveals that using a vanilla variant of memory networks doesn’t lead to competitive results on single-hop machine reading comprehension tasks. We have identified two primary reasons for this. First, the use of the output classification layer that keeps weights for every answer candidate is detrimental. This appears to pose a problem in the presence of a large number of unseen answers at test time. We see that when we remove the instances with unseen answers from the test set, the results are improved dramatically. We have found that replacing the original output classification layer with a pointing mechanism or performing answer anonymization are better capable of dealing with the problem of unseen answers. Secondly, the aggregation of different window vectors with barely varying attention weights leads to a poor input representation for the final layer. Incorporating a stronger signal about the most compatible passage window, for example via a one-hot attention feature, leads to a perceivable improvement.

A few remarks about the limitations of our study are in place. The scope of our results covers a generic memory network only. For other memory network-based architectures with more complex encoding and compatibility layers the results

might be different. While using pretrained word embeddings leads to an observable improvement, it is likely that other types of representations, such as pretrained contextualized embeddings from language models, may result in even stronger improvements. Finally, we have carried out the analysis on two cloze-style datasets only, which may limit the generalizability of our findings to other QA and machine reading datasets.

Acknowledgments

This work was carried out in the framework of the Accumulate IWT SBO project (nr. 150056), funded by the government agency for Innovation by Science and Technology.

References

- Ankesh Anand, Eugene Belilovsky, Kyle Kastner, Hugo Larochelle, and Aaron C. Courville. 2018. Blindfold Baselines for Embodied QA. *arXiv preprint arXiv:1811.05013*.
- 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 (Volume 1: Long Papers)*, pages 2358–2367. Association for Computational Linguistics.
- Jifan Chen and Greg Durrett. 2019a. How to learn (and how not to learn) multi-hop reasoning with memory networks. ICLR 2019 Conference Withdrawn Submission <https://openreview.net/pdf?id=B1lf43A5Y7>.
- Jifan Chen and Greg Durrett. 2019b. Understanding dataset design choices for multi-hop reasoning. 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 4026–4032, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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. Association for Computational Linguistics.
- Bhuvan Dhingra, Hanxiao Liu, Zhilin Yang, William Cohen, and Ruslan Salakhutdinov. 2017. Gated-attention readers for text comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1832–1846. Association for Computational Linguistics.

- Shi Feng, Eric Wallace, and Jordan Boyd-Graber. 2019. Misleading failures of partial-input baselines. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5533–5538, Florence, Italy. Association for Computational Linguistics.
- Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, et al. 2016. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538(7626):471.
- Kai Hakala, Suwisa Kaewphan, Tapio Salakoski, and Filip Ginter. 2016. Syntactic analyses and named entity recognition for pubmed and pubmed central — up-to-the-minute. In *Proceedings of the 15th Workshop on Biomedical Natural Language Processing*, pages 102–107. Association for Computational Linguistics.
- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Su-leyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *NIPS*.
- Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2016. The Goldilocks Principle: Reading Children’s Books with Explicit Memory Representations. In *ICLR*.
- Drew A Hudson and Christopher D Manning. 2018. Compositional attention networks for machine reasoning. In *ICLR*.
- 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.
- Divyansh Kaushik and Zachary C. Lipton. 2018. How much reading does reading comprehension require? a critical investigation of popular benchmarks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5010–5015, Brussels, Belgium. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Victor Zhong, Romain Paulus, and Richard Socher. 2016. Ask me anything: Dynamic memory networks for natural language processing. In *International Conference on Machine Learning*, pages 1378–1387.
- Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. 2016. Key-value memory networks for directly reading documents. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1400–1409. Association for Computational Linguistics.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *ICLR*.
- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end memory networks. In *NIPS*.
- Simon Šuster and Walter Daelemans. 2018. CliCR: a dataset of clinical case reports for machine reading comprehension. 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 1551–1563, New Orleans, Louisiana. Association for Computational Linguistics.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *NIPS*.
- Hai Wang, Takeshi Onishi, Kevin Gimpel, and David McAllester. 2017. Emergent predication structure in hidden state vectors of neural readers. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*. Association for Computational Linguistics.
- Shuohang Wang and Jing Jiang. 2017. Machine comprehension using match-1stm and answer pointer. In *ICLR*.
- Dirk Weissenborn. 2016. Separating answers from queries for neural reading comprehension. *arXiv preprint arXiv:1607.03316*.
- Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. Constructing datasets for multi-hop reading comprehension across documents. *Transactions of the Association for Computational Linguistics*, 6:287–302.
- Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. 2015. Towards AI-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*.
- Jason Weston, Sumit Chopra, and Antoine Bordes. 2014. Memory networks. *arXiv preprint arXiv:1410.3916*.
- Caiming Xiong, Stephen Merity, and Richard Socher. 2016. Dynamic memory networks for visual and textual question answering. In *International conference on machine learning*, pages 2397–2406.
- 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*.