

Text Understanding with the Attention Sum Reader Network

Rudolf Kadlec, Martin Schmid, Ondrej Bajgar & Jan Kleindienst *

IBM Watson

V Parku 4, Prague, Czech Republic

{rudolf_kadlec, martin.schmid, obajgar, jankle}@cz.ibm.com

Abstract

Several large cloze-style context-question-answer datasets have been introduced recently: the CNN and Daily Mail news data and the Children’s Book Test. Thanks to the size of these datasets, the associated text comprehension task is well suited for deep-learning techniques that currently seem to outperform all alternative approaches. We present a new, simple model that uses attention to directly pick the answer from the context as opposed to computing the answer using a blended representation of words in the document as is usual in similar models. This makes the model particularly suitable for question-answering problems where the answer is a single word from the document. Our model outperforms models previously proposed for these tasks by a large margin.

1 Introduction

Most of the information humanity has gathered up to this point is stored in the form of plain text. Hence the task of teaching machines how to understand this data is of utmost importance in the field of Artificial Intelligence. One way of testing the level of text understanding is simply to ask the system questions for which the answer can be inferred from the text. A well-known example of a system that could make use of a huge collection of unstructured documents to answer questions is for instance IBM’s Watson system used for the Jeopardy challenge (Ferrucci et al., 2010).

Cloze style questions (Taylor, 1953), i.e. questions formed by removing a phrase from a sentence, are an appealing form of such questions. While the task is easy to evaluate, one can vary the context, the question sentence or the specific phrase missing in the question to dramatically change the task structure and difficulty.

One way of altering the task difficulty is to vary the word type being replaced, as in (Hill et al., 2015). The complexity of such variation comes from the fact that

the level of context understanding needed in order to correctly predict different types of words varies greatly. While predicting prepositions can easily be done using relatively simple models with very little context knowledge, predicting named entities requires a deeper understanding of the context.

Also, as opposed to selecting a random sentence from a text (as done in (Hill et al., 2015)), the questions can be formed from a specific part of a document, such as a short summary or a list of tags. Since such sentences often paraphrase in a condensed form what was said in the text, they are particularly suitable for testing text comprehension (Hermann et al., 2015).

An important property of cloze style questions is that a large amount of such questions can be automatically generated from real world documents. This opens the task to data-hungry techniques such as deep learning.

In the first part of this article we introduce the task at hand and the main aspects of the relevant datasets. Then we present our own model to tackle the problem. Subsequently we compare the model to previously proposed architectures and finally describe the experimental results on the performance of our model.

2 Task and datasets

In this section we briefly introduce the task that we are seeking to solve and relevant large-scale datasets that have recently been introduced for this task.

2.1 Formal Task Description

The task consists of answering a cloze style question, the answer to which is dependent on the understanding of a context document provided with the question. The model is also provided with a set of possible answers from which the correct one is to be selected. This can be formalized as follows:

The training data consist of tuples $(\mathbf{q}, \mathbf{d}, a, A)$, where \mathbf{q} is a question, \mathbf{d} is a document that contains the answer to question \mathbf{q} , A is a set of possible answers and $a \in A$ is the ground truth answer. Both \mathbf{q} and \mathbf{d} are sequences of words from vocabulary V . We also assume that all possible answers are words from the vocabulary, that is $A \subseteq V$, and that the ground truth answer a appears in the document, that is $a \in \mathbf{d}$.

*We would also like to thank Tim Klinger for providing us with masked softmax code that we used in our implementation.

	CNN			Daily Mail			CBT CN			CBT NE		
	train	valid	test	train	valid	test	train	valid	test	train	valid	test
# queries	380,298	3,924	3,198	879,450	64,835	53,182	120,769	2,000	2,500	108,719	2,000	2,500
Max # options	527	187	396	371	232	245	10	10	10	10	10	10
Avg # options	26.4	26.5	24.5	26.5	25.5	26.0	10	10	10	10	10	10
Avg # tokens	762	763	716	813	774	780	470	448	461	433	412	424
Vocab. size	118,497			208,045			53,185			53,063		

Table 1: Statistics on the 4 data sets used to evaluate the model. CBT CN stands for CBT Common Nouns and CBT NE stands for CBT Named Entities. CBT had a fixed number of 10 options for answering each question. Statistics were taken from (Hermann et al., 2015) and the statistics provided with the CBT data set.

2.2 Datasets

We will now briefly summarize important features of the datasets.

2.2.1 News Articles — CNN and Daily Mail

The first two datasets¹ (Hermann et al., 2015) were constructed from a large number of news articles from the CNN and Daily Mail websites. The main body of each article forms a context, while the cloze style question is formed from one of short highlight sentences, appearing at the top of each article page. Specifically, the question is created by replacing a named entity from the summary sentence (e.g. “*Producer X will not press charges against Jeremy Clarkson, his lawyer says.*”).

Furthermore the named entities in the whole dataset were replaced by anonymous tokens which were further shuffled for each example so that the model cannot build up any world knowledge about the entities and hence has to genuinely rely on the context document to search for an answer to the question.

2.2.2 Children’s Book Test

The third dataset², the Children’s Book Test (CBT) (Hill et al., 2015), is built from books that are freely available thanks to Project Gutenberg³. Each context document is formed by 20 consecutive sentences taken from a children’s book story. Due to the lack of summary, the cloze style question is then constructed from the subsequent (21st) sentence.

One can also see how the task complexity varies with the type of the omitted word (named entity, common noun, verb, preposition). (Hill et al., 2015) have shown that while standard LSTM language models have human level performance on predicting verbs and prepositions, they lack behind on named entities and common nouns. In this article we therefore focus only on predicting the first two word types.

¹The CNN and Daily Mail datasets are available at <https://github.com/deepmind/rc-data>

²The CBT dataset is available at <http://www.thespermwhale.com/jaseweston/babi/CBTest.tgz>

³<https://www.gutenberg.org/>

Basic statistics about the CNN, Daily Mail and CBT datasets are summarized in Table 1.

3 Our Model — Attention Sum Reader

3.1 Motivation

Our model called the *Attention Sum Reader (AS Reader)* is tailor-made to leverage the fact that the answer is a word from the context document. This is a double-edged sword. While it achieves state-of-the-art results on all of the mentioned datasets (where this assumption holds true), it cannot produce an answer which is not contained in the document. Intuitively, our model is structured as follows:

1. We compute a vector embedding of the query.
2. We compute a vector embedding of each individual word in the context of the whole document (*contextual embedding*).
3. Using a dot product between the question embedding and the contextual embedding of each occurrence of a candidate answer in the document, we select the most likely answer.

3.2 Formal Description

Our model uses one word embedding function and two encoder functions. The word embedding function e translates words into vector representations. The first encoder function is a document encoder f that encodes every word from the document \mathbf{d} in the context of the whole document. We call this the *contextual embedding*. For convenience we will denote the contextual embedding of the i -th word in \mathbf{d} as $f_i(\mathbf{d})$. The second encoder g is used to translate the query \mathbf{q} into a fixed length representation of the same dimensionality as $f_i(\mathbf{d})$. Both encoders use word embeddings computed by e as their input. Then we compute a weight for every word in the document as the dot product of its contextual embedding and the query embedding. This weight might be viewed as an attention over the document \mathbf{d} .

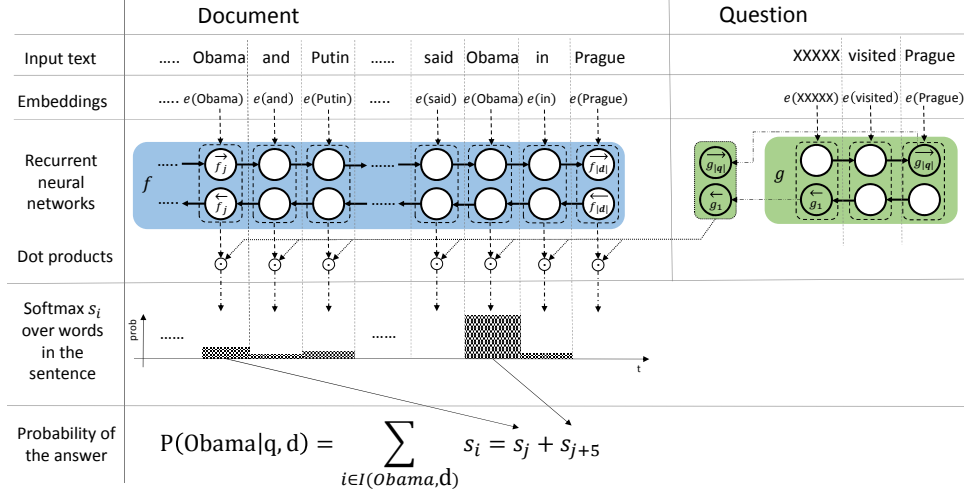


Figure 1: Structure of the model.

To form a proper probability distribution over the words in the document, we normalize the weights using the *softmax* function. This way we model probability s_i that the answer to query \mathbf{q} appears at position i in the document \mathbf{d} . In a functional form this is:

$$s_i \propto \exp(f_i(\mathbf{d}) \cdot g(\mathbf{q})) \quad (1)$$

Finally we compute the probability that word w is a correct answer as:

$$P(w|\mathbf{q}, \mathbf{d}) = \sum_{i \in I(w, \mathbf{d})} s_i \quad (2)$$

where $I(w, \mathbf{d})$ is a set of positions where w appears in the document \mathbf{d} . We call this mechanism *pointer sum attention* since we use attention as a pointer over discrete tokens in the context document and then we directly sum the word's attention across all the occurrences. This differs from the usual use of attention in sequence-to-sequence models (Bahdanau et al., 2015) where attention is used to blend representations of words into a new embedding vector. Our use of attention was inspired by Pointer Networks (PtrNets) (Vinyals et al., 2015).

A high level structure of our model is shown in Figure 1.

3.3 Model instance details

In our model the document encoder f is implemented as a bidirectional Gated Recurrent Unit (GRU) network (Cho et al., 2014) whose hidden states form the contextual word embeddings, that is $f_i(\mathbf{d}) = \vec{f}_i(\mathbf{d}) \parallel \overleftarrow{f}_i(\mathbf{d})$, where \parallel denotes vector concatenation and \vec{f}_i and \overleftarrow{f}_i denote forward and backward contextual embeddings from the respective recurrent networks. The query encoder g is implemented by another bidirectional GRU network. This time the last hidden state of the forward network is concatenated with the last hidden state of the backward network to form the

query embedding, that is $g(\mathbf{q}) = \vec{g}_{|\mathbf{q}|}(\mathbf{q}) \parallel \overleftarrow{g}_1(\mathbf{q})$. The word embedding function e is implemented in a usual way as a look-up table \mathbf{V} . \mathbf{V} is a matrix whose rows can be indexed by words from the vocabulary, that is $e(w) = V_w, w \in V$. Therefore, each row of \mathbf{V} contains embedding of one word from the vocabulary. During training we jointly optimize parameters of f , g and e .

4 Related Work

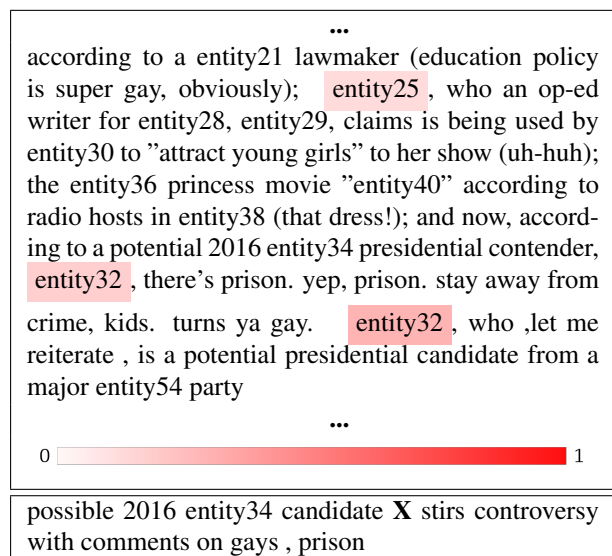
Two recent deep neural network architectures (Hermann et al., 2015; Hill et al., 2015) were applied to the task of text comprehension. Both these architectures use an attention mechanism that allows them to highlight places in the document that might be relevant to answering the question. We will now briefly describe these architectures and compare them to our approach.

4.1 Attentive and Impatient Readers

Attentive and *Impatient Readers* were proposed in (Hermann et al., 2015). The simpler Attentive Reader is very similar to our architecture. It also uses bidirectional document and query encoders to compute an attention in a similar way we do. The more complex Impatient Reader computes attention over the document after reading every word of the query. However, empirical evaluation has shown that both models perform almost identically on the CNN and Daily Mail datasets.

The key difference between the Attentive Reader and our model is that the Attentive Reader uses attention to compute a fixed length representation r of the document \mathbf{d} that is equal to a weighted sum of contextual embeddings of words in \mathbf{d} , that is $r = \sum_i s_i f_i(\mathbf{d})$. A joint query and document embedding is then a non-linear function of r and the query embedding $g(\mathbf{q})$. This joint embedding is in the end compared against all candidate answers $a' \in A$ using the dot product $e(a') \cdot r$, in the end the scores are normalized by softmax. That is: $P(a'|\mathbf{q}, \mathbf{d}) \propto \exp(e(a') \cdot r)$.

Table 2: Attention in an example where our system selected the correct answer. Note that the attention is focused only on named entities.



In contrast to the Attentive Reader, we select the answer from the context directly using the computed attention rather than using such attention for a weighted sum of the individual representations (see Eq. 2). The motivation for such simplification is the following.

Consider a context “Bayern Munchen plays FC Barcelona tonight. Last time it played Liverpool”⁴ and question “X is one of the candidates for tonight’s victory.”

Since both Bayern Munchen and FC Barcelona are equally good candidates, the attention mechanism might put the same attention on both these candidates in the context. Because the resulting representation is computed as a weighted sum when using the above blending mechanism, the vector will represent something in between these two good candidates. It could very well be that the candidate with closest representation to the resulting blended embedding is Liverpool.

By using a sum rather than an average, our architecture would correctly propose Bayern and Barcelona as equally good candidates rather than potentially proposing a third one.

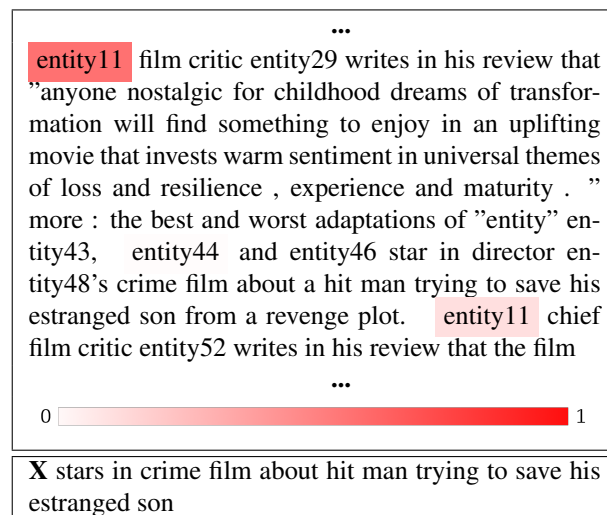
4.2 Memory Networks

MemNNs (Weston et al., 2014) were applied to the task of text comprehension in (Hill et al., 2015).

The best performing memory networks model setup - window memory - uses windows of fixed length (8) centered around the candidate words as memory cells. Due to this limited context window, the model is unable to capture dependencies out of scope of this window.

⁴In the text comprehension task as defined in (Hermann et al., 2015) these named entities will be anonymized. However, this problem can still emerge with anonymous word embeddings.

Table 3: Attention over an example where our system failed to select the correct answer (entity43). The system was probably misled by the co-occurring word ‘film’. Namely, entity11 occurs 7 times in the whole document and 6 times it is together with the word ‘film’. On the other hand, the correct answer occurs only 3 times in total and only once together with ‘film’.



Furthermore, the representation within such window is computed simply as the sum of embeddings of words in that window. By contrast, in our model the representation of each individual word is computed using a recurrent network, which not only allows it to capture context from the entire document but also the embedding computation is much more flexible than a simple sum.

To improve the initial accuracy, a heuristic approach called *self supervision* is used in (Hill et al., 2015) to help the network to select the right supporting “memories”. Plain MemNNs without this heuristic are not competitive on these machine reading tasks. Our model does not need any similar heuristics.

4.3 Pointer Networks

Our model architecture was inspired by Ptr-Nets (Vinyals et al., 2015) in using an attention mechanism to select the answer in the context rather than to blend words from the context into an answer representation. While a Ptr-Net consists of an encoder as well as a decoder, which uses the attention to select the output at each step, our model outputs the answer in a single step. Furthermore, the pointer networks assume that no input in the sequence appears more than once, which is not the case in our settings.

4.4 Summary

Our model combines the best features of the architectures mentioned above. We use recurrent networks to “read” the document and the query as does Attentive and Impatient readers and we use attention in a way similar to Ptr-Nets. We also use summation of atten-

Table 6: Average duration of one epoch of training on the four datasets.

Dataset	Time per epoch
CNN	10h 22min
Daily Mail	25h 42min
CBT Named Entity	1h 5min
CBT Common Noun	0h 56min

tion weights in a way similar to MemNNs (Hill et al., 2015).

5 Evaluation

In this section we evaluate our model on the CNN, Daily Mail and CBT datasets. We show that despite its simplicity the model achieves state-of-the-art performance on each of these datasets.

5.1 Training Details

To train the model we used stochastic gradient descent with the ADAM update rule (Kingma and Ba, 2015) and learning rate of 0.001 or 0.0005. We used negative log-likelihood as the training objective.

The initial weights in the word embedding matrix were drawn randomly uniformly from the interval $[-0.1, 0.1]$. Weights in the GRU networks were initialized by random orthogonal matrices (Saxe et al., 2014) and biases were initialized to zero. The gradient clipping (Pascanu et al., 2012) threshold was set to 10 and the batch size to 32.

During training we randomly shuffled all examples in each epoch. To speedup training, we always pre-fetched 10 batches worth of examples and sorted them according to the length of the document. This way each batch contained documents of roughly the same length.

For the CNN and Daily Mail datasets, for each batch we randomly reshuffled the assignment of named entities to the corresponding word embedding vectors to match the procedure proposed in (Hermann et al., 2015). This guaranteed that word embeddings of named entities were used only as semantically meaningless labels not encoding any intrinsic features of the represented entities. This forced the model to truly deduce the answer from the single context document associated with the question.

During training we evaluated the model performance after each epoch and stopped the training when the error on the validation set started increasing.

The models usually converged after two epochs of training. Time needed to complete a single epoch of training on each dataset on an Nvidia K40 GPU is shown in Table 6.

The hyperparameters, namely the recurrent hidden layer dimension and the source embedding dimension, were chosen by grid search. We started with a range of 128 to 384 for both parameters and subsequently

kept increasing the upper bound by 128 until we started observing a consistent decrease in validation accuracy. The region of the parameter space that we explored together with the parameters of the model with best validation accuracy are summarized in Table 7.

Table 7: Dimension of the recurrent hidden layer and of the source embedding for the best model and the range of values that we tested.

Dataset	Rec. Hid. Layer			Embedding		
	min	max	best	min	max	best
CNN	128	512	384	128	512	128
Daily Mail	128	1024	512	128	512	384
CBT NE	128	512	384	128	512	384
CBT CN	128	1536	256	128	512	384

Our model was implemented using Theano (Bastien et al., 2012) and Blocks (van Merriënboer et al., 2015).

5.2 Evaluation Method

We evaluated the proposed model both as a single model and using ensemble averaging.

For single models we are reporting results for the best model as well as the average of accuracies for the best 20% of models with best performance on validation data since single models display considerable variation of results due to random weight initialization, even for identical hyperparameter values. Single model performance may consequently prove difficult to reproduce.

What concerns ensembles, we used simple averaging of the answer probabilities predicted by ensemble members⁵.

The ensemble models were chosen either as the top 70% of all trained models or using the following algorithm: We started with the best performing model according to validation performance. Then in each step we tried adding the best performing model that had not been previously tried. We kept it in the ensemble if it did improve its validation performance and discarded it otherwise. This way we gradually tried each model once. We call the resulting model a *greedy ensemble*.

5.3 Results

Performance of our models on the CBT dataset is summarized in Table 5, Table 4 shows results on the CNN and Daily Mail datasets. The tables also list performance of other published models that were evaluated on these datasets. Ensembles of our models set the new state-of-the-art results on all evaluated datasets. However, even our best single models outperform the best previously reported results.

⁵Attempts to use the *Constrained Optimization By Linear Approximation* (COBYLA) method (Powell, 1994) to optimize the weights lead to overfitting on the validation data with respect to which the optimization was done.

Table 4: Results of our AS Reader on the CNN and Daily Mail datasets. Results for models marked with \dagger are taken from (Hermann et al., 2015), results of models marked with \ddagger are taken from (Hill et al., 2015). Performance of \ddagger models was evaluated only on CNN dataset.

	CNN		Daily Mail	
	valid	test	valid	test
Deep LSTM Reader \dagger	55.0	57.0	63.3	62.2
Attentive Reader \dagger	61.6	63.0	70.5	69.0
Impatient Reader \dagger	61.8	63.8	69.0	68.0
MemNNs (single model) \ddagger	63.4	66.8	NA	NA
MemNNs (ensemble) \ddagger	66.2	69.4	NA	NA
AS Reader (single model)	68.6	69.5	75.0	73.9
AS Reader (avg for top 20%)	68.4	69.9	74.5	73.5
AS Reader (avg ensemble)	73.9	75.4	78.1	77.1
AS Reader (greedy ensemble)	74.5	74.8	78.7	77.7

Table 5: Results of our AS Reader on the CBT datasets. Results marked with \ddagger are taken from (Hill et al., 2015). (*) Human results were collected on 10% of the test set.

	Named entity		Common noun	
	valid	test	valid	test
Humans (query) (*)	NA	52.0	NA	64.4
Humans (context+query) (*)	NA	81.6	NA	81.6
LSTMs (context+query) \ddagger	51.2	41.8	62.6	56.0
MemNNs (window memory + self-sup.) \ddagger	70.4	66.6	64.2	63.0
AS Reader (single model)	73.8	68.6	68.8	63.4
AS Reader (avg for top 20%)	73.3	68.4	67.7	63.2
AS Reader (avg ensemble)	74.5	70.6	71.1	68.9
AS Reader (greedy ensemble)	76.2	71.0	72.4	67.5

Table 8 then measures accuracy as the proportion of test cases where the ground truth was among the top k answers proposed by the greedy ensemble model for $k = 1, 2, 5$.

CNN. On the CNN dataset our single model with best validation accuracy achieves a test accuracy of 69.5%, this is 0.1% better than ensemble of models reported in (Hill et al., 2015). The average performance of the top 20% models according to validation accuracy is 69.9% which is even 0.5% better than the single best-validation model. This shows that there were many models that performed better on test set than the best-validation model. Fusing multiple models then gives a significant further increase in accuracy. Our simple-average ensemble outperforms the accuracy of the best previously reported ensemble (Hill et al., 2015) by 6%.

Daily Mail. On the Daily Mail dataset our single model with accuracy of 73.9% outperforms the best previous result achieved by the Attentive Reader (Hermann et al., 2015) by 4.9% absolute and our averaging ensemble is by 8.7% absolute better.

CBT. In named entity prediction our best single model with accuracy of 68.6% performs 2% absolute

better than the MemNN with self supervision, the averaging ensemble performs 4% absolute better than the best previous result. In common noun prediction our single models is 0.4% absolute better than MemNN however the ensemble improves the performance to 69% which is 6% absolute better than MemNN.

Table 8: Proportion of test examples for which the top k answers proposed by the greedy ensemble included the correct answer.

Dataset	$k = 1$	$k = 2$	$k = 5$
CNN	74.8	85.5	94.8
Daily Mail	77.7	87.6	94.8
CBT NE	71.0	86.9	96.8
CBT CN	67.5	82.5	95.4

6 Analysis

To further analyze the properties of our model, we examined the dependence of accuracy on the length of the context document (Figure 2), the number of candidate

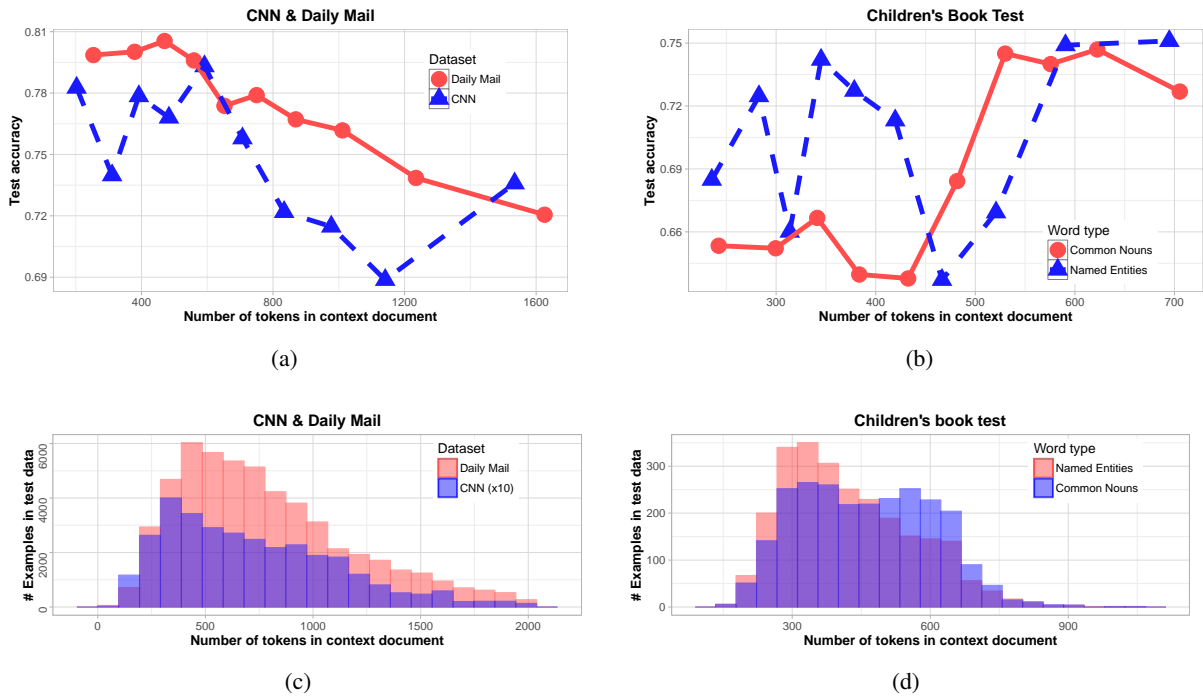


Figure 2: Sub-figures (a) and (b) plot the test accuracy against the length of the context document (for CNN the count was multiplied by 10). The examples were split into ten buckets of equal size by their context length. Averages for each bucket are plotted on each axis. Sub-figures (c) and (d) show distributions of context lengths in the four datasets. The number of examples was multiplied by 10 for the CNN dataset.

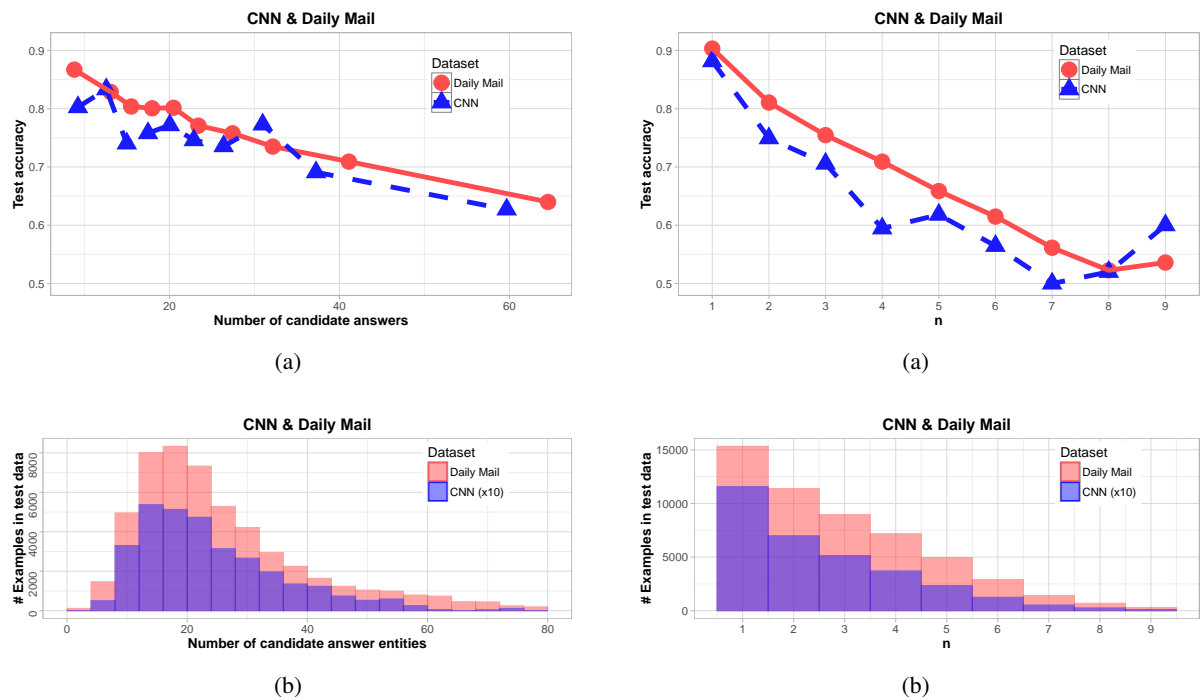


Figure 3: Subfigure (a) illustrates how the model accuracy decreases with an increasing number of candidate named entities. Subfigure (b) shows the overall distribution of the number of candidate answers in the news datasets. The number of examples was multiplied by 10 for the CNN dataset.

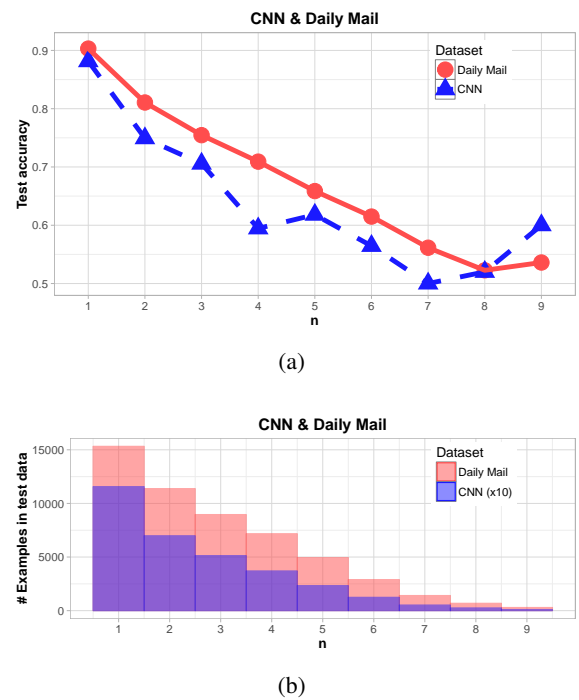


Figure 4: Subfigure (a) shows the model accuracy when the correct answer is among n most frequent named entities for $n \in [1, 10]$. Subfigure (b) shows the number of test examples for which the correct answer was among n most frequent entities. The number of examples was multiplied by 10 for the CNN dataset.

answers (Figure 3) and the frequency of the correct answer in the context (Figure 4).

On the CNN and Daily Mail datasets, the accuracy decreases with increasing document length (Figure 2a). We hypothesize this may be due to multiple factors. Firstly long documents may make the task more complex. Secondly such cases are quite rare in the training data (Figure 2b) which motivates the model to specialize on shorter contexts. Finally the context length is correlated with the number of named entities, i.e. the number of possible answers which is itself negatively correlated with accuracy (see Figure 3).

On the CBT dataset this negative trend seems to disappear (Fig. 2c). This supports the later two explanations since the distribution of document lengths is somewhat more uniform (Figure 2d) and the number of candidate answers is constant (10) for all examples in this dataset.

The effect of increasing number of candidate answers on the model’s accuracy can be seen in Figure 3a. We can clearly see that as the number of candidate answers increases, the accuracy drops. On the other hand, the amount of examples with large number of candidate answers is quite small (Figure 3b).

Finally, since the summation of attention in our model inherently favours frequently occurring tokens, we also visualize how the accuracy depends on the frequency of the correct answer in the document. Figure 4a shows that the accuracy significantly drops as the correct answer gets less and less frequent in the document compared to other candidate answers. On the other hand, the correct answer is likely to occur frequently (Fig. 4a).

6.1 Comparison to Weighted Average Blending

We hypothesized in Section 4.1 that the fact that the Attentive Reader uses attention to create a blended representation potentially harms its performance. In order to verify this intuition, we implemented blending into our model, bringing it closer to the Attentive Reader (Hermann et al., 2015).

In this modified model we compute attention weights s_i in the same way as in our original model (see Eq. 1). However, we replace Eq. 2 with the following equations:

$$r = \sum_i s_i e(w_i) \quad (3)$$

$$P(a' | \mathbf{q}, \mathbf{d}) \propto \exp(r \cdot e(a')) \quad (4)$$

where r is the blended response embedding and $a' \in A$ is a possible candidate response. This change in the architecture indeed lead to a significant decrease in accuracy across all four datasets.

Namely, on each CBT dataset the difference was almost 15% while on CNN and Daily Mail the decrease was over 6% and 2% respectively. Besides the training time for the models with blending was several times

longer than our attention sum architecture both measured by the time per epoch and by the number of epochs required for the model to converge.

7 Conclusion

In this article we presented a new neural network architecture for natural language text comprehension. While our model is simpler than previously published models, it gives a new state-of-the-art accuracy on all the evaluated datasets.

References

- [Bahdanau et al.2015] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. *ICLR*.
- [Bastien et al.2012] Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, James Bergstra, Ian J. Goodfellow, Arnaud Bergeron, Nicolas Bouchard, and Yoshua Bengio. 2012. Theano: new features and speed improvements. Deep Learning and Unsupervised Feature Learning NIPS 2012 Workshop.
- [Cho et al.2014] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. *EMNLP*, jun.
- [Ferrucci et al.2010] David Ferrucci, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya a. Kalyanpur, Adam Lally, J. William Murdock, Eric Nyberg, John Prager, Nico Schlaefler, and Chris Welty. 2010. Building Watson: An Overview of the DeepQA Project. *AI Magazine*, 31(3):59–79.
- [Hermann et al.2015] Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems*, pages 1684–1692.
- [Hill et al.2015] Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2015. The goldilocks principle: Reading children’s books with explicit memory representations. *arXiv preprint arXiv:1511.02301*.
- [Kingma and Ba2015] Diederik P. Kingma and Jimmy Lei Ba. 2015. Adam: a Method for Stochastic Optimization. *International Conference on Learning Representations*, pages 1–13.
- [Pascanu et al.2012] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2012. On the difficulty of training recurrent neural networks. *Proceedings of The 30th International Conference on Machine Learning*, pages 1310–1318.

- [Powell1994] Michael J. D. Powell, 1994. *Advances in Optimization and Numerical Analysis*, chapter A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation, pages 51–67. Springer Netherlands, Dordrecht.
- [Saxe et al.2014] Andrew M Saxe, James L McClelland, and Surya Ganguli. 2014. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *ICLR*.
- [Taylor1953] Wilson L Taylor. 1953. Cloze procedure: a new tool for measuring readability. *Journalism and Mass Communication Quarterly*, 30(4):415.
- [van Merriënboer et al.2015] Bart van Merriënboer, Dzmitry Bahdanau, Vincent Dumoulin, Dmitriy Serdyuk, David Warde-farley, Jan Chorowski, and Yoshua Bengio. 2015. Blocks and Fuel : Frameworks for deep learning. pages 1–5.
- [Vinyals et al.2015] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in Neural Information Processing Systems*, pages 2674–2682.
- [Weston et al.2014] Jason Weston, Sumit Chopra, and Antoine Bordes. 2014. Memory networks. *arXiv preprint arXiv:1410.3916*.