

Reading Comprehension using Entity-based Memory Network

Xun Wang^{1,2}, Katsuhito Sudoh¹, Masaaki Nagata¹,
Tomohide Shibata², Kawahara Daisuke², and Kurohashi Sadao²

¹ NTT Communication Science Laboratories, Kyoto, Japan,
wang.xun, sudoh.katsuhito, nagata.masaaki@lab.ntt.co.jp

² Kyoto University, Kyoto, Japan
shibata,dk, kuro@i.kyoto-u.ac.jp

Abstract. This paper introduces a novel neural network model for question answering, the *entity-based memory network*. It enhances neural networks' ability of representing and calculating information over a long period by keeping records of entities contained in text. The core component is a memory pool which comprises entities' states. These entities' states are continuously updated according to the input text. Questions with regard to the input text are used to search the memory pool for related entities and answers are further predicted based on the states of retrieved entities. Compared with previous memory network models, the proposed model is capable of handling fine-grained information and more sophisticated relations based on entities. We formulated several different tasks as question answering problems and tested the proposed model. Experiments reported satisfying results.

Keywords: Text Comprehension, Entity Memory Network, Question Answering

1 Introduction

A question answering (QA) task is to predict an answer for a given question with regard to related information. It can be formulated as a map $f : \{related_text, question\} \rightarrow \{answer\}$ [10].

QA is an important task in natural language processing (NLP). It is an efficient way of obtaining information from text and a natural way of interacting with computers. More than that, almost all problems in NLP can be formulated as QA tasks. Some ones, like information retrieval and dialog system, are by nature question answering tasks. Other problems, like machine translation, pos tagging, co-reference resolution and so on, can also be formulated as question answering tasks. Take the co-reference resolution for example, given a piece of text, we raise questions like "What does XX refer to?" and expect the system to give correct answers. Similarly, we can model POS tagging as a question answering task by asking "What are the parts of speech?"

Formulating these tasks as QA provides us with the convenience of solving several different problems using one scheme. This point is of vital importance. As we know, existing work on different NLP tasks are highly differentiated, each designed for an (or

a class of) unique task(s) with unique features and unique architectures. It is almost impossible to develop a comprehensive system which can conduct several different tasks without damaging the performance.

Although challenging, such comprehensive systems are of great interest to the artificial intelligence community in their ability of comprehension. Formulating different tasks as QA problems and resolving them using a unified scheme is to some extent, closer to how humans process languages. It differs from previous work in that comprehension of text is needed to serve as the basis for answering various questions.

Developing a comprehensive system with a unified scheme faces challenges in feature representations and model design. Feature representation converts text into features which can be easily computed by models. Models are designed accordingly to process the input features and generate the desired output. A comprehensive model requires us to develop feature representations which are capable of storing all the information contained in text as different tasks may need different information, and to develop a model which is capable of paying attention to different aspects of the information carried by features with regard to the problems raised and generating the desired results. The two are the core challenges to be met and overcome towards a comprehensive system.

Now with the deep neural networks, it becomes probable to develop such a multi-purpose system in a unified scheme. All deep learning models rely on distributed representations representing various features as vectors. These vectors are believed to have encoded all the semantic and syntactic information in themselves. By replacing the various features used in traditional models with vector representations, we can solve the problem of feature representations. But existing deep neural network models are often developed for a certain problem or a certain class of problems. In other words, they are in no sense different from traditional methods in being highly differentiated.

A recent breakthrough is to employ memories in neural networks, the memory network. Memory network and its variants rely on storing sentences in a memory pool for later retrieval to answer questions. One problem with memory networks is that using sentence vectors as elementary units of information makes it impossible to explore fine-grained information.

To overcome the above problems and to fully explore the potential of distributed representations and the neural networks, we introduce a novel model named the *entity-based memory network*. Entities refer to anything that exist in reality or are purely hypothetical. We assume that text can be projected to a world of entities. The key of conducting comprehension and reasoning over text is to identify its containing entities and analyze the states of these entities and the relations between them. The entity-based memory network we proposed is capable of keeping a memory of entities when answering questions.

The proposed model is tested on several datasets, including the toy bAbI dataset [27], large movie review dataset [15] and the machine comprehension test dataset [18]. Results show we have achieved satisfying results using the entity-based memory network. The rest of the paper is organized as follows: Section 2 describes our approaches and elaborates the details. Section 3 reviews previous work that uses memories. Section 4 presents the experiments and the analysis. Section 5 concludes the paper.

2 Approaches

2.1 Overview

Firstly we use an example to illustrate how the model works. Table 1 shows a piece of text which contains 4 sentences and 2 questions. There are 7 entities in total, all of them underlined.

1	<u>Mary</u> moved to the <u>bathroom</u> .
2	<u>John</u> went to the <u>hallway</u> .
3	Where is Mary? <u>Bathroom</u> . 1
4	<u>Daniel</u> went back to the <u>hallway</u> .
5	<u>Sandra</u> moved to the <u>garden</u> .
6	Where is Daniel? <u>Hallway</u> . 4

Table 1. An example from bAbI, a toy dataset for question answering.

This text is elaborated around the 7 entities. It describes how their states change (i.e., the change of a character’s location) when the story goes on. Note that here all the entities are concrete concepts that exist in reality. It is also possible to talk about abstract concepts.

The core of the proposed model are entities. We take Sentence 1 (S_1) as input and extract the entities it contains {Mary, bathroom}. Vectors representing the states of these entities are initialized using some pre-learned word embeddings $\{\overrightarrow{Mary}, \overrightarrow{bathroom}\}$ and stored in a memory pool. Meanwhile, we turn S_1 into a vector ($\overrightarrow{S_1}$) using an autoencoder model³. Then we use the sentence vector $\overrightarrow{S_1}$ to update the entities’ states $\{\overrightarrow{Mary}, \overrightarrow{bathroom}\}$. The goal is to reconstruct $\overrightarrow{S_1}$ solely from $\{\overrightarrow{Mary}, \overrightarrow{bathroom}\}$. In the same way, we process the following text (S_2) and its containing entities (John, hallway) until encounter a question (S_3). S_3 is converted into a vector ($\overrightarrow{S_3}$) following the same method that processes previous input text. Then taking $\overrightarrow{S_3}$ as input, we retrieve related entities from the memory which now stores all the entities (Mary, bathroom, John, hallway) that appear before S_3 . The related entities’ states are then used to produce a feature vector. In this case, (Mary and bathroom) are related to the question and their states are used for constructing the feature vector. Note the current states of the two entities (Mary and bathroom) are different from their initial values due to S_1 . Based on the feature vector, we then use another neural network model to predict the answer to S_3 .

The model monitors the entities involved in text and keeps updating their states according to the input. Whenever we have a question with regard to the text, we check the states of entities and predict an answer accordingly. The proposed model comprises of 4 modules, as is shown in Fig. 1. Each module is designed for a unique purpose and together they construct the *entity-based memory network* model.

³ Note that the sentence vector is not used to answer question directly and it is also plausible to use other models to learn sentence representation.

1. I: Input module. Take as input a sentence and turn it into a vector. Meanwhile, extract all the entities it contains. The question is also processed using this module.
2. G: Generalization module. Update the states of related entities according to the input. For entities that are not contained in the memory pool, create a new memory slot for each of them and initialize these slots using pre-learned word embeddings.
3. O: Output feature module. It is triggered whenever a question arrives. Retrieve related entities according to the input question and then produce an output feature vector accordingly.
4. R: Response module. Generate the response according to the output feature vector.

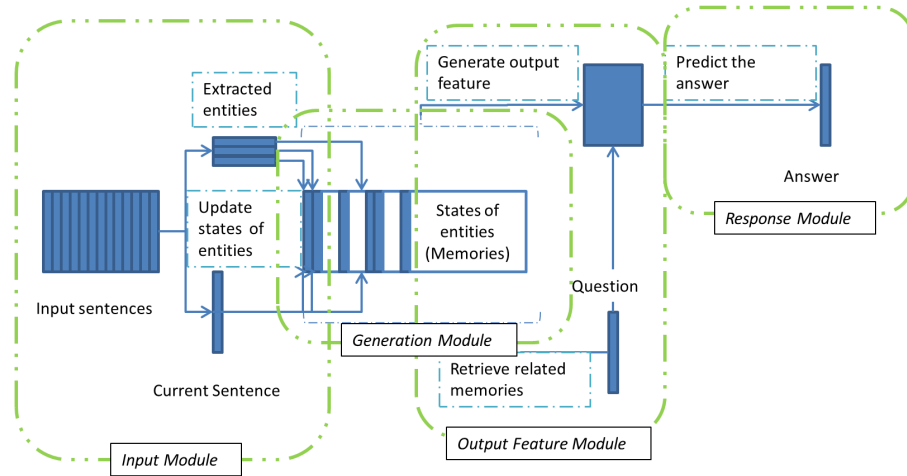


Fig. 1. Architecture of the entity-based memory network. The model is divided into four modules which are shown in the figure using squares.

2.2 Entity-based Memory Network Model

Here we present a formal description of the proposed model. Assume we have text S_1, S_2, \dots, S_n whose entities are annotated in advance as e_1, e_2, \dots, e_m .

Input Module We firstly turn each sentence S_i into its vector representation:

$$S_i = f_1(S_i) \quad (1)$$

Generalization Module For a sentence S_i , we collect all the entities it contains $\{e_1^i, \dots, e_k^i, \dots, e_j^i\}$. These entities' states $\{e_k^i\}$ are simultaneously updated according to S_i as follows:

$$\begin{aligned} S'_i &= f_2(e_1^i, \dots, e_k^i, \dots, e_j^i) \\ \{e_k^i\} &= \arg \min_{\{e_k^i\}} (|S'_i - S_i|) \end{aligned} \quad (2)$$

f_2 is to reconstruct S_i only using the states of S_i 's containing entities $\{e_k^i\}$. $\{e_k^i\}$ are updated to minimize the difference between S'_i and S_i . Recall that S_i is generated using f_1 with the whole sentence S_i as input. We compress the information carried by S_i into a vector S_i and then unfold it into $\{e_k^i\}$.

After processing these sentences, we construct a memory pool which consists of entities whose states are regarded as capable of representing the information carried by the input text.

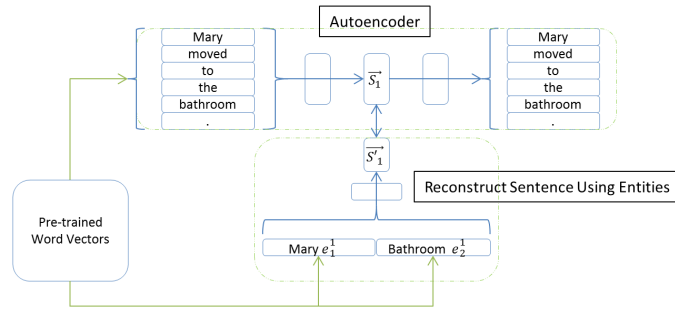


Fig. 2. the Generalization Module: Using S_1 as an example, the autoencoder is used to convert the sentence into a vector S_1 and the entities contained in S_1 are used to reconstruct the sentence vector.

Output Feature Module Question q is turned into a vector $q = f_1(q)$ and then q is used to retrieve related entities from the memory pool.

$$\begin{cases} Q_0 = q \\ Q_{j-1} = g(Q_{j-2}, e_{j-2}) \\ p(e_j, Q_{j-1}) = h(e_j, Q_{j-1}) \\ O_j = u(O_{j-1}, p(e_j, Q_{j-1}) * e_j) \end{cases} \quad (3)$$

$p(e_j, Q_{j-1})$ is the probability (or score) of e_i being selected to compose the feature vector for answering q . In Q , we consider the entity selected in the previous iteration. Q is kept updated using e and p . O_* is the output feature vector.

After several iterations, we use the final O_m as the output feature vector O . Note that if the O_* does not change much between iterations, we will omit the remaining entities. This early-stop strategy helps reduce the time cost.

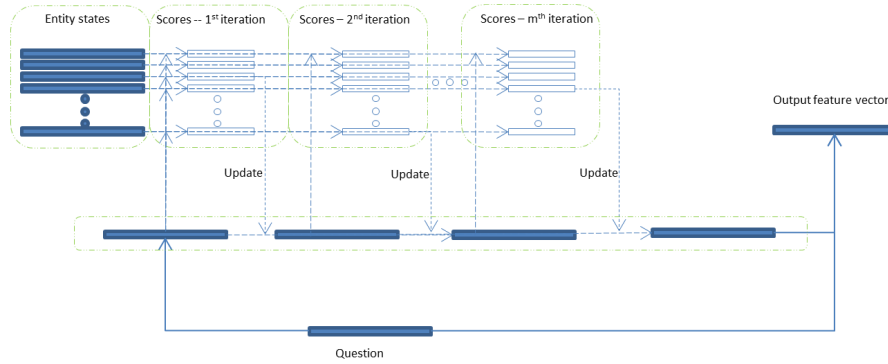


Fig. 3. the Output Feature Module: In each iteration, entities are assigned different scores which indicate their importance in constructing the output feature vector.

Response Module Then we decide the answer using $a(q) = v(O)$. $a(q)$ produces a vector whose each item corresponds to one word in the vocabulary. $a(q)_i$ indicates the probability of $word_i$ being used as the correct answer. We choose the one with the highest probability. Models like recurrent neural network can be used to output a sentence as the answer.

2.3 Implementation

This is a supervised model and requires annotated data for the training. The training data contains the input text, questions and answers. Also we need all the entities and entities that are related to the answer labeled.

We define the function form for training as follows: As for f_1 , many models, like the recurrent neural network, recursive neural network and so on [16,20,9], can be used to convert a sentence into a vector. Here we use an Long Short-Term Memory (LSTM) autoencoder [13] which takes a word sequence as input and outputs the same sequence.

f_2 takes a list of entity states as input and tries to reconstruct S_i . We use the Gated Recurrent Unit (GRU) [2].

$$\begin{aligned} S_i^k &= \tanh(GRU(S_i^{k-1}, e_k^i)) \\ S_i^j &= S_i^j \end{aligned} \quad (4)$$

A GRU can be represented as the follows:

$$\begin{cases} z_t^j = \delta(W_z * x_t + U_z * h_{t-1})^j \\ \bar{h}_t^j = \tanh(W * x_t + U * (r_t \circ h_{t-1}))^j \\ r_t^j = \delta(W_r * x_t + U_r * h_{t-1}) \\ h_t^j = (1 - z_t^j)h_{t-1}^j + z_t^j\bar{h}_t^j \end{cases} \quad (5)$$

◦ represents an element-wise multiplication. z_t^j and r_t^j are two gates controlling the impact of historical h_{t-1}^j on the current h_t^j . The GRU takes x_t as input and updates the state of the neuron to h_t^j . Compared with LSTM which it often replaces, it simplifies the computation while still keeps a memory of previous states. Therefore it takes less time to train GRU than LSTM.

Our goal is to minimize the loss $|S'_i - S_i|$. Using the stochastic gradient descent, we are able to train f_2 and also update $\{e_k^i\}$. Note that the input module and the generalization module do not interact with the remaining. Thus they can be trained in advance.

The output feature module checks the memory pool repeatedly to select entities to form a feature vector:

$$\begin{cases} Q_{j-1} = \tanh(\text{GRU}(Q_{j-2}, e_{j-2})) \\ p(e_j, \mathbf{Q}_{j-1}) = \text{sigmoid}(W * \text{GRU}(e_j, \mathbf{Q}_{j-1}) + b) \\ O_j = \tanh(\text{GRU}(O_{j-1}, p(e_j, \mathbf{Q}_{j-1}) * e_j)) \end{cases} \quad (6)$$

To generate the final answer, we use a simple neural network which takes the feature vector O as input and predict a word as output. $p_w = v(O) = \text{softmax}(\tanh(W' * O + b))$. The word with the highest probability is selected. Suppose a sentence is to be generated, we use the GRU to update O and then generate the sentence $\{w_*\}$ as follows:

$$\begin{cases} \mathbf{p}_w^{i-1} = \text{softmax}(\tanh(W' * O_{i-1} + b)) \\ w_{i-1} = \arg \max \mathbf{p}_w^{i-1} \\ \mathbf{O}_i = \tanh(\text{GRU}(O_{i-1}, \mathbf{w}_{i-1})) \end{cases} \quad (7)$$

Similar to [28], we use the stochastic gradient descent algorithm to minimize the loss function shown in Equation 8 over parameters. For an input S_i and a given question q annotated with the correct answer $word_k$ and related entities $\{e_{c_j}\}$, the loss function is as follows:

$$\sum_{i \neq c_j} \max(0, \gamma - (p(e_{c_j}, q) - p(e_i, q))) + \sum_{l \neq k} \max(0, \gamma - (p_{word_k} - p_{word_l})) + |\Theta^2| \quad (8)$$

Here γ is the margin and $|\Theta^2|$ is the squared sum of all parameters which is used for regularization. Note that Θ does not include parameters of f_1 and f_2 . Their parameters and states of entities are learned as described in Section 2.2. Word vectors used to initialize entity states and words in autoencoder come from GloVe [17]. The dimension is set to be 50.

2.4 Data Annotation

The model requires entities to be annotated in advance. In this work, we treat each noun and pronoun as an entity. Different words are regarded as different entities for simplicity. This strategy saves us the effort of entity resolution which is a challenge for many languages. It also makes possible the application of the proposed model to entity

resolution ⁴. For datasets with related entities annotated, we can use the loss function described above. But annotating the related entities is time and labor-costing. Most datasets available are not annotated. The weakly supervised learning can be applied to such data by trimming the loss function to $\sum_{l \neq k} \max(0, \gamma - (p_{word_k} - p_{word_l})) + |\Theta^2|$. For unannotated data, a fully supervised training is also possible if we regard entities contained in questions as related entities or if we can use other methods to identify entities that are believed to be related.

3 Memories in Deep Neural Networks

In this section we will briefly introduce some neural models considering memories.

3.1 Memory at Neuron Level (LSTM & GRU)

Recurrent neural network is widely adopted in various tasks. It uses hidden states to handle historical inputs. The long short-term memory network (LSTM) is regarded as an improvement of the traditional recurrent neural networks (RNN).

A long short-term memory [5] cell is composed of four main elements: an input gate, a neuron with a self-recurrent connection, a forget gate and an output gate. The input gate controls the impact of the input value on the state of the memory cell and the output gate controls the impact of the state of the memory cell on the output. The self-recurrent connection controls the evolution of the state of the memory cell and the forget gate decides whether to keep or reset the histories of the memory cell's states. These elements serve different purposes and work together to make LSTM cells much more powerful than traditional neural cells. LSTM is widely used in various tasks in NLP and other machine learning fields such as machine translation [22,4], open-domain dialogue generation [11,12,24], or parsing [23].

LSTM uses additional memory units to control the gates. It increases the complexity of the neural networks. The Gated Recurrent Unit (GRU) simplifies the structure and reports performance on par with LSTM but costs less time to train. An empirical analysis of GRU and LSTM model shows that they have similar performance on sequence modeling [2]. In our model, we use the GRUs as an effective method to deal with information over time series. We have an elaborated description about how to use them in our work in Section 2.

3.2 Memory at Layer Level (Memory Network)

In LSTM, the neuron looks at the input in a relatively small scale. Normally the neuron takes one word vector as input each time thus it cannot look beyond the level of words. Layer-level memory networks are designed to keep memories of sentence vectors.

The Memory Network (MNN) [28] contains four parts: the input module which converts sentences into vectors, the memory which keeps all sentence vectors a retrieval

⁴ We treat each mentions of entities as different one when processing the text and ask questions about which of these mentions refer to the same entities.

module and a response module. Whenever a question comes, the question is turned into a vector and the question vector is used to retrieve the memory for related sentences. The response module is used to predict an answer based on the related sentences. The core component is the memory pool that stores all the input sentences so that they can be retrieved later to answer questions. This model contains several neural networks which are jointly optimized according to the task. Experiments on a toy dataset show that this model is able to answer simple questions according to the input text. This proposed has been successfully applied to many question answering tasks [3,14,26].

Later [8] propose the Dynamic Memory Network (DMNN) which introduces the attention mechanism into the memory network model. When retrieving memories, the location of the next related sentence is predicted according to the related sentences identified in the previous iterations. Using the attention mechanism, they obtain further improvements. Some other work [21,1] introduce other variants of MNN by introducing additional memory network modules. These work focus on storing sentence vectors for later retrieval with no exceptions. Most of them have been tested on the toy dataset bAbI [27] and are reported to have achieved satisfying results. When further tested on some practical tasks, these models also show the ability to produce results as good as existing state-of-the-art systems or even better results.

Compared with LSTM, memory networks store sentence vectors as memories and have the superiority of processing information from a large scale. Experiment results they reported on a series of tasks are concrete proofs. But there is also a problem with the memory networks as we have stated. Taking sentence vectors as input means that it is difficult to further analyze and take advantages of relations between smaller text units, such as entities. For example, when an entity e_a of sentence A interacts with another entity e_b of sentence B, we have to take the whole sentences A and B into consideration rather than just focus on e_a and e_b . This inevitably brings about noise and damages the comprehension of text. The failure of obtaining fine-grained information prevents further improvements.

4 Experiments

To verify the effectiveness of the proposed model, we conduct experiments on several datasets, including a toy QA data set bAbI, the large movie review dataset for sentiment classification and the Machine Comprehension dataset (MC Test).

4.1 bAbI

bAbI [27] is a toy data set developed for comprehension-based question answering. The example shown in Table 1 is extracted from the bAbI dataset. It contains 20 topics, each of which contains short stories, simple questions with regard to the stories and answers. The data is generated with a simulation which behaves like a classic text adventure game. According to some pre-set rules, stories are generated in a controlled context.

Previous work reports extremely satisfying results using memory networks for most topics (around 90% for most of them). However, we notice an interesting thing that all

of them with no exception fail on the problem of path-finding which is to predict a simple path like "north, west" given the locations of several subjects. Another one is the positional reasoning. The Memory Network [27] reports accuracies of 36% and 65% for the two topics. The Dynamic Memory Network [8] reports accuracies of 35% and 60%. The proposed model (Entity-MNN) reports accuracies of 53% and 67% respectively. It is still far from satisfying but the improvements on the two tasks indicates the superiority of the entity-based memory network.

Results on the toy dataset is not as convincing as that on practical tasks. Given how the bAbI data is generated, it is easy to achieve a 100% accuracy if we do simple reverse engineering to identify the entities and rules. The good results of memory networks, including our model, can not be solely attributed to their ability of comprehension. It may be partly due to their ability of inducting the entities and rules from text.

4.2 Large Movie Review Dataset

We further tested our model on the Large Movie Review Dataset [15], which is a collection of 50,000 reviews from IMDB, about 30 reviews per movie. Each review is assigned a score from 1 (very negative) to 10 (very positive). The ratio of positive samples to negative samples is 50:50. Following the previous work [15], we only consider polarized samples with scores no greater than 4 or no smaller than 7.

For each review, we present it as a short story and then add a question "what is the opinion?". The answer is either "negative" or "positive". In this way we turn this task into a question answering problem. Note that although here the answer to a question is either "negative" or "positive", we do not put any constraints on the output. It is treated in the same way as open domain question answering and the system is expected to learn to predict the output by itself.

We do not use the full dataset as the training takes a long time. We randomly select 10K samples (5K negative samples + 5K positive samples) for training and another 10K for test. We obtain an accuracy of 97.2% on the subset which is higher than the 89% reported by [15], 93.4% by [6] and 95% by [7].

Sys.	Acc.%
Maas' 11 [15]	89
Johnson' 14 [6]	93.4
Johnson' 15 [7]	95
EntityMNN	97.2

Table 2. Results on Large Movie Dataset

4.3 Machine Comprehension Test

The machine comprehension test (MCTest) dataset [18] has 500 stories and 2000 questions (MC500). All of them are multiple choice reading comprehension questions. The stories are for children so limited world knowledge is required. An additional smaller dataset with 160 stories and 640 questions (MC160) is also included in the MCTest data and used in our work.

Since the proposed model does not consider the form of multiple choice questions, we need to convert MCTest data into suitable formats firstly. When answering a multiple choice question, one is provided with several alternatives of which at least one is correct. These alternatives can be regarded as information known.

For a question, we replace the “Wh-” words using each alternative and Each alternative is turned to a new declarative sentences. These generated declarative sentences are generally understandable though may not be grammatically correct. Then we use the proposed system to decide whether the generated sentences are correct or wrong. However, we do not distinguish between questions with only one answer and those with more than one answers as these newly generated sentences are treated separately. In other words, all questions are treated as having multiple answers.

The MCTest contains only hundreds of stories and is usually used for test only as statistical models normally require a large amount of training data. However, we still obtain satisfying results using this dataset. Table 3 demonstrates the effectiveness of the entity-based model on the MCTest dataset. We outperform the previous state-of-the-art [25,19] on both MC160 and MC500. Our model does not employ rich semantic features as others do, and hence is easy to be migrated to languages aside from English.

Sys.	Acc.(%) MC160			Acc.(%) MC500		
	Single	Multiple	Average	Single	Multiple	Average
Richardson’13 [18]	76.8	62.5	69.2	68.0	59.5	63.3
Wang’15 [25]	84.2	67.9	75.3	72.1	67.9	69.9
Sachan’16 [19]	-	-	-	72.0	68.9	70.3
EntityMNN	Average=76.1			Average=76.6		

Table 3. Results on Machine Comprehension Test

4.4 Analysis

The proposed model is designed based on the assumption that entities are the core of text. By updating the states of entities, information carried by text is encoded into entities. Thus all questions which are related to the text can be answered based on entities solely. Entities enable us to break a sentence into smaller text units and analyze text from a smaller scale. Therefore the entity-based model is more sophisticated and powerful than those based on sentences as has been proven in our experiments.

A shortcoming with such a model is that, it cannot handle text that contains very few entities. Also hidden entities are not considered. As we know, pro-drop languages, like Japanese and Chinese, tend to omit certain classes of pronouns when they are inferable. The proposed model will encounter problems when dealing with such text.

5 Conclusion

This work presents the entity-based memory network model for text comprehension. All the information conveyed by text is encoded into the states of its containing entities and questions regarded to the text are answered using these entities. Experiments on several tasks have proven the effectiveness of the proposed model.

The proposed model is based on the assumption that entities can express all the information of text. In future research, we will further explore its ability by considering more components in text. not merely entities.

References

1. Bordes, A., Usunier, N., Chopra, S., Weston, J.: Large-scale simple question answering with memory networks. arXiv preprint arXiv:1506.02075 (2015)
2. Chung, J., Gulcehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. arXiv preprint arXiv:1412.3555 (2014)
3. Dodge, J., Gane, A., Zhang, X., Bordes, A., Chopra, S., Miller, A., Szlam, A., Weston, J.: Evaluating prerequisite qualities for learning end-to-end dialog systems. arXiv preprint arXiv:1511.06931 (2015)
4. Dyer, C., Ballesteros, M., Ling, W., Matthews, A., Smith, N.A.: Transition-based dependency parsing with stack long short-term memory. arXiv preprint arXiv:1505.08075 (2015)
5. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural computation* 9(8), 1735–1780 (1997)
6. Johnson, R., Zhang, T.: Effective use of word order for text categorization with convolutional neural networks. arXiv preprint arXiv:1412.1058 (2014)
7. Johnson, R., Zhang, T.: Semi-supervised convolutional neural networks for text categorization via region embedding. In: NIPS. pp. 919–927 (2015)
8. Kumar, A., Irsoy, O., Su, J., Bradbury, J., English, R., Pierce, B., Ondruska, P., Gulrajani, I., Socher, R.: Ask me anything: Dynamic memory networks for natural language processing. arXiv preprint arXiv:1506.07285 (2015)
9. Le, Q.V., Mikolov, T.: Distributed representations of sentences and documents. In: ICML. vol. 14, pp. 1188–1196 (2014)
10. Lehnert, W.G.: *The process of question answering: A computer simulation of cognition*. Lawrence Erlbaum Associates (1978)
11. Li, J., Galley, M., Brockett, C., Gao, J., Dolan, B.: A diversity-promoting objective function for neural conversation models. arXiv preprint arXiv:1510.03055 (2015)
12. Li, J., Galley, M., Brockett, C., Gao, J., Dolan, B.: A persona-based neural conversation model. arXiv preprint arXiv:1603.06155 (2016)
13. Li, J., Luong, M.T., Jurafsky, D.: A hierarchical neural autoencoder for paragraphs and documents. arXiv preprint arXiv:1506.01057 (2015)
14. Li, J., Miller, A.H., Chopra, S., Ranzato, M., Weston, J.: Dialogue learning with human-in-the-loop. arXiv preprint arXiv:1611.09823 (2016)
15. Maas, A.L., Daly, R.E., Pham, P.T., Huang, D., Ng, A.Y., Potts, C.: Learning word vectors for sentiment analysis. In: ACL-HLT. pp. 142–150 (2011)
16. Mikolov, T., Karafiát, M., Burget, L., Cernocký, J., Khudanpur, S.: Recurrent neural network based language model. In: Interspeech. vol. 2, p. 3 (2010)
17. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: EMNLP. vol. 14, pp. 1532–43 (2014)
18. Richardson, M., Burges, C.J., Renshaw, E.: Mctest: A challenge dataset for the open-domain machine comprehension of text. In: EMNLP. vol. 3, p. 4 (2013)
19. Sachan, M., Xing, E.P.: Machine comprehension using rich semantic representations. In: ACL (2016)
20. Socher, R., Lin, C.C., Manning, C., Ng, A.Y.: Parsing natural scenes and natural language with recursive neural networks. In: ICML. pp. 129–136 (2011)
21. Sukhbaatar, S., Weston, J., Fergus, R., et al.: End-to-end memory networks. In: NIPS. pp. 2440–2448 (2015)
22. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In: NIPS. pp. 3104–3112 (2014)
23. Vinyals, O., Kaiser, Ł., Koo, T., Petrov, S., Sutskever, I., Hinton, G.: Grammar as a foreign language. In: Advances in Neural Information Processing Systems. pp. 2773–2781 (2015)

24. Vinyals, O., Le, Q.: A neural conversational model. arXiv preprint arXiv:1506.05869 (2015)
25. Wang, H., McAllester, M.B.K.G.D.: Machine comprehension with syntax, frames, and semantics. ACL, Volume 2: Short Papers p. 700 (2015)
26. Weston, J.: Dialog-based language learning. arXiv preprint arXiv:1604.06045 (2016)
27. Weston, J., Bordes, A., Chopra, S., Rush, A.M., van Merriënboer, B., Joulin, A., Mikolov, T.: Towards ai-complete question answering: A set of prerequisite toy tasks. arXiv preprint arXiv:1502.05698 (2015)
28. Weston, J., Chopra, S., Bordes, A.: Memory networks. arXiv preprint arXiv:1410.3916 (2014)