

A CONTEXT-AWARE ATTENTION NETWORK FOR INTERACTIVE QUESTION ANSWERING

Huayu Li^{1*}, Martin Renqiang Min², Yong Ge³, Asim Kadav²

Department of Computer Science, UNC Charlotte¹

Machine Learning Group, NEC Laboratories America²

Management Information Systems, University of Arizona³

hli38@uncc.edu, renqiang@nec-labs.com,

yongge@email.arizona.edu, asim@nec-labs.com

ABSTRACT

We develop a new model for Interactive Question Answering (IQA), using Gated-Recurrent-Unit recurrent networks (GRUs) as encoders for statements and questions, and another GRU as a decoder for outputs. Distinct from previous work, our approach employs context-dependent word-level attention for more accurate statement representations and question-guided sentence-level attention for better context modeling. Employing these mechanisms, our model accurately understands when it can output an answer or when it requires generating a supplementary question for additional input. When available, user’s feedback is encoded and directly applied to update sentence-level attention to infer the answer. Extensive experiments on QA and IQA datasets demonstrate quantitatively the effectiveness of our model with significant improvement over conventional QA models.

1 INTRODUCTION

The ultimate goal of Question Answering (QA) research is to build intelligent systems capable of naturally communicating with humans, which poses a major challenge for natural language processing and machine learning. Inspired by recent success of sequence-to-sequence models with an encoder-decoder framework (Sutskever et al., 2014; Cho et al., 2014), researchers have attempted to apply variants of such models with explicit memory and attention to QA tasks, aiming to move a step further from machine learning to machine reasoning (Sainbayar et al., 2015; Kumar et al., 2016; Xiong et al., 2016). Similarly, all these models employ encoders to map statements and questions to fixed-length feature vectors, and a decoder to generate outputs. Empowered by the adoption of memory and attention, they have achieved remarkable success on several challenging datasets, including the recently acclaimed Facebook bAbI dataset.

However, previous models suffer from the following important limitations. First, they fail to model context-dependent meaning of words. Different words may have different meanings in different contexts, which increases the difficulty of extracting the essential semantic logic flow of each sentence in different paragraphs. Second, many existing models only work in ideal QA settings and fail to address the uncertain situations under which models require additional user input to gather complete information to answer a given question. As shown in Table 1, the example on the left is an ideal QA problem. We can clearly understand what the question is and then locate the relevant sentences to generate the answer. However, it is hard to answer the question in the right example, because there are two types of bedrooms mentioned in the story and we do not know which bedroom the user refers to. These scenarios with incomplete information naturally appear in human conversations, and thus, effectively handling them is a key capability of intelligent QA models.

To address the challenges presented above, we propose a Context-aware Attention Network (CAN) to learn fine-grained representations for input sentences, and develop a mechanism to interact with the user for comprehensively understanding a given question. Specifically, we employ two-level attention applied at word level and sentence level to compute representations of all input sentences.

*Most of this work was done when the author was an intern at NEC Labs America.

<p>The office is north of the kitchen. The garden is south of the kitchen. Q: What is north of the kitchen? A: Office</p>	<p>The master bedroom is east of the garden. The guest bedroom is east of the office. Q: What is the bedroom east of? A: Unknown</p>
--	---

Table 1: Two examples of QA problems. Left is an ideal QA example, where the question is very clear. Right is an example with incomplete information, where the question is ambiguous and it is difficult to provide an answer only using the input statements.

The context information extracted from the input story is allowed to influence the attention over each word, and governs the word semantic meaning contributing to a sentence representation. In addition, an interactive mechanism is activated to generate a supplementary question for the user when the model feels that it does not have complete information to answer a given question. User’s feedback is then encoded and exploited to attend over all input sentences to infer the answer. Our proposed model CAN can be viewed as an encoder-decoder approach augmented with two-level attention and an interactive mechanism, rendering our model self-adaptive, as illustrated in Figure 1.

Our contributions in this paper are as follows: (i) We develop a new encoder-decoder model called CAN for question answering with two-level attention. Due to the new attention mechanism, our model avoids the necessity of multiple-hop attention, required by previous QA models, and knows when it can readily output an answer and when it needs additional information. (ii) We augment the encoder-decoder framework for QA with an interactive mechanism for handling user’s feedback, which immediately changes sentence-level attention to infer the final answer without additional model training. (iii) We introduce a new dataset based on the bAbI dataset, namely ibAbI, for IQA tasks. (iv) Extensive experiments show that our approach outperforms state-of-the-art models on both QA and IQA datasets. Specifically, our approach achieves 40% improvement over traditional QA models (e.g., MemN2N and DMN+) on IQA datasets.

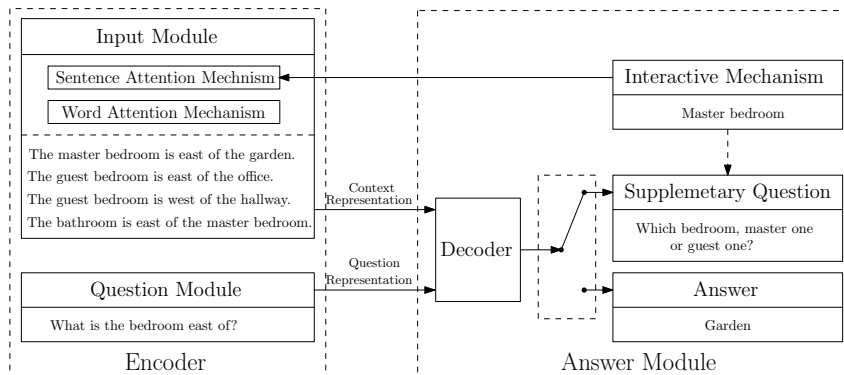


Figure 1: An example of QA problem using CAN.

2 RELATED WORK

Recent work on QA has been heavily influenced by research on various models with attention and/or memory. Most of these models employ an encoder-decoder framework, and have been successfully applied to image classification (Seo et al., 2016), image captioning (Xu et al., 2015; Mnih et al., 2014), machine translation (Cho et al., 2014; Bahdanau et al., 2015; Luong et al., 2015), document classification (Yang et al., 2016), and textual/visual QA (Sainbayar et al., 2015; Yang et al., 2015; Lu et al., 2016; Kumar et al., 2016; Xiong et al., 2016). For textual QA in the form of statements-question-answer triplets, Sainbayar et al. (2015) utilizes an external memory module. It maps each input sentence to an input representation space regarded as a memory component. The output representation is calculated by summarizing over input representations with different attention weights. This single-layer memory can be extended to multi-layer memory by reasoning the content and the question multiple times. Instead of simply stacking the memory layers, Kumar et al. (2016) have introduced a dynamic memory network (DMN) to update the memory vectors through a modified GRU, in which the gate weight is trained in a supervised fashion. To improve DMN by train-

ing without supervision, Xiong et al. (2016) encode input sentences with a bidirectional GRU and then utilize an attention-based GRU to summarize these input sentences. Neural Turing Machine (NTM) (Graves et al., 2014), a model with content and location-based memory addressing mechanisms, has also been used for QA tasks recently. There is other recent work about QA using external resources (Wu et al., 2015; Fader et al., 2014; Savenkov & Emory, 2016; Hermann et al., 2015; Golub & He, 2016), and exploring dialog tasks (Weston, 2016; Bordes & Weston, 2016).

Our model in this paper also addresses textual QA in the form of statements-question-answer triplets, but it differs from prior work in two aspects. First, in our attention network, the word attention are context-dependent for generating accurate sentence representations and the sentence attention are question-guided for generating context representation. Second, this new attention mechanism helps our model understand when it can readily output an answer and when it can generate a supplementary question for activating the user interaction. Incorporating user’s feedback does not require additional model training and this property makes our model highly self-adaptive.

3 GATED RECURRENT UNIT NETWORKS

Gated Recurrent Unit (GRU) (Cho et al., 2014) is the basic building block of our model for IQA. GRU has been widely adopted for many NLP tasks, such as machine translation (Bahdanau et al., 2015) and language modeling (Zaremba et al., 2014). GRU improves Long Short-term Memory (LSTM) (Hochreiter & Schmidhuber, 1997) by removing the cell component and making each hidden state adaptively capture the dependencies over different time scales using reset and update gates. For each time step t with input \mathbf{x}^t and previous hidden state \mathbf{h}^{t-1} , we compute the updated hidden state $\mathbf{h}^t = GRU(\mathbf{h}^{t-1}, \mathbf{x}^t)$ by,

$$\begin{aligned} \mathbf{r}^t &= \sigma(\mathbf{U}_r \mathbf{x}^t + \mathbf{W}_r \mathbf{h}^{t-1} + \mathbf{b}_r), & \mathbf{z}^t &= \sigma(\mathbf{U}_z \mathbf{x}^t + \mathbf{W}_z \mathbf{h}^{t-1} + \mathbf{b}_z), \\ \tilde{\mathbf{h}}^t &= \tanh(\mathbf{U}_h \mathbf{x}^t + \mathbf{W}_h (\mathbf{r}^t \odot \mathbf{h}^{t-1}) + \mathbf{b}_h), & \mathbf{h}^t &= \mathbf{z}^t \odot \mathbf{h}^{t-1} + (\mathbf{1} - \mathbf{z}^t) \odot \tilde{\mathbf{h}}^t, \end{aligned}$$

where σ is the sigmoid activation function, \odot is an element-wise product, $\mathbf{U}_r, \mathbf{U}_z, \mathbf{U}_h \in \mathbb{R}^{K \times D}$, $\mathbf{W}_r, \mathbf{W}_z, \mathbf{W}_h \in \mathbb{R}^{K \times K}$, $\mathbf{b}_r, \mathbf{b}_z, \mathbf{b}_h \in \mathbb{R}^{K \times 1}$, K is the hidden size and D is the input size.

4 CONTEXT-AWARE ATTENTION NETWORK

In this section, we first illustrate the proposed model CAN (§ 4.1), including the question module (§ 4.2), the input module (§ 4.3) and the answer module (§ 4.4). We then describe each of these modules in detail. Finally, we elaborate the training procedure of CAN (§ 4.5).

4.1 FRAMEWORK

Given a story represented by N input sentences (or statements), i.e., (l_1, \dots, l_N) , and a question q , our goal is to generate an answer a . Each sentence l_t includes a sequence of N_t words, denoted as $(w_1^t, \dots, w_{N_t}^t)$, and a question with N_q words is represented as $(w_1^q, \dots, w_{N_q}^q)$. Let V denote the size of dictionary, including the words from each l_t , q and a , and end-of-sentence (EOS) symbols.

The whole framework of our model is shown in Figure 2, consisting of the following three key parts:

- **Question Module:** The question module encodes the target question into a vector representation.
- **Input Module:** The input module encodes a set of input sentences into a vector representation.
- **Answer Module:** The answer module generates an answer based on the outputs of question and input modules. Unlike traditional QA models, it has two choices, either to output an answer immediately or to interact with the user for further information. Hence, if the model lacks sufficient evidence for answer prediction based on the existing knowledge at current timestamp, an interactive mechanism is enabled. Specifically, the model generates a supplementary question, and the user needs to provide a feedback, which is utilized to estimate an answer.

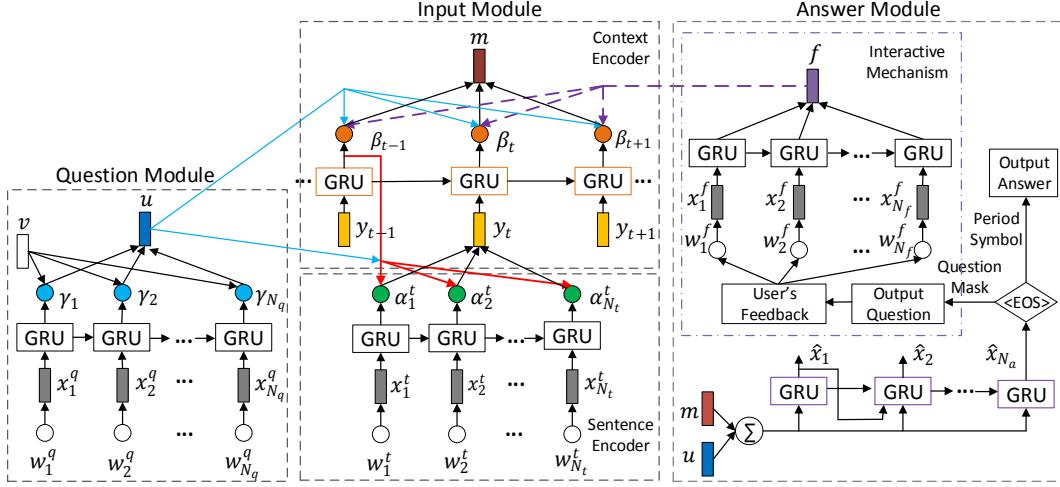


Figure 2: The illustration of the proposed model, consisting of a question module, an input module and an answer module. The question module maps the question sentence into a sentence level space. The input module generates a context representation based on input sentences. The answer module has a binary choice, either to generate an answer immediately or to take an interactive mechanism.

4.2 QUESTION MODULE

Suppose a question is a sequence of N_q words, we encode each word w_j into a K_w -dimensional vector space \mathbf{x}_j^q using an embedding matrix $\mathbf{W}_w \in \mathbb{R}^{K_w \times V}$, i.e., $\mathbf{x}_j^q = \mathbf{W}_w[w_j]$, where $[w_j]$ is a one-hot vector associated with word w_j . The sequence order within a sentence significantly affects each word's semantic meaning due to its dependence on the previous words. Thus, a GRU is employed by taking each word vector \mathbf{x}_j^q as input and updating the hidden state $\mathbf{g}_j^q \in \mathbb{R}^{K_h \times 1}$ as:

$$\mathbf{g}_j^q = GRU_w(\mathbf{g}_{j-1}^q, \mathbf{x}_j^q), \quad (1)$$

where the subscript of GRU is used to distinguish other GRUs used in the following sections. The hidden state \mathbf{g}_j^q can be regarded as the annotation vector of word w_j by incorporating the word order information. We also explore a variety of encoding schema, such as LSTM and RNN. However, LSTM is prone to overfitting due to large number of parameters, and RNN has a poor performance because of exploding and vanishing gradients (Bengio et al., 1994).

In addition, each word contributes differently to the representation of a question. For example, in a question 'Where is the football?', 'where' and 'football' play a critical role in summarizing this sentence. Therefore, an attention mechanism is introduced to generate a question representation by focusing on the important words for their semantic meaning. A positive weight γ_j is placed on each word to indicate the relative importance of contribution to the representation of the question. Specifically, this weight is measured as the similarity of corresponding word annotation vector \mathbf{g}_j^q and a word level latent vector $\mathbf{v} \in \mathbb{R}^{K_h \times 1}$ for question which is jointly learned during the training process. The question representation $\mathbf{u} \in \mathbb{R}^{K_c \times 1}$ is then generated by a weighted summation of the word annotation vectors and corresponding important weights, where we also use one-layer MLP to transfer it from sentence-level space into context-level space,

$$\gamma_j = \text{softmax}(\mathbf{v}^T \mathbf{g}_j^q), \quad (2) \quad \mathbf{u} = \mathbf{W}_{ch} \sum_{j=1}^{N_q} \gamma_j \mathbf{g}_j^q + \mathbf{b}_c^{(q)}, \quad (3)$$

where softmax is defined as $\text{softmax}(x_i) = \frac{\exp(x_i)}{\sum_{j'} \exp(x_{j'})}$, $\mathbf{W}_{ch} \in \mathbb{R}^{K_c \times K_h}$, and $\mathbf{b}_c^{(q)} \in \mathbb{R}^{K_c \times 1}$.

4.3 INPUT MODULE

Input module aims at generating a representation for input sentences, including a sentence encoder and a context encoder. Sentence encoder computes a sentence representation, and context encoder calculates a representation of input sentences, both of which are introduced in the following sections.

4.3.1 SENTENCE ENCODER

For each input sentence l_t , containing a sequence of N_t words (w_1, \dots, w_{N_t}) , similar to the question module, each word w_i is embedded into word space $\mathbf{x}_i^t \in \mathbb{R}^{K_w \times 1}$ with the embedding matrix \mathbf{W}_w , and a recurrent neural network is used to capture the context information from the words which have already been generated in the same sentence. Let $\mathbf{h}_i^t \in \mathbb{R}^{K_h \times 1}$ denote the hidden state which can be interpreted as the word annotation in the input space. A GRU retrieves each word annotation by taking word vector as input and relying on previous hidden state,

$$\mathbf{h}_i^t = GRU_w(\mathbf{h}_{i-1}^t, \mathbf{x}_i^t). \quad (4)$$

In Eq.(4), each word annotation vector takes the sequence order into consideration to learn its semantic meaning based on previous information within a sentence through a recurrent neural network. A question answering system is usually given multiple input sentences which often form a story together. A single word has different meaning in the different stories. Learning a single sentence context at which a word is located is insufficient to understand the meaning of this word, in particular when the sentence is placed in a story context. In other words, only modeling a sequence of words prior to a word within a sentence may lose some important information which results in the failure of the generation of sentence representation. Hence, we take the whole context into account as well to appropriately characterize each word and well understand this sentence's meaning. Suppose $\mathbf{s}_{t-1} \in \mathbb{R}^{K_c \times 1}$ is the annotation vector of previous sentence l_{t-1} , which will be introduced in the next section. To incorporate context information generated by previous sentences, we feed word annotation \mathbf{h}_i^t and previous sentence annotation \mathbf{s}_{t-1} through a two-layer MLP, where a context-aware word vector $\mathbf{e}_i^t \in \mathbb{R}^{K_c \times 1}$ is obtained as follows:

$$\mathbf{e}_i^t = \sigma(\mathbf{W}_{ee} \tanh(\mathbf{W}_{es} \mathbf{s}_{t-1} + \mathbf{W}_{eh} \mathbf{h}_i^t + \mathbf{b}_e^{(1)}) + \mathbf{b}_e^{(2)}), \quad (5)$$

where $\mathbf{W}_{ee}, \mathbf{W}_{es} \in \mathbb{R}^{K_c \times K_c}$ and $\mathbf{W}_{eh} \in \mathbb{R}^{K_c \times K_h}$ are weight matrices, and $\mathbf{b}_e^{(1)}, \mathbf{b}_e^{(2)} \in \mathbb{R}^{K_c \times 1}$ are the bias terms. It is worth noting that \mathbf{s}_{t-1} is dependent on its previous sentence. Recursively, this sentence relies on its previous one as well. Hence, our model is able to encode the previous context. In addition, the sentence representation will focus on those words which are able to address the question. Inspired by this intuition, another word level attention mechanism is introduced to attend informative words about the question for generating a sentence's representation. As the question representation is utilized to guide the word attention, a positive weight α_i^t associated with each word is computed as the similarity of the question vector \mathbf{u} and the corresponding context-aware word vector \mathbf{e}_i^t . Then the sentence representation $\mathbf{y}_t \in \mathbb{R}^{K_h \times 1}$ is generated by aggregating the word annotation vectors with different weights,

$$\alpha_i^t = \text{softmax}(\mathbf{u}^T \mathbf{e}_i^t), \quad \mathbf{y}_t = \sum_{i=1}^{N_t} \alpha_i^t \mathbf{h}_i^t. \quad (6)$$

4.3.2 CONTEXT ENCODER

Suppose a story is comprised of a sequence of sentences, i.e., (l_1, \dots, l_N) , each of which is encoded as a K_h -dimensional vector \mathbf{y}_t through a sentence encoder. As input sentences have a sequence order, simply using their sentence vectors for context generation cannot effectively capture the entire context of the sequence of sentences. To address this issue, a sentence annotation vector is introduced to capture the previous context and this sentence's own meaning using a GRU. Given the sentence vector \mathbf{y}_t and the state \mathbf{s}_{t-1} of previous sentence, we get annotation vector $\mathbf{s}_t \in \mathbb{R}^{K_c \times 1}$ as:

$$\mathbf{s}_t = GRU_s(\mathbf{s}_{t-1}, \mathbf{y}_t). \quad (7)$$

A GRU can learn a sentence's meaning based on previous context information. However, just relying on GRU at sentence level using simple word embedding vectors makes it difficult to learn the precise semantic meaning for each word in the story. Hence, we introduce a context-aware attention mechanism shown in Eq.(5) to properly encode each word for the generation of sentence representation, which guarantees that each word is reasoned under the specific context.

Once the sentence annotation vectors $(\mathbf{s}_1, \dots, \mathbf{s}_N)$ are obtained as described above, a sentence level attention mechanism is enabled to emphasize those sentences that are highly relevant to the question. We can estimate the attention weight β_t with the similarity of the question and the

corresponding sentence. Hence, the context representation \mathbf{m} is retrieved by summing over all sentence representations associated with corresponding attention weights, and given by:

$$\beta_t = \text{softmax}(\mathbf{u}^T \mathbf{s}_t), \quad (8) \quad \mathbf{m} = \sum_{t=1}^N \beta_t \mathbf{s}_t. \quad (9)$$

Similar to bidirectional RNN, our model can be extended to use another sentence-level GRU that moves backward through time beginning from the end of the sequence.

4.4 ANSWER MODULE

The answer module utilizes a decoder to generate an answer, where it has two output cases according to the understanding ability of both the question and the context. One is to generate the answer immediately after receiving the context and question information. Another one is to generate a supplementary question and then use the user’s feedback to predict the answer. This process is taken by an interactive mechanism.

4.4.1 ANSWER GENERATION

Given the question representation \mathbf{u} and the context representation \mathbf{m} , another GRU is used as the decoder to generate a sentence as the answer. To fuse \mathbf{u} and \mathbf{m} together, we sum these vectors rather than concatenating them to reduce the total number of parameters. Suppose $\hat{\mathbf{x}}_{k-1} \in \mathbb{R}^{K_w \times 1}$ is the predicted word vector in last step, GRU updates the hidden state $\mathbf{z}_k \in \mathbb{R}^{K_o \times 1}$ as follows,

$$\hat{\mathbf{x}}_k \stackrel{\mathbf{W}_w}{=} \text{softmax}(\mathbf{W}_{od} \mathbf{z}_k + \mathbf{b}_o), \quad \mathbf{z}_k = \text{GRU}_d(\mathbf{z}_{k-1}, [\mathbf{m} + \mathbf{u}; \hat{\mathbf{x}}_{k-1}]) \quad (10)$$

where $\stackrel{\mathbf{W}_w}{=}$ denotes the predicted word vector through the embedding matrix \mathbf{W}_w . Note that we require that each sentence ends with a special EOS symbol, including question mask and period symbol, which enables the model to define a distribution over sentences of all possible lengths.

Output Choices. In practice, the system is not always able to answer question immediately based on its current knowledge due to the lack of some crucial information bridging the gap between question and context knowledge, i.e., incomplete issue. Therefore, we allow the decoder to make a binary choice, either to generate an answer immediately, or to enable an interactive mechanism. Specifically, if the model has sufficiently strong evidence for a successful answer prediction based on the well-learned context representation and question representation, the decoder will directly output the answer. Otherwise, the system generates a supplementary question for user, where an example is shown in Table 2. At this time, this user needs to offer a feedback which is then encoded to update the sentence-level attentions for answer generation. This procedure is our interactive mechanism.

Problem	The master bedroom is east of the garden. The guest bedroom is east of the office. Target Question: What is the bedroom east of?
Interactive Mechanism	System: Which bedroom, master one or guest one? (SQ) User: Master bedroom (User’s Feedback) System: Garden (Predicted Answer)

Table 2: An example of interactive mechanism. “SQ” denotes supplementary question.

The sentence generated by the decoder ends with a special symbol, either a question mask or a period symbol. Hence, this special symbol is utilized to make a decision. In other words, if EOS symbol is a question mask, the generated sentence is regarded as a supplementary question and an interactive mechanism is enabled; otherwise the generated sentence is the estimated answer and the prediction task is done. In the next section, we will introduce the details of interactive mechanism.

4.4.2 INTERACTIVE MECHANISM

The interactive process is summarized as follows: 1) The decoder generates a supplementary question; 2) The user provides a feedback; 3) The feedback is used for answer prediction for the target question. Suppose the feedback contains a sequence of words, denoted as $(w_1^f, \dots, w_{N_f}^f)$. Similar to the input module, each word w_d^f is embedded to a vector \mathbf{x}_d^f through an embedding matrix

\mathbf{W}_w . Then the corresponding annotation vector $\mathbf{g}_d^f \in \mathbb{R}^{K_h \times 1}$ is retrieved via a GRU by taking the embedding vector as input, and shown as follows:

$$\mathbf{g}_d^f = GRU_w(\mathbf{g}_{d-1}^f, \mathbf{x}_d^f). \quad (11)$$

Based on the annotation vectors, a representation $\mathbf{f} \in \mathbb{R}^{K_h \times 1}$ can be obtained by a simple attention mechanism where each word is considered to contribute equally, and given by:

$$\mathbf{f} = \frac{1}{N_f} \sum_{d=1}^{N_f} \mathbf{g}_d^f. \quad (12)$$

Our goal is to utilize the feedback representation \mathbf{f} to generate an answer for the target question. The provided feedback improves the ability to answer the question by distinguishing the relevance of each input sentence to the question. In other words, the similarity of specific input sentences in the provided feedback make these sentences more likely to address the question. Hence, we refine the attention weight of each sentence shown in Eq.(9) after receiving the user’s feedback, given by,

$$\mathbf{r} = \tanh(\mathbf{W}_{rf}\mathbf{f} + \mathbf{b}_r^{(f)}), \quad (13) \quad \beta_t = \text{softmax}(\mathbf{u}^T \mathbf{s}_t + \mathbf{r}^T \mathbf{s}_t) \quad (14)$$

where $\mathbf{W}_{rf} \in \mathbb{R}^{K_c \times K_h}$ and $\mathbf{b}_r^{(f)} \in \mathbb{R}^{K_c \times 1}$ are the weight matrix and bias vector, respectively. Eq.(13) is an one-layer neural network to transfer feedback representation to context space. After obtaining the newly learned attention weights, we update the context representation using the soft-attention operation shown in Eq.(9). This updated context representation and question representation will be used as the input for decoder to generate an answer. Note that for simplifying the problem, we allow the decoder to only generate at most one supplementary question. In addition, one advantage of using the user’s feedback to update the attention weights of input sentences is that we do not need to re-train the encoder again once a feedback is entering the system.

4.5 TRAINING PROCEDURE

During training, all modules share an embedding matrix. There are three different GRUs employed for sentence encoding, context encoding and answer/supplementary question decoding. In other words, the same GRU is used to encode the question, input sentences and the user’s feedback. The second one is applied to generate context representation and the third one is used as decoder. Training can be treated as a supervised classification problem to minimize the cross-entropy error of the answer sequence and the supplementary question sequence.

5 EXPERIMENTS

In this section, we evaluate our approach with many baseline methods based on various datasets.

5.1 EXPERIMENTAL SETUP

Datasets. In this paper, we use two types of datasets to evaluate the performance of our approach. One is traditional QA dataset, where we use Facebook bAbI English 10k dataset (Weston et al., 2015). It contains 20 different types of tasks with emphasis on different forms of reasoning and induction. The second is the newly designed IQA dataset, where we extend bAbI to add interactive QA and denote it as ibAbI. Overall, we generate three ibAbI datasets based on task 1 (single supporting fact), task 4 (two argument relations), and task 7 (counting). Specifically, the former two datasets focus on solving ambiguous actors/objects problem, and the latter one is to ask further information that assists answer prediction. Table 3 shows three examples for our three ibAbI tasks.

In addition, we also mix IQA data and corresponding QA data together with different IQA ratios, where the IQA ratio is ranging from 0.3 to 1 (with step as 0.1) and denoted as R_{IQA} . For example, in task 1, we randomly pick $R_{IQA} \times 100$ percent data from ibAbI task 1, and then randomly select the remaining data from bAbI task 1. $R_{IQA} = 1$ indicates that the whole dataset only consists of IQA problems; otherwise (i.e., ranging from 0.3 to 0.9) it consists of both types of QA problems. Overall, we have three tasks for ibAbI dataset, and eight sub-datasets for each task. In the experiments, 10k examples are used as training and another 1k examples are used as testing.

<p>IQA task 1: John journeyed to the garden. Daniel moved to the kitchen.</p> <p>Q: Where is he? SQ: Who is he? FB: Daniel A: Kitchen</p>	<p>IQA task 4: The master bedroom is east of the garden. The guest bedroom is east of the office. The guest bedroom is west of the hallway. The bathroom is east of the master bedroom.</p> <p>Q: What is the bedroom east of? SQ: Which bedroom, master one or guest one? FB: Master bedroom A: Garden</p>	<p>IQA task 7: John grabbed the bread. John grabbed the milk. John grabbed the apple. Sandra went to the bedroom.</p> <p>Q: How many special objects is John holding? SQ: What objects are you referring to? FB: Milk, bread A: Two</p>
--	--	--

Table 3: Examples of three different tasks on the generated ibAbI datasets. “Q” indicates the target question. “SQ” is the supplementary question. “FB” refers to user’s feedback. “A” is the answer.

Experiment Settings. We train our models using the Adam optimizer (Kingma & Ba, 2014). Xavier initialization is used for all parameters except for word embeddings, which utilize random uniform initialization ranging from $-\sqrt{3}$ to $\sqrt{3}$. The learning rate is set as 0.001. The grid search method is utilized to find optimal parameters, such as batch size and hidden size.

5.2 BASELINE METHODS

To demonstrate the effectiveness of our approach CAN, we compare it with the following models:

- **DMN+:** Xiong et al. (2016) improve Dynamic Memory Networks (Kumar et al., 2016) by using stronger input and memory modules, where a bidirectional GRU is adopted to generate representations for statements and a neural network is used to update episodic memory multiple times.
- **MemN2N:** This is an extension of Memory Network with weak supervision as proposed in Sainbayar et al. (2015). Here, an external memory module is used to encode the input statements and a recurrent attention mechanism is used to read the memory for answer prediction.
- **EncDec:** We extend the encoder-decoder framework (Cho et al., 2014) to solve QA tasks as a baseline method. Specifically, EncDec uses a GRU to encode statements and questions, the end of hidden states is used as context representation, and another GRU to generate the output.

5.3 PERFORMANCE OF QUESTION ANSWERING

In this section, we evaluate model’s ability for answer prediction based on traditional QA dataset (i.e., bAbI-10k). For this task, our model (denoted as CAN+QA) does not use the interactive mechanism. As the output answers for this dataset only contain a single word, we adopt test error rate as evaluation metric. For DMN+ and MemN2N methods, we select the best performance over bAbI dataset reported in (Xiong et al., 2016). The results of various models across 20 tasks are reported in Table 4. We summarize the main observations as follows:

- Our approach is better than all baseline methods in each individual task. For example, it reduces the error by 4% compared to DMN+ in task 17, and compared to MemN2N, it reduces 18.4% and 4.8% error in task 17 and 18 respectively. We can achieve a better result primarily because our approach can model the semantic logic flow for statements. Table 5 shows two examples in task 17 and 18, where MemN2N predicts incorrectly while CAN+QA can make correct predictions. In these two examples, the semantic logic determines the relationship between two objects mentioned in the question, such as chest and suitcase. In addition, Kumar et al. (2016) has shown that memory networks with multiple hops are better than the one with single hop. Our strong results illustrate that our approach has more accurate context modeling even without multiple hops.
- EncDec performs the worst amongst all models over all tasks. EncDec concatenates the statements and questions as a single input, resulting in the difficulty of training the GRU. For example, EncDec is not good on task 2 and 3 because these two tasks have longer inputs than other tasks.
- The results of DMN+ and MemN2N are much better than EncDec. It is not surprising that they can outperform EncDec, because they are specifically designed for question answering and do not suffer from the problem mentioned above by treating input sentences separately.
- All models perform poorly on task 16. Xiong et al. (2016) points out that MemN2N with a simple update for memory could achieve a near perfect error rate of 0.4 while a more complex method will lead to a much worse result. This shows that a sophisticated modeling method makes it

difficult to achieve a good performance in certain simple tasks with such limited data. This can be a possible reason for the poor performance of our model on this specific task as well.

In addition, different from MemN2N, we use a GRU to capture the semantic logic flow of input sentences, where the sentence-level attention can weaken the influence of unrelated sentences in a long story. Table 6 shows two examples of our results with long stories. From the attention weights, we can see our model can correctly search relevant sentences in a long story to address a question.

Task	CAN+QA	DMN+	MemN2N	EncDec
1 - Single Supporting Fact	0.0	0.0	0.0	52.0
2 - Two Supporting Facts	0.1	0.3	0.3	66.1
3 - Three Supporting Facts	0.2	1.1	2.1	71.9
4 - Two Arg. Relations	0.0	0.0	0.0	29.2
5 - Three Arg. Relations	0.4	0.5	0.8	14.3
6 - Yes/No Questions	0.0	0.0	0.1	31.0
7 - Counting	0.3	2.4	2.0	23.6
8 - Lists/Sets	0.0	0.0	0.9	28.8
9 - Simple Negation	0.0	0.0	0.3	39.1
10 - Indefinite Knowledge	0.0	0.0	0.0	45.0
11 - Basic Coreference	0.0	0.0	0.1	31.7
12 - Conjunction	0.0	0.0	0.0	35.0
13 - Compound Coref.	0.0	0.0	0.0	8.7
14 - Time Reasoning	0.0	0.2	0.1	67.2
15 - Basic Deduction	0.0	0.0	0.0	62.2
16 - Basic Induction	43.0	45.3	51.8	54.0
17 - Positional Reasoning	0.2	4.2	18.6	43.1
18 - Size Reasoning	0.5	2.1	5.3	9.0
19 - Path Finding	0.0	0.0	2.3	89.6
20 - Agents Motivations	0.0	0.0	0.0	2.3

Table 4: Performance comparison of various models in terms of test error rate (%) in QA dataset.

<p>The red square is below the triangle. The pink rectangle is to the left of the red square.</p> <p>Q: Is the triangle above the pink rectangle? A: yes</p>	<p>The box is bigger than the suitcase. The suitcase fits inside the container. The box of chocolates fits inside the container. The container fits inside the chest. The chocolate fits inside the suitcase. Q: Is the chest bigger than the suitcase? A: yes</p>
--	--

Table 5: Examples of bAbI task 17 (left) and 18 (right), where our model predicts correct answers while MemN2N makes wrong predictions.

Story	Support	Weight	Story	Support	Weight
Line 1: Mary journeyed to the office. ...		0.00	Line 1: John went back to the kitchen. ...		
Line 48: Sandra grabbed the apple there. Line 49: Sandra dropped the apple. Line 50: ...	yes yes	0.13 0.85	Line 13 : Sandra grabbed the apple there. ... Line 29: Sandra left the apple. Line 30: ...	yes yes	0.14 0.79
What is Sandra carrying? Answer: nothing Prediction: nothing			What is Sandra carrying? Answer: nothing Prediction: nothing		

Table 6: Examples of our model’s results on QA tasks. Supporting facts are shown in the datasets which our model does not use during training. “Weight” indicates the attention weight for sentence. Our model can locate correct supporting sentences for long stories.

5.4 PERFORMANCE OF INTERACTIVE QUESTION ANSWERING

In this section, we evaluate the performance of various models based on IQA dataset (as described in Section 5.1). For testing, we simulate the interactive procedure by taking the predefined feedback as user’s input for the generated supplementary question, and then generating an answer. All baseline methods do not have interactive part, so they take both statements and question as input and then estimate an answer. We compare our approach (CAN+IQA) with baseline methods in terms of test error rate shown in Table 7. From the results, we can achieve the following conclusions:

- Our method significantly outperforms all baseline methods. Specifically, we can achieve 0% test error rate in task 1 and task 4 with $R_{IQA} = 1.0$; while the best result of baseline methods can only get 40.5% test error rate. CAN+IQA benefits from more accurate context modeling, which allows it to correctly understand when to output an answer or require additional information. For those QA problems with incomplete information, it is necessary to gather the additional information from users. Randomly guessing may harm model’s performance, which makes conventional QA models difficult to converge. But our approach uses an interactive procedure to obtain user’s feedback and allows the model to provide the correct answer.
- For the baseline methods, DMN+ and MemN2N perform similarly and do better than EncDec. Their similar performance (which are worse than our approach) is due to the limitation that they could not learn the accurate meaning of statements and questions with limited resource and then have trouble training the models. But they are superior over EncDec as they treat each input sentence separately instead of modeling very long inputs.

In addition, we also quantitatively evaluate the quality of supplementary question generated by our approach where the details can be found in Appendix A.

Methods	R_{IQA} = 1.0	R_{IQA} = 0.9	R_{IQA} = 0.8	R_{IQA} = 0.7	R_{IQA} = 0.6	R_{IQA} = 0.5	R_{IQA} = 0.4	R_{IQA} = 0.3
IQA Task 1								
CAN+IQA	0.00	0.00	0.10	0.50	0.60	0.70	2.10	0.40
DMN+	42.2	42.1	33.0	28.9	25.0	19.9	17.3	11.6
MemN2N	40.5	38.9	34.4	30.0	23.9	18.4	16.9	13.9
EncDec	53.6	54.3	52.9	53.5	51.8	50.1	45.1	44.8
IQA Task 4								
CAN+IQA	0.00	1.30	0.10	0.60	1.10	1.40	1.20	1.20
DMN+	53.5	56.1	50.4	40.7	34.5	27.4	23.4	16.6
MemN2N	50.4	50.1	41.8	36.1	29.5	25.3	18.7	15.8
EncDec	55.9	54.9	52.5	49.2	45.9	38.9	30.4	24.2
IQA Task 7								
CAN+IQA	0.30	2.10	2.50	1.80	2.00	0.70	0.20	0.30
DMN+	54.1	50.3	47.7	42.3	38.1	33.9	27.7	27.6
MemN2N	54.6	52.0	46.3	40.8	36.1	32.4	23.3	19.6
EncDec	55.5	50.9	48.6	44.7	39.1	32.3	31.9	26.6

Table 7: Performance comparison of various models in terms of test error rate (%) based on interactive question answering datasets with different IQA ratios.

5.5 QUALITATIVE ANALYSIS OF INTERACTIVE MECHANISM

In this section, we qualitatively show the attention weights over input sentences generated by our model on both QA and IQA data. We train our model (CAN+IQA) on task 1 of ibAbI dataset with $Q_{IQA} = 0.9$, and randomly select one IQA example from the testing data. Then we do the prediction on this IQA problem. In addition, we change this instance to a QA problem by replacing the question “Where is she?” with “Where is Sandra?”, and then do the prediction as well. The prediction results on both QA and IQA problems are shown in Table 8. From the results, we observe the following: 1) The attention that uses user’s feedback focuses on the key relevant sentence while the attention without feedback only focuses on an unrelated sentence. This happens because utilizing user’s feedback allows the model to understand a question better and locate the relevant input sentences. This illustrates the effectiveness of an interactive mechanism on addressing questions that require additional information. 2) The attention on both two problems can finally focus on the relevant sentences, showing the usefulness of our model for solving different types of QA problems.

6 CONCLUSION

In this paper, we present a self-adaptive model, CAN, which learns more accurate representations for statements and questions. More importantly, our model is aware *what it knows* and *what it does not know* within the context of the story, and takes an interactive mechanism to answer a question. Hence, our model takes an important step towards having a natural and intelligent conversation

Input Sentences	Support	QA Data	IQA Data	
			Before IM	After IM
Mary journeyed to the kitchen.	yes	0.00	0.99	0.00
Sandra journeyed to the kitchen.		0.00	0.00	0.00
Mary journeyed to the bedroom.		0.00	0.00	0.00
Sandra moved to the bathroom.		0.00	0.00	0.00
Sandra travelled to the office.		0.99	0.00	0.99
Mary journeyed to the garden.		0.00	0.00	0.00
Daniel travelled to the bathroom.		0.00	0.00	0.00
Mary journeyed to the kitchen.		0.00	0.00	0.00
John journeyed to the office.		0.00	0.00	0.00
Mary moved to the bathroom.		0.00	0.00	0.00
			Q: Where is Sandra? A: Office	Q: Where is she? SQ: Who is she? FB: Sandra A: Office

Table 8: Examples of sentence attention weights obtained by our model in both QA and IQA data. “Before IM” indicates the sentence attention weights over input sentences before the user provides a feedback. “After IM” indicates the sentence attention weights updated by user’s feedback. The attention weights with value as 0.00 are very small. The results show that our approach can attend the key relevant sentences for both QA and IQA problems.

with humans. In the future, we plan to employ more powerful attention mechanisms with explicit unknown state modeling and multi-round feedback-guided fine-tuning to make the model fully self-aware, self-adaptive, and self-taught. We also plan to expand our results to harder co-reference and interactive visual QA tasks with uncertainty modeling.

REFERENCES

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015.
- Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL workshop*, 2005.
- Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *Trans. Neur. Netw.*, 5(2):157–166, 1994. ISSN 1045-9227.
- Antoine Bordes and Jason Weston. Learning end-to-end goal-oriented dialog. *CoRR*, abs/1605.07683, 2016.
- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP*, pp. 1724–1734, 2014.
- Anthony Fader, Luke Zettlemoyer, and Oren Etzioni. Open question answering over curated and extracted knowledge bases. In *KDD*, pp. 1156–1165, 2014.
- David Golub and Xiaodong He. Character-level question answering with attention. *CoRR*, abs/1604.00727, 2016.
- Alex Graves, Greg Wayne, and Ivo Danihelka. Neural Turing machines. *CoRR*, abs/1410.5401, 2014.
- Karl Moritz Hermann, Tomáš Kociský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. Teaching machines to read and comprehend. In *NIPS*, pp. 1693–1701, 2015.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8): 1735–1780, 1997.

- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Victor Zhong, Romain Paulus, and Richard Socher. Ask me anything: Dynamic memory networks for natural language processing. In *ICML*, pp. 1378–1387, 2016.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. *CoRR*, abs/1606.00061, 2016.
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. Effective approaches to attention-based neural machine translation. *CoRR*, abs/1508.04025, 2015.
- Volodymyr Mnih, Nicolas Heess, Alex Graves, and Koray Kavukcuoglu. Recurrent models of visual attention. In *NIPS*, pp. 2204–2212, 2014.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: A method for automatic evaluation of machine translation. In *Association for Computational Linguistics*, pp. 311–318, 2002.
- Sukhbaatar Sainbayar, Szymon Antoniak, Weston Jason, and Fergus Rob. End-to-end memory networks. In *NIPS*, pp. 2440–2448, 2015.
- Denis Savenkov and Eugene Agichtein Emory. When a knowledge base is not enough: Question answering over knowledge bases with external text data. In *SIGIR*, pp. 235–244, 2016.
- Paul Hongsuck Seo, Zhe Lin, Scott Cohen, Xiaohui Shen, and Bohyung Han. Hierarchical attention networks. *CoRR*, abs/1606.02393, 2016.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *NIPS*, pp. 3104–3112, 2014.
- Jason Weston. Dialog-based language learning. *NIPS*, 2016.
- Jason Weston, Antoine Bordes, Sumit Chopra, and Tomas Mikolov. Towards ai-complete question answering: A set of prerequisite toy tasks. *CoRR*, abs/1502.05698, 2015.
- Qi Wu, Peng Wang, Chunhua Shen, Anton van den Hengel, and Anthony R. Dick. Ask me anything: Free-form visual question answering based on knowledge from external sources. *CoRR*, abs/1511.06973, 2015.
- Caiming Xiong, Stephen Merity, and Richard Socher. Dynamic memory networks for visual and textual question answering. In *ICML*, pp. 2397–2406, 2016.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C. Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. *CoRR*, abs/1502.03044, 2015.
- Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alexander J. Smola. Stacked attention networks for image question answering. *CoRR*, abs/1511.02274, 2015.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alexander J. Smola, and Eduard H. Hovy. Hierarchical attention networks for document classification. In *HLT*, pp. 1480–1489, 2016.
- Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. Recurrent neural network regularization. *CoRR*, abs/1409.2329, 2014.

A SUPPLEMENTARY QUESTION ANALYSIS

We quantitatively evaluate the quality of supplementary question generated by our model based on IQA dataset. All baseline methods are designed to only predict an answer, none of them can generate a question. Thus we design another baseline method to generate supplementary question based on EncDec, and denote it as $EncDec^*$. Specifically, in training procedure, $EncDec^*$ takes statements and questions as input. If the supplementary question is available, it is used as output; otherwise the corresponding answer is viewed as output. Similar to our approach, the EOS symbol is used to determine whether the generated sentence is question or not, where the question ends with question mark and the answer ends with period symbol.

To test model’s ability to generate supplementary question, we define some following metrics. Suppose the number of problems is N , and the number of problems having supplementary question is N_s . Then $N_a = N - N_s$ is the number of remaining problems. Let $SQueAcc = \frac{\hat{N}_s}{N_s}$ is the fraction of IQA problems which can be correctly estimated, and $AnsAcc = \frac{\hat{N}_a}{N_a}$ is the fraction of remaining problems which can be correctly estimated as QA problem. So $SQueAnsAcc = \frac{\hat{N}_s + \hat{N}_a}{N}$ is the overall accuracy. In addition, the widely used BLEU (Papineni et al., 2002) and METEOR (Banerjee & Lavie, 2005) are also adopted to evaluate the quality of generated supplementary question. The results of our method and baseline method are presented in Table 9.

From the results, we can observe that 1) Two models can almost correctly determine whether it is time to output a question or not; 2) Two models are able to generate the correct supplementary questions whose contents exactly match with the ground truth. There is no surprise that $EncDec^*$ also performs well in generating question, because it is specifically designed for only outputting questions. Thus, if given enough training data, $EncDec^*$ could predict good questions. However, the limitation is that it cannot predict a supplementary question and an answer at the same time. Different from $EncDec^*$, our approach can accurately know when to output an answer or when to generate a supplementary question.

IQA Ratio	Method	SQueAcc	AnsAcc	SQueAnsAcc	BLEU-1	BLEU-4	METEOR
IQA Task 1							
0.8	CAN+IQA	100%	100%	100%	100%	100%	100%
	$EncDec^*$	100%	99.5%	99.9%	100%	100%	100%
IQA Task 4							
0.8	CAN+IQA	100%	100%	100%	100%	100%	100%
	$EncDec^*$	100%	100%	100%	100%	100%	100%
IQA Task 7							
0.8	CAN+IQA	100%	100%	100%	100%	100%	100%
	$EncDec^*$	100%	100%	100%	100%	100%	100%

Table 9: Performance comparison of the generated supplementary question quality with R_{IQA} as 0.8. Both two methods achieve 100% under all metrics in all tasks with other different R_{IQA} values.