

Multi Document Reading Comprehension

Avi Chawla

avi.chawla.cse16@iitbhu.ac.in

Indian Institute Of Technology (BHU) Varanasi
Varanasi, India

ABSTRACT

Reading Comprehension (RC) is a task of answering a question from a given passage or a set of passages. In the case of multiple passages, the task is to find the best possible answer to the question. Recent trials and experiments in the field of Natural Language Processing (NLP) have proved that machines can be provided with the ability to not only process the text in the passage and understand its meaning to answer the question from the passage, but also can surpass the Human Performance on many datasets such as Stanford's Question Answering Dataset (SQuAD) [38]. This paper presents a study on Reading Comprehension and its evolution in Natural Language Processing over the past few decades. We shall also study how the task of Single Document Reading Comprehension acts as a building block for our Multi-Document Reading Comprehension System. In the latter half of the paper, we'll be studying about a recently proposed model for Multi-Document Reading Comprehension – RE³QA [22] that is comprised of a Reader, Retriever, and a Re-ranker based network to fetch the best possible answer from a given set of passages.

KEYWORDS

Reading Comprehension, Neural Networks, Question Answering, Natural Language Processing.

1 INTRODUCTION

The ability to read and understand the unstructured text, and then answer the questions about it is an ordinary skill among literate humans. But for machines, it is not easy. Teaching our machines to read natural human language has always been a long-standing goal expected from Natural Language Processing [21]. Throughout many years of efforts, researchers in the field of NLP have used more and more powerful NLP tools to analyze and understand different aspects of human texts. A point that arises here is why we should build these tools, and what is the need behind building such tools? Basically, how can a machine's reading ability be evaluated is the primary question that comes to mind. The answer to the question lies in the task of Reading Comprehension. Just like human knowledge is tested based on asking questions, we should similarly ask machines about what it has understood from the text.




Figure 1: Diagram representing a basic Reading Comprehension System

Reading Comprehension is not a new topic. As earlier as 1977, Researchers already realized the importance of Reading Comprehension. *Wendy Lehnert* [1] stated the following line in his book, "The Process of Question Answering":

"Only when we can ask a program to answer a question about what it reads will we be able to begin to access that program's comprehension."

Reading Comprehension can be cast as a Question Answering task, which is also an application of Natural Language Processing(NLP) wherein a model is supposed to answer a question put up by the user. The machine answers the question by arbitrarily using some complex, unstructured, semi-structured knowledge bases and returns the correct answer to its user. Question Answering is another of the big successes in using Deep Learning inside Natural Language Processing, and it's a technology that has some obvious commercial usage.

S.No.	Question Category	Description
1	Verification Questions	Intends a 'Yes' or 'No' answer.
2	Factoid-type Questions	What, Which, When, Who or How?
3	Casual Questions	Why or How?
4	Hypothetical Questions	What would happen if ...?
5	Complex Questions	What are the reasons for ...?
6	Definition Questions	What is the definition of ...?
7	Quantification Questions	How much is ...?

Table 1: A few question types generally posed by humans.

Question Answering systems can be classified into two types:

- (1) *Closed-Domain Question Answering* [26]: This is an easier type of Question Answering wherein the question belongs to a particular domain/topic. Questions in this domain are generally asked for descriptive or factoid information.
- (2) *Open-Domain Question Answering* [9]: In contrast to the category mentioned above, here the question can belong to any class and thus, can pose anything in the world. These systems generally require much data to answer the question accurately, e.g., Wikipedia Dump.

Here, the goal is to understand the passage and answer the question from the passage. NLP systems in such a scenario are expected to leverage the power of Natural Language Understanding, Natural Language Inference or Entailment, and Information Retrieval, etc. Thus, reading comprehension is challenging for us.

There are two main paradigms for Question Answering. These two paradigms cover the majority of Question Answering Systems that we use or build today. These are:

- (1) *IR-based Approaches*: These systems concentrate on finding an answer to a question by looking at strings of text. Google Search is a common example of this approach. [2] and [16] also presents a detailed survey on the same.
- (2) *Knowledge-based or Hybrid Approaches*: These systems build the answer from understanding a part of the text. Natural Language Understanding plays a crucial role in building such systems. Examples include IBM Watson, Apple Siri. A detailed study on such systems is given by [11] and [56]

A comprehensive study between the IR-based approaches and Knowledge-based approaches is presented in [30]. Question asked to a Question Answering can be of many times, ranging from Factoid Questions to Complex Narrative Questions. A few types of Question Categories are shown in Table 1.

2 HISTORY OF QUESTION ANSWERING

Two of the oldest Question Answering system were called BASEBALL [15] and LUNAR [52].

BASEBALL was designed way back in 1961 to answer questions related to US Baseball League for little more than a year. It accepted question in English language and answered from an already stored knowledge base.

On the other hand, LUNAR was developed in 1971 and was capable of answering questions about moon rocks and soil that the Apollo 11 Mission had gathered in July 1969.

Both the systems, BASEBALL and LUNAR, were Closed-Domain Question Answering systems as they answered questions related to a particular topic only and not about anything in the world at that time.

SHRDLU [34] was another such domain-restricted question answering system that was developed in late 1968. Researchers also called this system the "Blocks World" wherein a user interacted in English with SHRDLU to move various objects in the "Blocks World." It comprised of different basic geometrical objects such as Cone, Cylinder, Cube, etc.




Figure 2: SHRDLU system working

Simmons [42] also did an exploration of answering questions using dependency parses matching techniques of question and the answer. This is demonstrated in the figure below.

Question: What do worms eat?



Figure 3: Question

Potential Answers:




Figure 4: Potential Answers

Considering the above example, the system took the question and parsed it into a dependency format to establish a relationship between the words. The next step was to find a matching dependency structure in the extensive database of answers. Figure 4 above shows some of the potential answers that the system fetched through its matching algorithm.

At the start of this decade, IBM's Watson [14] [7] came at the forefront of a new era of Cognitive Computing, beating humans in the very famous Jeopardy Challenge, back in 2011. Cognitive Computing is a very different type of computing that is very different from the programmed systems that have preceded it. With today's world of big data and the need for more complex evidence-based decisions, approaches such as Decision Tress [41] or Tabulated Answer Retrieval based approach often breaks or fails to keep up with the available information. Cognitive Computing enables machines to create a new kind of value and find answers through insights drawn from vast amounts of data. Watson and its Cognitive capabilities adhere to some of the critical abilities and elements of human expertise, leading to a system that solves a problem as a human does. When we seek to understand something and try to make a decision, we go through four key steps:

- (1) We *observe* the visible phenomena and bodies of evidence.
- (2) We mix what we know and see to *interpret* its meaning and generate a hypothesis about what it means.
- (3) We *evaluate* which hypothesis is right and wrong.
- (4) Finally, we *decide*, choosing the option that seems best and acting accordingly.

Just as humans become experts by going through Observation, Evaluation, and Decision-making, Cognitive Systems like Watson use similar processes to explain the information they read. Watson can do all this at a massive speed and scale.

3 EVOLUTION

Although the importance of Reading Comprehension has been recognized for over 45 years, Reading Comprehension is a very new domain to explore in the Natural Language Processing Research Community. This is evident from the following information:

Pre-2015

In this period, i.e., before 2015, we didn't have any statistical NLP systems capable of reading a simple passage and answering questions from it. Considering the Dataset aspect of this period, before 2015, we hardly had two datasets consisting of passage, question, and answer pairs. There were:

- (1) MCTest [40]: 2600 pairs of question, passage and answer.
- (2) ProcessBank [6]: 500 pairs of question, passage and answer.

If we look at the system or model aspect of this period, most of the models built for experimentation in this period were hand-built systems like linear classifiers with some commonly used linguistic features known at that time.

Post-2015

In this period, some significant changes came in the field of Reading Comprehension. Talking about the Dataset aspect of this period, post-2015, we had enormous datasets with us to build Reading Comprehension or Question Answering Systems. These datasets were collected from various sources like News Articles or Wikipedia or children's books. A few examples of these datasets are:

- (1) CNN Daily Mail [17]
- (2) SQuAD [38]
- (3) LAMBADA [35]
- (4) Who Did What (WDW) [33]
- (5) Children's Book Test (CBT) [19]
- (6) MSMACRO [32]
- (7) Maluuba News QA [45]
- (8) TriviaQA [23]

- (9) WikiReading [18]
- (10) SearchQA [13]
- (11) CMU Race [25]

One common thing among all these datasets was that they all contained more than 1 lakh pairs of Passage, Question, and Answer. If one looks at the system side of this period, a revolution in the model designs and architectures took place. We saw that all the hand-built systems of the Pre-2015 period evolved into robust end-to-end neural network architectures. Not only did it outperform the earlier-used linear classifier based models but also started achieving near-human performance on many datasets mentioned above.

4 APPLICATIONS OF QUESTION ANSWERING IN DAILY-LIFE

These days it seems like every major tech company either has or is working on a Virtual Assistant to enhance the customer experience. Researchers in Artificial Intelligence have been dreaming of computers we could talk to for decades now. It appears that now, we have started to get to the point where these friendly, helpful voices are beginning to feel unavoidable. Recent innovations in the field of Artificial Intelligence have led to the development of numerous Virtual Assistants that help us significantly in our day to day lives. These Virtual Assistants are based on Machine Learning Techniques like Question Answering Systems. Each Virtual Assistant and its closely related chat-bot systems run just a bit differently, but they all have the same result, making our lives much more comfortable. These chat-bots eliminate the availability of humans on a variety of tasks like the management of a company's Customer Assistance portal, guiding school students with their study plan, and many more. Virtual Assistants and the technology around them are helping us do the things we do every day just quicker and more comfortable. In recent times, these technologies have now even surpassed the Question Answering System. These systems have become proficient in making real-time conversations with the users and asking follow-up questions to understand better what we want. Some popularly known examples of Virtual Assistants that we use daily are Google Assistant by Google, Siri by Apple, Alexa by Amazon, and Cortana by Microsoft. A comprehensive overview of such Virtual Assistants and their working is presented in [46].

5 RELATED WORKS TO QUESTION ANSWERING

Machine Reading Comprehension is a very old problem in the field of NLP. With massive collections of full-text documents, i.e., the Web, simply returning relevant documents is of limited use. So instead, what we look for these days is the direct answer to our questions. In this regard, many recent Natural Language Processing Systems have heavily

leveraged the power of Deep Learning and other massive NLP architectures for providing a solution to the Question Answering problem of NLP.

Statistical Question Answering

A majority of basic and traditional approaches to answering question answering involve techniques such as rule-based algorithms or the employment of linear classifiers on a set of hand-engineered features. Lynette Hirschman proposed Deep Read [20] built on a training corpus of short stories in the books of students of 3rd to 6th grade. The task aimed at reading a story and answering a question based on that. The proposed system made use of a simple bag-of-words approach technique for matching the question with the passage to find the most relevant answer. [40] proposed two baselines for the task of Question Answering. One of them uses some simple techniques such as a sliding window to match the bags-of-words from the passage. The other made use of word-distances between words present in the question and in the document. Some statistical approaches have also focused on learning a structured representation of the entities in the data and storing the relations in the document as a knowledge-base [5]. The next step focuses on the conversion of the question to a structured query. This allows the model to match the content of the question to that of the knowledge-base. Other statistical approaches have revolved around the usage of semantic features and the syntactic features of words such as POS [49]. One of the finest and competitive baseline came through the efforts of [8]. They proposed an approach using a set of carefully crafted lexical, syntactic, and word order features of the Question and Answering data. Christopher Burges [40] proposed a challenge problem for Artificial Intelligence and released a Reading Comprehension training corpus called the MCTest [40].

Neural Question Answering

Deep Learning has always been the key to many extraordinary advancements done in the field of NLP. Neural Models find a vast application in the field of Question Answering also. For instance, the use of Neural-attention models has been widely applied in recent times for machine comprehension in NLP. [17] proposed an Attentive Reader model in addition to a large cloze-style question answering dataset CNN/Daily Mail that led the foundation of many Deep Learning-based architectures in the future. Another dataset was released that originated from the Children's book data [19]. As a baseline, they proposed a window-based memory network. The use of pointer networks has also gained popularity for answer prediction [24]. [43] also made attempts in using the neural attention model for the task of machine comprehension. One major revolution in Neural QA came about through

S.No.	Dataset Name	Description	Release Year	Training Samples	Testing Samples
1	CliCR [44]	Created from reports of Clinical Cases	2018	91,344	7,184
2	CNN [17]	Created from News Articles	2016	380,298	3,198
3	Daily Mail [17]	Created from News Articles	2016	879,450	53,182
4	CoQA [39]	Meant for Building Conversational Question Answering system	2019	1,27,000*	-
5	MS MARCO [32]	Obtained from Bing’s Search Results	2016	1,010,916	-
6	NewsQA [45]	Created from CNN News Articles	2016	119,633	-
7	RACE [25]	Created from English examinations in China	2016	87,866	4,934
8	SQuAD [38]	Created from Wikipedia Data	2016	129,941	5,915

Table 2: A few question types generally posed by humans.

the release of the SQuAD dataset [38], which comprises of Question, Passage and Answer triplets. The answer is always a continuous span to words in the passage, and the questions appear to be more realistic in this dataset. Since then, many attempts have been made to build a Question-Answering system on SQuAD. [50] proposed an end-to-end neural network-based approach that comprises of a Match-LSTM encoder. [55] make use of a dynamic chunk reader, which is a neural reading comprehension model aimed at the extraction of a set of potential answers from the passage and then ranks these answers to answer the question precisely.

The scope of Question Answering is not restricted to only text-based Question Answering. Many attempts have also been made in the past to propose models for Visual Question Answering. For instance, [28] proposed a hierarchical co-attention model for this task. The performance of this model on the COCO-VQA dataset [4] achieved the state of the art results. This followed the development of another VQA model that made use of a co-attention mechanism to compute a conditional representation of the image given the question and vice versa [29]. Inspired by the above idea, [4] proposed a Dynamic Co-attention Network (DCN) and an year later a Dynamic Co-attention Network+ (DCN+) [53]. The architecture comprises of a Co-attentive encoder and a Dynamic Decoder. The answer is predicted with the help of a Pointer Network [48].

In recent times, NLP has witnessed a considerable work in the direction of generating Contextualised Word Embeddings (CWEs) [12] [37] [3]. CWEs have been inspired by the idea of Sense Embeddings, which suggests using different Word Embeddings for each sense of the word. Word Embeddings models such as Word2Vec [31] and Glove [36] ignore polysemy of words and provide only a single embedding for each word. In contrast, CWEs provide a separate embedding for each word’s occurrence in the context it appears. One such revolution in the field of CWEs can be realized by the efforts made by the [12] and their proposed BERT model. The model is based on a Transformer [47] architecture that makes use of

Attention Mechanism to find the Contextualised Embeddings for words. Recent approaches in Question Answering have also focused on using these models in place of the previously used Static pre-trained Embeddings.

6 DATASETS

In the past few years, researchers in the field of NLP have put in many efforts of curate Reading Comprehension datasets. Most of these datasets are composed of a Passage, Question, and Answer triplets and consists of approximately one lakh training examples. A brief overview and statistics of some of the popularly used datasets for Reading Comprehension are given in Table 2.

7 MULTI DOCUMENT READING COMPREHENSION

The task of natural Reading Comprehension is of minimal practical use. It does not allow the system to search in multiple documents to search for an answer. The approaches discussed in the Related Work Section primarily focus on such a scenario of Question Answering. In contrast, Multi-Document Reading Comprehension refers to the task of finding the correct answer to a question by searching in more than one documents. This task has more applications as an end product than the Single document Reading Comprehension.

Most of the approaches proposed in this area have heavily relied on the working mechanism of a natural Reading Comprehension system. This system functions very in a fashion remarkably similar to that of the natural one. It just adds an extra step of ranking the answers obtained from various sources. A basic pipeline to this task is explained below in the following steps:

- (1) *Document Selection*: Out of the myriad of documents available in the knowledge base, this step involves the selection of top n -documents that match to the question.

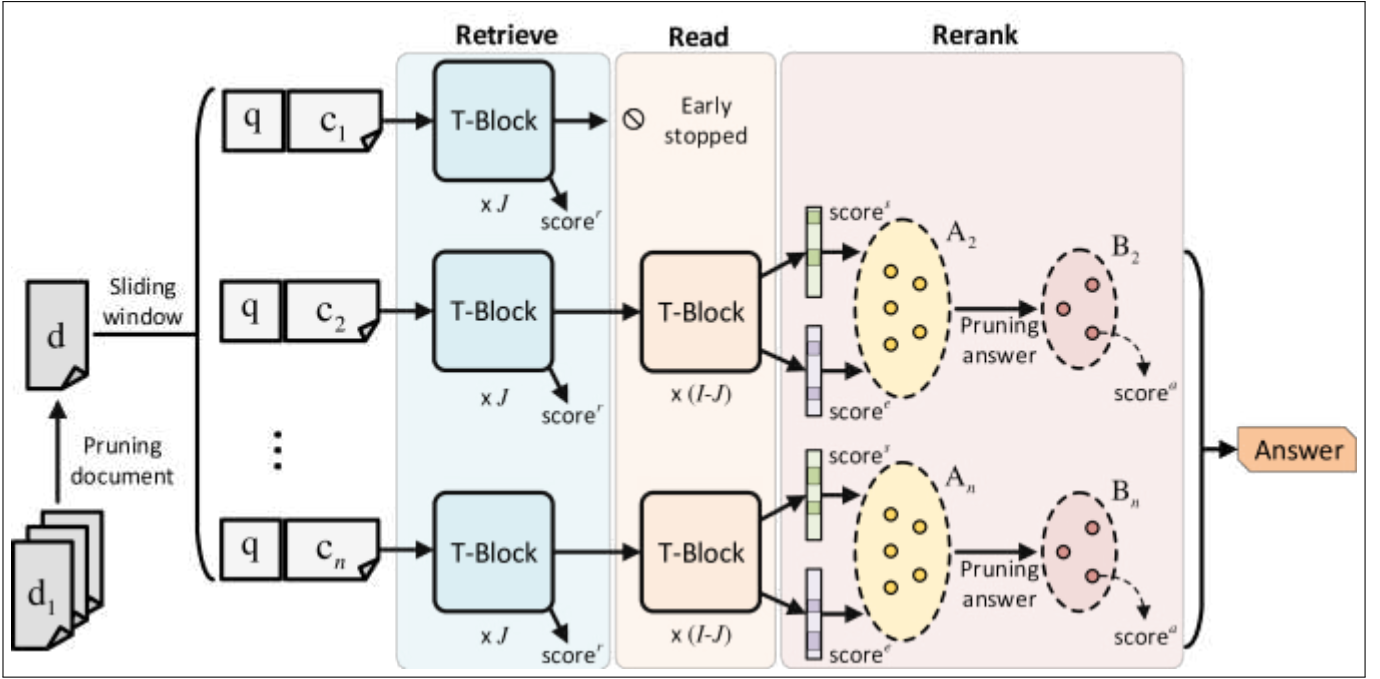


Figure 5: Figure shows the flow of final answer prediction from the RE³QA approach to Multi-Documen Reading Compre-hension

- (2) *Finding Candidate Answers*: In this step, the model tries to find some potential answers to the question from the selected sets of documents.
- (3) *Answer Ranking and Selection*: This step aims to rank the answers extracted from the previous step according to their relevance to the question and return the best answer to the question.

In recent times, researchers working in the area on Multiple Document Reading Comprehension have gained considerable success in building efficient systems that can find answers to question hidden in volumes of data [54] [22]. These systems leverage the power of highly efficient and effective algorithms and architectures at each step. [10] focused on the construction of a coarse-to-fine grained framework that answers the question from a retrieved document summary. [27] proposed a pipeline system consisting of a paragraph selector and a paragraph reader. [51] extended [10] and introduced a network for cross-passage answer verification. This paper studies one such recent approach that makes use of the end-to-end network for efficient answer retrieval from multiple documents. This approach is based on Retriever, Reader, and Re-ranker model that works in the exact similar to the pipeline described above. The approach is described in the following section.

8 RE³QA: RETRIEVER, READER, AND RE-RANKER BASED QUESTION ANSWERING

Figure 1 gives an overview of the RE³QA model for Multi-Documen Reading Comprehension taken from [22]. The model comprises of three sub-networks. These are:

- (1) *Retriever*: Retriever aims to select a few relevant documents from volumes of data available to us as a knowledge base.
- (2) *Reader*: Reader's works on the output of Retriever, and it is used to extract candidate answers from the retrieved documents.
- (3) *Re-ranker*: Re-ranker aims to re-score multiple candidate answers generated by the Reader and return the best answer to the question.

In contrast to previously adopted approaches, the RE³QA model integrates the Retriever, Reader, and Re-ranker components into a unified network instead of a pipeline and is trained using an end-to-end strategy. One more improvement is brought to the model by providing it Contextualised Word Embeddings (CWEs) of BERT rather than pre-trained Embeddings such as Glove Vectors. The use of BERT has already been found to be very useful on many NLP task, and the use of BERT Contextualised Embeddings in RE³QA is what accounts for its exceptional performance gains. This

helps the model better understand the text and establish relationships between the question and the knowledge base to be searched.

The three sub-networks that govern the entire answer generation process run through a five-staged process. These are described in the following sub-sections:

Document Pruning

This is the first of five steps involved in the answer retrieval process for this model. As the name suggests, this step involves the selection of potential documents from the entire database which may contain the answer to the question.




Figure 6

The input to this step is the Question Q and the set of Documents $D \in \{D_1, D_2, D_3, \dots, D_N\}$. Paragraphs in each Document are ranked according to TF-IDF scores with the question and top-K paragraphs are selected for further steps. These paragraphs are then concatenated to form a single document d .

Segment Encoding

As the name suggests, this part involves splitting of the Document d obtained in the previous step into small components. To achieve this, the model slides a window of length l with a stride r over the Document d . This results in the formation of a set of text-segments $C \in \{C_1, C_2, C_3, \dots, C_N\}$.

In the next step, these segments are individually concatenated with the Question Q and encoded with the help of the BERT model. Input to the Transformer Block is $QC = ['[CLS]'; Q; '[SEP]'; C_i; '[SEP]']$ where $'[CLS]'$ is called the Classification Token and $'[SEP]'$ is a token used for separating sentences. The length of vector QC is L_X . This is then passed into the TransformerBlock which returns I hidden layers. The value of I for bert-small is 12 and for bert-large is 24. Each hidden layer is of size $L_X * D_H$ where D_H is the size of hidden layer.

$$h_i = \text{TransformerBlock}(h(i-1)) \text{ where } i \in [1, I]$$

$$h_i \in \mathbb{R}^{L_X * D_H}$$

Early Stopped Retriever

The above encoding is very effective and appealing but there is a lot of computational efficiency involved in it as well. This

becomes evident from the fact that average number of text-segments C for each question is roughly 20. Each segment is concatenated with the Question Q and passed through at least 12 hidden layers for encoding it. Each hidden layer is at least of size 768. The Early Stopped Retriever applies another level of pruning. In a nutshell, it ranks all the text-segments using the hidden representations obtained in Segment Encoding step. This is done using a series of Multi-Layer Perceptron and a self-aligning layer with the hidden states. To train this component, it is trained on an objective function which computes the loss over those ranked segments that don't contain any word of the final answer. From this step, we select M text-segments according to the ranks obtained for further processing.

Distantly-Supervised Reader

The aim of the Reader is to generate multiple candidate answers. This is achieved by projecting the last hidden layer of each text-segment obtained in the previous step into two scores. These scores determine the potential start and end of the answer. The objective function defines a loss that is totally similar to that of the Early Stopped Retriever model. The module returns the start and the end position of the preliminary answers from each of the M text-segments. These are: $A = [A_1, A_2, A_3, \dots, A_M]$.

Answer Re-ranker

This module ranks the potential answers from the previous module with the help of their span representations. The objective function defined here is also similar to that described in the earlier modules. This objective function takes into account whether a particular span representation contains at least one token from the original answer or not. Those representations that don't contain any such token are used in the loss function.

In the above sub-sections, the objective of each of the three-modules, i.e. the Retriever, Reader and the Re-ranker are defined. Rather than separately training each component, the training strategy involves an end-to-end strategy for training them together. This is done by adding the three loss functions and train the entire network in based on this loss function using Gradient Descent.

Datasets

Experimentation is done on 3 prominent and widely used Question Answering datasets. These are:

- (1) TriviaQA-Wikipedia: This dataset consists of Top-10 search documents from Wikipedia.

- (2) SQuAD Document: It is variant of the original SQuAD dataset. Difference is that it maps the question to the entire Wikipedia Page instead of a specific paragraph.
- (3) SQuAD Open: This dataset is basically the entire Wikipedia-domain.

Results

The RE³QA model, when trained with the appropriate data and training strategy, outperforms the previously adopted approaches and achieves state-of-the-art results on the three-datasets written above. The results for each of these models is shown in the following table:

Dataset	RE ³ QA Large	RE ³ QA Base
TriviaQA-Wikipedia	83.0	79.9
SQuAD Document	87.20	84.81
SQuAD Open	50.2	48.4

Table 3: F1(%) scores obtained by using the BERT-base and the BERT-large model on each of the three datasets. There are taken directly from [22].

9 CONCLUSION

This paper presented a study on the evolution of Question Answering as an AI problem and its importance in our day to day lives. We also discussed the history of Question Answering over the past few decades and how the advancements in the field have led to the development of plenty of Chat-bots using Question Answering techniques. In addition to this, we discussed two significant periods of Question Answering, i.e., the Pre-2015 and the Post-2015. While coming to the end of our study, we saw one of the recently proposed approaches to Multi-Document Reading Comprehension that made use of a unified network to answer questions by looking into multiple documents at a time. The model discussed is currently the state-of-the-art in this domain on three challenging datasets of Question Answering.

REFERENCES

- [1] 1980. The Process of Question Answering - A Computer Simulation of Cognition. *American Journal of Computational Linguistics* 6, 3-4 (1980). <https://www.aclweb.org/anthology/J80-3006>
- [2] Zahra Abbasiyantaeb and Saeedeh Momtazi. 2020. Text-based Question Answering from Information Retrieval and Deep Neural Network Perspectives: A Survey. *ArXiv abs/2002.06612* (2020).
- [3] Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. Contextual String Embeddings for Sequence Labeling. In *COLING 2018, 27th International Conference on Computational Linguistics*. 1638–1649.
- [4] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. VQA: Visual Question Answering. *CoRR abs/1505.00468* (2015). [arXiv:1505.00468](http://arxiv.org/abs/1505.00468) <http://arxiv.org/abs/1505.00468>
- [5] Jonathan Berant and Percy Liang. 2014. Semantic Parsing via Paraphrasing. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Baltimore, Maryland, 1415–1425. <https://doi.org/10.3115/v1/P14-1133>
- [6] Jonathan Berant, Vivek Srikumar, Pei-Chun Chen, Abby Vander Linden, Brittany Harding, Brad Huang, Peter Clark, and Christopher D. Manning. 2014. Modeling Biological Processes for Reading Comprehension. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Doha, Qatar, 1499–1510. <https://doi.org/10.3115/v1/D14-1159>
- [7] E. Brown. 2013. Watson: The Jeopardy! Challenge and beyond. In *2013 IEEE 12th International Conference on Cognitive Informatics and Cognitive Computing*. 2–2.
- [8] Danqi Chen, Jason Bolton, and Christopher D. Manning. 2016. A Thorough Examination of the CNN/Daily Mail Reading Comprehension Task. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, 2358–2367. <https://doi.org/10.18653/v1/P16-1223>
- [9] Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading Wikipedia to Answer Open-Domain Questions. *CoRR abs/1704.00051* (2017). [arXiv:1704.00051](http://arxiv.org/abs/1704.00051) <http://arxiv.org/abs/1704.00051>
- [10] Eunsol Choi, Daniel Hewlett, Jakob Uszkoreit, Illia Polosukhin, Alexandre Lacoste, and Jonathan Berant. 2017. Coarse-to-Fine Question Answering for Long Documents. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Vancouver, Canada, 209–220. <https://doi.org/10.18653/v1/P17-1020>
- [11] Wanyun Cui, Yanghua Xiao, Haixun Wang, Yangqiu Song, Seung-won Hwang, and Wei Wang. 2019. KBQA: Learning Question Answering over QA Corpora and Knowledge Bases. *CoRR abs/1903.02419* (2019). [arXiv:1903.02419](http://arxiv.org/abs/1903.02419) <http://arxiv.org/abs/1903.02419>
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [13] Matthew Dunn, Levent Sagun, Mike Higgins, V. Ugur Güney, Volkan Cirik, and Kyunghyun Cho. 2017. SearchQA: A New Q&A Dataset Augmented with Context from a Search Engine. *CoRR abs/1704.05179* (2017). [arXiv:1704.05179](http://arxiv.org/abs/1704.05179) <http://arxiv.org/abs/1704.05179>
- [14] D.A. Ferrucci. 2012. Introduction to “This is Watson”. *IBM Journal of Research and Development* 56 (05 2012), 1:1–1:15. <https://doi.org/10.1147/JRD.2012.2184356>
- [15] Bert F. Green, Alice K. Wolf, Carol Chomsky, and Kenneth Laughery. 1961. Baseball: An Automatic Question-Answerer. In *Papers Presented at the May 9-11, 1961, Western Joint IRE-AIEE-ACM Computer Conference* (Los Angeles, California) (*IRE-AIEE-ACM '61 (Western)*). Association for Computing Machinery, New York, NY, USA, 219–224. <https://doi.org/10.1145/1460690.1460714>
- [16] Poonam Gupta and Vishal Gupta. 2012. A Survey of Text Question Answering Techniques. *International Journal of Computer Applications* 53 (2012), 1–8.
- [17] Karl Moritz Hermann, Tomás Kociský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching Machines to Read and Comprehend. *CoRR abs/1506.03340* (2015). [arXiv:1506.03340](http://arxiv.org/abs/1506.03340) <http://arxiv.org/abs/1506.03340>
- [18] Daniel Hewlett, Alexandre Lacoste, Llion Jones, Illia Polosukhin, Andrew Fandrianto, Jay Han, Matthew Kelcey, and David Berthelot. 2016. WikiReading: A Novel Large-scale Language Understanding Task over

- Wikipedia. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, 1535–1545. <https://doi.org/10.18653/v1/P16-1145>
- [19] Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2016. The Goldilocks Principle: Reading Children’s Books with Explicit Memory Representations. *CoRR* abs/1511.02301 (2016).
- [20] Lynette Hirschman, Marc Light, Eric Breck, and John D. Burger. 1999. Deep Read: A Reading Comprehension System. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, College Park, Maryland, USA, 325–332. <https://doi.org/10.3115/1034678.1034731>
- [21] L. Virginia Holland. 2006. Kenneth Burke’s Theory of Communication. *Journal of Communication* 10, 4 (02 2006), 174–184. <https://doi.org/10.1111/j.1460-2466.1960.tb00541.x> arXiv:<https://academic.oup.com/joc/article-pdf/10/4/174/22386174/jnlcom0174.pdf>
- [22] Minghao Hu, Yuxing Peng, Zhen Huang, and Dongsheng Li. 2019. Retrieve, Read, Rerank: Towards End-to-End Multi-Document Reading Comprehension. *CoRR* abs/1906.04618 (2019). arXiv:1906.04618 <http://arxiv.org/abs/1906.04618>
- [23] Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension. *CoRR* abs/1705.03551 (2017). arXiv:1705.03551 <http://arxiv.org/abs/1705.03551>
- [24] Rudolf Kadlec, Martin Schmid, Ondrej Bajgar, and Jan Kleindienst. 2016. Text Understanding with the Attention Sum Reader Network. *CoRR* abs/1603.01547 (2016). arXiv:1603.01547 <http://arxiv.org/abs/1603.01547>
- [25] Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. RACE: Large-scale ReAding Comprehension Dataset From Examinations. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Copenhagen, Denmark, 785–794. <https://doi.org/10.18653/v1/D17-1082>
- [26] Sweta Lende and M. Raghuwanshi. 2016. CLOSED DOMAIN QUESTION ANSWERING SYSTEM USING NLP TECHNIQUES. (01 2016). <https://doi.org/10.5281/zenodo.49808>
- [27] Yankai Lin, Haozhe Ji, Zhiyuan Liu, and Maosong Sun. 2018. Denoising Distantly Supervised Open-Domain Question Answering. 1736–1745. <https://doi.org/10.18653/v1/P18-1161>
- [28] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2016. Hierarchical Question-Image Co-Attention for Visual Question Answering. *CoRR* abs/1606.00061 (2016). arXiv:1606.00061 <http://arxiv.org/abs/1606.00061>
- [29] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2016. Hierarchical Question-Image Co-Attention for Visual Question Answering. *ArXiv* abs/1606.00061 (2016).
- [30] Rosy Madaan, Ashok Sharma, Ashutosh Dixit, Deepti Kapri, and Renu Mudgal. 2012. A comparative study of web based and IR based question answering systems. *International Journal of Advances in Computing and Information Technology* 1 (04 2012), 188–194. <https://doi.org/10.6088/ijacit.12.10024>
- [31] Tomas Mikolov, Kai Chen, Greg S. Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. <http://arxiv.org/abs/1301.3781>
- [32] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A Human Generated MACHine Reading COMprehension Dataset. *CoRR* abs/1611.09268 (2016). arXiv:1611.09268 <http://arxiv.org/abs/1611.09268>
- [33] Takeshi Onishi, Hai Wang, Mohit Bansal, Kevin Gimpel, and David A. McAllester. 2016. Who did What: A Large-Scale Person-Centered Cloze Dataset. *CoRR* abs/1608.05457 (2016). arXiv:1608.05457 <http://arxiv.org/abs/1608.05457>
- [34] Santiago Ontañón. 2018. SHRDLU: A Game Prototype Inspired by Winograd’s Natural Language Understanding Work. In *AIIDE*.
- [35] Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Quan Ngoc Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. 2016. The LAMBADA dataset: Word prediction requiring a broad discourse context. *CoRR* abs/1606.06031 (2016). arXiv:1606.06031 <http://arxiv.org/abs/1606.06031>
- [36] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*. 1532–1543. <http://www.aclweb.org/anthology/D14-1162>
- [37] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proc. of NAACL*.
- [38] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100, 000+ Questions for Machine Comprehension of Text. *CoRR* abs/1606.05250 (2016). arXiv:1606.05250 <http://arxiv.org/abs/1606.05250>
- [39] Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. CoQA: A Conversational Question Answering Challenge. *Transactions of the Association for Computational Linguistics* 7 (March 2019), 249–266. https://doi.org/10.1162/tacl_a_00266
- [40] Matthew Richardson, Christopher J.C. Burges, and Erin Renshaw. 2013. MCTest: A Challenge Dataset for the Open-Domain Machine Comprehension of Text. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Seattle, Washington, USA, 193–203. <https://www.aclweb.org/anthology/D13-1020>
- [41] Lior Rokach and Oded Maimon. 2005. *Decision Trees*. Vol. 6. 165–192. https://doi.org/10.1007/0-387-25465-X_9
- [42] R. F. Simmons. 1965. Answering English Questions by Computer: A Survey. *Commun. ACM* 8, 1 (Jan. 1965), 53–70. <https://doi.org/10.1145/363707.363732>
- [43] Alessandro Sordani, Philip Bachman, and Yoshua Bengio. 2016. Iterative Alternating Neural Attention for Machine Reading. *CoRR* abs/1606.02245 (2016). arXiv:1606.02245 <http://arxiv.org/abs/1606.02245>
- [44] Simon Suster and Walter Daelemans. 2018. CliCR: A Dataset of Clinical Case Reports for Machine Reading Comprehension. *CoRR* abs/1803.09720 (2018). arXiv:1803.09720 <http://arxiv.org/abs/1803.09720>
- [45] Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2016. NewsQA: A Machine Comprehension Dataset. *CoRR* abs/1611.09830 (2016). arXiv:1611.09830 <http://arxiv.org/abs/1611.09830>
- [46] Amrita Tulshan and Sudhir Dhage. 2019. *Survey on Virtual Assistant: Google Assistant, Siri, Cortana, Alexa: 4th International Symposium SIRS 2018, Bangalore, India, September 19–22, 2018, Revised Selected Papers*. 190–201. https://doi.org/10.1007/978-981-13-5758-9_17
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. *CoRR* abs/1706.03762 (2017). arXiv:1706.03762 <http://arxiv.org/abs/1706.03762>
- [48] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer Networks. *ArXiv* abs/1506.03134 (2015).
- [49] Di Wang and Eric Nyberg. 2015. A Long Short-Term Memory Model for Answer Sentence Selection in Question Answering. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language*

- Processing (Volume 2: Short Papers)*. Association for Computational Linguistics, Beijing, China, 707–712. <https://doi.org/10.3115/v1/P15-2116>
- [50] Shuohang Wang and Jing Jiang. 2016. Machine Comprehension Using Match-LSTM and Answer Pointer. *CoRR* abs/1608.07905 (2016). arXiv:1608.07905 <http://arxiv.org/abs/1608.07905>
 - [51] Yizhong Wang, Kai Liu, Jing Liu, Wei He, Yajuan Lyu, Hua Wu, Sujian Li, and Haifeng Wang. 2018. Multi-Passage Machine Reading Comprehension with Cross-Passage Answer Verification. *CoRR* abs/1805.02220 (2018). arXiv:1805.02220 <http://arxiv.org/abs/1805.02220>
 - [52] William Woods, Ronald Kaplan, and Bonnie Webber. 1972. The Lunar Science Natural Language Information System: Final Report. (01 1972).
 - [53] Caiming Xiong, Victor Zhong, and Richard Socher. 2017. DCN+: Mixed Objective and Deep Residual Coattention for Question Answering. *CoRR* abs/1711.00106 (2017). arXiv:1711.00106 <http://arxiv.org/abs/1711.00106>
 - [54] Ming Yan, Jiangnan Xia, Chen Wu, Bin Bi, Zhongzhou Zhao, Ji Zhang, Luo Si, Rui Wang, Wei Wang, and Haiqing Chen. 2018. A Deep Cascade Model for Multi-Document Reading Comprehension. *CoRR* abs/1811.11374 (2018). arXiv:1811.11374 <http://arxiv.org/abs/1811.11374>
 - [55] Yang Yu, Wei Zhang, Kazi Saidul Hasan, Mo Yu, Bing Xiang, and Bowen Zhou. 2016. End-to-End Reading Comprehension with Dynamic Answer Chunk Ranking. *CoRR* abs/1610.09996 (2016). arXiv:1610.09996 <http://arxiv.org/abs/1610.09996>
 - [56] Shuguang Zhu, Xiaodian Cheng, and Sen Su. 2020. Knowledge-based question answering by tree-to-sequence learning. *Neurocomputing* 372 (2020), 64–72.