

BiPaR: A Bilingual Parallel Dataset for Multilingual and Cross-lingual Reading Comprehension on Novels

Yimin Jing, Deyi Xiong* and Yan Zhen

School of Computer Science and Technology, Soochow University, Suzhou, China

{yymmjing, corezhen}@gmail.com

dyxiong@suda.edu.cn

Abstract

This paper presents *BiPaR*, a **bilingual parallel** novel-style machine reading comprehension (MRC) dataset, developed to support multilingual and cross-lingual reading comprehension. The biggest difference between BiPaR and existing reading comprehension datasets is that each triple (Passage, Question, Answer) in BiPaR is written parallelly in two languages. We collect 3,667 bilingual parallel paragraphs from Chinese and English novels, from which we construct 14,668 parallel question-answer pairs via crowdsourced workers following a strict quality control procedure. We analyze BiPaR in depth and find that BiPaR offers good diversification in prefixes of questions, answer types and relationships between questions and passages. We also observe that answering questions of novels requires reading comprehension skills of coreference resolution, multi-sentence reasoning, and understanding of implicit causality, etc. With BiPaR, we build monolingual, multilingual, and cross-lingual MRC baseline models. Even for the relatively simple monolingual MRC on this dataset, experiments show that a strong BERT baseline is over 30 points behind human in terms of both EM and F1 score, indicating that BiPaR provides a challenging testbed for monolingual, multilingual and cross-lingual MRC on novels. The dataset is available at <https://multinlp.github.io/BiPaR/>.

1 Introduction

Machine reading comprehension is to evaluate how well computer systems understand natural language texts, where machines read a given text passage and answer questions about the passage. It has been regarded as a crucial technology for many applications such as question answering, dialogue systems (Nguyen et al., 2016; Chen et al.,

*Corresponding author

The Duke of the Mount Deer / 《鹿鼎记》

Madame ...This is the Leader's **Five Dragon Disc**, White Dragon Marshal,' she said. ...You are to return it when you have completed your mission.' 'Yes,' replied Trinket, ... I want you to stay here for a while,' Madame Hong ordered. The rest of you may leave.' **Rootless, Black Dragon, and Yellow Dragon** saluted and left.

洪夫人... “**白龙使**, 这是教主的五龙令... 立功之后, 将令缴回。” 韦小宝应道: “是。”... 三人暂留, 余人退去。” **无根道人和黑龙使、黄龙使**三人行礼退出。

Q_{en}-Q_{zh}: What did Madame Hong give to Trinket? / 洪夫人将什么东西给了韦小宝?

A_{en}-A_{zh}: Five Dragon Disc / 五龙令

Monolingual MRC

Q_{en}: What did Madame Hong give to Trinket?

A_{en}: **Five Dragon Disc**

Multilingual MRC

Q_{en}-A_{en}: Who did Madame Hong allow to leave first? /

Rootless, Black Dragon, and Yellow Dragon

Q_{zh}-A_{zh}: 韦小宝在神龙教中担任什么职位? / **白龙使**

Cross-lingual MRC

Q_{en}: Who did Madame Hong allow to leave first?

A_{zh}: **无根道人和黑龙使、黄龙使**

Figure 1: Illustration of BiPaR with the monolingual, multilingual and cross-lingual MRC on the dataset.

2017; Liu et al., 2018; Wang et al., 2018) and so on. In order to enable machine to understand texts, large-scale reading comprehension datasets have been developed, such as CNN/Daily Mail (Hermann et al., 2015), SQuAD (Rajpurkar et al., 2016), MS MACRO (Nguyen et al., 2016), hotpotQA (Yang et al., 2018), CoQA (Reddy et al., 2019), etc.

The majority of such datasets, unfortunately, are only for monolingual text understanding. To the best of our knowledge, there is no publicly available bilingual parallel reading comprehension dataset, which is exactly what BiPaR is mainly developed for, as illustrated in Figure 1. BiPaR

Dataset	Language	Parallel	Answer Type	Domain
CNN/DM (Hermann et al., 2015)	EN	×	Fill in entity	News
HLF-RC (Cui et al., 2016)	ZH	×	Fill in word	Fairy/News
SQuAD (Rajpurkar et al., 2016)	EN	×	Span of words	Wikipedia
CMRC2018 (Cui et al., 2018)	ZH	×	Span of words	Wikipedia
MS MACRO (Nguyen et al., 2016)	EN	×	Manual summary	Web doc.
DuReader (He et al., 2018)	ZH	×	Manual summary	Web doc./CQA
BiPaR (this paper)	EN-ZH	✓	Span of words	Novels (Kongfu novels, science fictions, etc.)

Table 1: Comparison of BiPaR with several existing reading comprehension datasets.

produces over 14K Chinese-English parallel questions from nearly 4K bilingual parallel novel passages with consecutive word spans from these passages as answers, following SQuAD (Rajpurkar et al., 2016). Table 1 shows that BiPaR has two significant differences in comparison with existing datasets: (1) each (Passage, Question, Answer) triple is bilingual parallel, and (2) passages and questions are from novels. BiPaR’s bilinguality and novel-based questions provide an interesting corner and unexplored territory for MRC.

With an in-depth analysis on the manually created questions, we observe that answering these novel-style questions requires challenging skills: coreference resolution, multi-sentence reasoning, and understanding of implicit causality, etc. Further monolingual MRC experiments on BiPaR demonstrate that, the English BERT_large model (Devlin et al., 2019) achieves an F1 score of 56.5%, which is 35.4 points behind human performance (91.9%), and the Chinese BERT_Base model achieves an F1 score of 64.1%, which is 28.0 points behind human performance (92.1%), indicating there is a huge gap to be bridged for MRC on novels.

More interestingly, the bilinguality of BiPaR supports multilingual and cross-lingual MRC tasks on this dataset in addition to the traditional monolingual MRC. It is more cost-effective to build a single model that can handle machine reading comprehension on multiple languages, than building multiple MRC systems, one reading system for each language. Different from previous multilingual QA systems that are trained on independently developed datasets of different languages and domains, we are able to train a single multilingual MRC model to do MRC on different languages, by exploring BiPaR that is built parallelly on two different languages with alignments between triples of (Passage, Question, Answer) of the two languages, as shown in Figure 1.

Yet another interesting task that we can do with BiPaR is cross-lingual reading comprehension. We define two types of cross-lingual MRC on BiPaR: (1) using questions in one language to find answers from passages written in another language and (2) finding answers from passages of two different languages for questions in one language. The former is in essence similar to the early cross-lingual question answering (CLQA) (Aceves-Pérez et al., 2008; Peñas et al., 2009; Pérez et al., 2009). The intuitive approaches to CLQA are to translate the questions into the document’s language (Sutcliffe et al., 2005; de Pablo-Sánchez et al., 2005; Aceves-Pérez et al., 2007), which, however, suffers from translation errors. The BiPaR dataset provides a potential opportunity for building cross-lingual MRC that does not rely on machine translation.

To summarize, our contributions are threefold:

- We build the BiPaR, the first publicly available bilingual parallel dataset for MRC. The passages are novel paragraphs, originally written in Chinese or English and then translated into the other language. The questions are manually constructed undergoing a strict quality control procedure.
- We conduct an in-depth analysis on BiPaR, which reveals that MRC on novels is very challenging, requiring skills of coreference resolution, inter-sentential reasoning, implicit causality understanding, etc.
- We build monolingual, multilingual and cross-lingual MRC baseline models on BiPaR and provides baseline results as well as human performance on this dataset.

2 Related Work

MRC Datasets and Models Large-scale cloze-style datasets, such as CNN/Daily Mail (Hermann

et al., 2015), have been automatically developed in the early days of MRC. Several neural network models have been proposed and tested on these datasets, such as ASReader (Kadlec et al., 2016), StanfordAttentiveReader (Chen et al., 2016), AoAReader (Cui et al., 2017), etc. However, Chen et al. (2016) argue that such datasets may be noisy due to the automatic data creation method and co-reference errors. Rajpurkar et al. (2016) propose SQuAD, a dataset created from English Wikipedia, where questions are manually generated by crowdsourced workers, and answers are spans in the Wikipedia passages. Along with the development of this dataset, a variety of neural MRC models have been proposed, such as BiDAF (Seo et al., 2016), R-NET (Wang et al., 2017), ReasonNet (Shen et al., 2017), DCN (Xiong et al., 2016), QANet (Yu et al., 2018), SAN (Liu et al., 2018), etc. Recent years have witnessed substantial progress made on this dataset. However, there are some limitations on SQuAD, which lie in that questions are created based on single passages, that answers are limited to a single span in passages, and that most questions can be answered from a single supporting sentence without requiring multi-sentence reasoning (Chen, 2018). To address these limitations, a number of datasets have been built recently, such as MS MARCO (Nguyen et al., 2016), DuReader (He et al., 2018), TriviaQA (Joshi et al., 2017), RACE (Lai et al., 2017), NarrativeQA (Kocisky et al., 2018), SQuAD2.0 (Rajpurkar et al., 2018), hotpotQA (Yang et al., 2018), CoQA (Reddy et al., 2019), etc. These datasets and models are only for monolingual text understanding. By contrast, BiPaR, following these efforts of creating challenging MRC datasets, aims at setting up a new benchmark dataset for MRC on novels and bilingual/cross-lingual MRC.

Multilingual MRC and Datasets Previous studies on multilingual MRC are very limited. Asai et al. (2018) propose a multilingual MRC system by translating the target language into a pivot language via runtime machine translation. They still rely on SQuAD to train the MRC model of the pivot language. No multilingual MRC dataset is created except that a Japanese and French test set is created by manually translating the test set of SQuAD into the two languages.

Multilingual/Cross-lingual QA Answering questions in multiple languages or retrieving answers from passages that are written in a language different from questions is an important capability for QA systems. To achieve this goal, QA@CLEF¹ has organized a series of public evaluations for multilingual/cross-lingual QA (Magnini et al., 2006). Widely-used approaches to multilingual/cross-lingual QA are to build monolingual QA systems and then adapt them to multilingual/cross-lingual settings via machine translation (Lin et al., 2005). Such QA systems are prone to being affected by machine translation errors. Hence, various techniques have been proposed to reduce the errors of the machine translation module (Sutcliffe et al., 2005; de Pablo-Sánchez et al., 2005; Aceves-Pérez et al., 2007).

3 Dataset Creation

In this section, we elaborate on the three stages of our dataset creation process: collecting bilingual parallel passages, crowdsourcing question-answer pairs on those passages, and constructing multiple answers for the development and test set.

3.1 Bilingual Parallel Passage Collection

We select bilingual parallel passages from six Chinese and English novels with different topics, including Chinese martial arts, science fictions, fantasy literature, etc. These novels are either written in Chinese and translated into English or vice versa. Automatic paragraph alignments between Chinese and English are available for these novels. The number of words in each Chinese passage in parallel passages is limited in the range of [120, 600], to make the passage not too short or too long for crowdsourced workers to construct questions. As some of bilingual parallel passages are very difficult to understand and create parallel questions, we need to consider the appropriateness of these passages. The following rules are used to select bilingual parallel passages:

- The Chinese passage shall not contain poetry, couplets, or classical Chinese words/phrases.
- The passage shall not contain too much dialogue, where speakers are difficult to recognize without global context.
- The Chinese passage shall not contain full-passage description of Kongfu fighting. Such

¹<http://www.clef-initiative.eu/track/qaclef>

descriptions on fighting are hard to translate in a direct way into a target language (i.e., English in this dataset).

- In order to ensure the correct alignments of Chinese and English passages, if an English passage has over 10 words fewer than its Chinese counterpart, such automatically aligned bilingual passages shall not be selected as they are normally not translations of each other.

Following these selection rules, we finally collect 3,667 bilingual parallel passages. Table 2 provides the number of selected passages from each novel.

3.2 Question-Answer Pair Crowdsourcing

We then ask our bilingual crowdsourced workers to create questions and answers on these collected passages. We develop a crowdsourcing annotation system and 150 bilingual workers, 3 bilingual reviewers and 1 expert participate in the data annotation process. The collected bilingual parallel passages are divided into 150 groups and randomly assigned to the bilingual workers. They will create bilingual parallel questions and find corresponding answers to the questions after reading the bilingual parallel passages. Particularly, we encourage workers to create questions according to the following rules:

- For each parallel passage, at least three bilingual question-answer pairs are to be created.
- If the answers in Chinese and English are not parallel (i.e., not translations of each other), the corresponding questions shall be deleted and new questions shall be created.
- Answers have to be consecutive spans in passages.
- If possible, questions of how and why are preferred.
- It is not suggested to directly copy words from passages for creating questions.

In order to guarantee the quality of created question-answer pairs, we use a strict quality control procedure during data annotation. In particular, 30% annotated data from each group are randomly sampled and passed to the three reviewers

Novels	#Parallel Passages	#Avg EN/ZH tokens
I	1,948	251.3/195.9
II	75	223.3/194.6
III	490	206.4/190.1
IV	245	175.1/188.9
V	87	190.7/189.8
VI	822	202.4/212.3
Total	3,667	227.3/198.2

Table 2: Statistics on the selected passages (I: *The Duke of the Mount Deer* / 《鹿鼎记》, II: *Demi-Gods and Semi-Devils* / 《天龙八部》, III: *The Three-Body Problem* / 《三体》, IV: *The Great Gatsby* / 《了不起的盖茨比》, V: *The Old Man and the Sea* / 《老人与海》, VI: *Harry Potter* / 《哈利波特》).

who will review all answers created by the workers and correct answers if they consider they are wrong. Then, 5% of the reviewed data will be further sampled from each reviewer. The sampled data will be checked again by the expert. If the accuracy is lower than 95%, the corresponding workers and reviewers need to revise the answers again. This quality control loop is executed three times.

At last, we collect 14,668 question-answer pairs along with their corresponding passages. We randomly partition the annotated data into a training set (with 11,668 QA pairs), a development set (1,500 QA pairs), and a test set (1,500 QA pairs).

3.3 Multiple Answers Construction

In order to make evaluation more robust, we ask crowdsourced workers to create at least two additional answers for each question in the development and test sets, similar to SQuAD (Rajpurkar et al., 2016). But differently, we make the answers from the crowdsourced workers visible to each other, and encourage them to annotate different but reasonable answers. The reason for creating multiple answers is that we often encounter situations where multiple answers are correct. Consider the following example from SQuAD²:

P: Official corporal punishment, often by caning, remains commonplace in schools in some Asian, African and Caribbean countries. For details of individual countries see School corporal punishment.

Q: What countries is corporal punishment still a normal practice?

The ground truth answers of the question are *some Asian, African and Caribbean countries* /

²<https://rajpurkar.github.io/SQuAD-explorer/explore/1.1/dev/>

Phenomenon	Example	%
Relationship between a question and its passage		
Lexical match	<p>P_{en}: ... Doublet flew into the attack, flailing around her like the wind. She was too small to reach the bodies of her enemies, ...jabbing the Vital Points on the riders' legs.</p> <p>P_{zh}: ... 双儿出手如风，只是敌人骑在马上，她身子又矮，打不到敌人，... 便戳中敌人腿上的穴道。</p> <p>Q_{en}-Q_{zh}: Who did jab the Vital Points on the riders' legs? / 谁戳中敌人腿上的穴道?</p>	49
Paraphrasing	<p>P_{en}: ... Getting very difficult ter find anyone fer the Dark Arts job.</p> <p>P_{zh}: ... 现在找一个黑魔法防御术课老师很困难，人们都不大想干，觉得这工作不吉利。</p> <p>Q_{en}-Q_{zh}: What occupation do people avoid? / 人们不愿从事什么职业?</p>	27
Summary	<p>P_{en}: ... As soon as he opened the door to Ding Yi's brand-new three-bedroom apartment, ... The apartment was unfinished, with only a few pieces of furniture and little decoration, and the huge living room seemed very empty. The most eye-catching object was the pool table in the corner.</p> <p>P_{zh}: ... 推开丁仪那套崭新的三居室的房门，... 看到房间还没怎么装修，也没什么家具和陈设，宽大的客厅显得很空，最显眼的是客厅一角摆放的一张台球桌。</p> <p>Q_{en}-Q_{zh}: What does Ding Yi's three-bedroom look like now? / 丁仪的三居室现在是什么样?</p>	24
Reading comprehension skills required to answer questions		
Coreference resolution	<p>P_{en}: Trying hard to bear all this in mind, Harry took a ... <u>he</u> opened his mouth and immediately swallowed a lot of hot ash. "D-Dia-gon Alley," <u>he</u> coughed.</p> <p>P_{zh}: 哈利拼命把这些都记在心里，... <u>他</u>一张嘴，马上吸了一大口滚烫的烟灰。“对一对角巷。”他咳着说。</p> <p>Q_{en}-Q_{zh}: Where did <u>Harry</u> go? / 哈利去了哪儿?</p>	31
Multi-sentence reasoning	<p>P_{en}: and heard a woman's voice cry out from within it: 'Stop! Lay down your arms! We should all be friends here!' ... The cart stopped in front of them, and out jumped Fang Yi.</p> <p>P_{zh}: 车中一个女子声音叫道：“是自己人，别动手！”... 小车驶别眼前，车中跃出一人，正是方怡。</p> <p>Q_{en}-Q_{zh}: Who did stop the conflict? / 是谁制止了冲突?</p>	32
Implicit causality	<p>P_{en}: Harry, however, was shaken awake several hours earlier than he would have liked by Oliver Wood, Captain of the Gryffindor Quidditch team. "Whassamatter?" said Harry groggily. "Quidditch practice" said Wood.</p> <p>P_{zh}: 哈利一早就被格兰芬多魁地奇队长奥利弗伍德摇醒了，他本来还想再睡几个小时的。“什一什么事？”哈利迷迷糊糊地说。“魁地奇训练”伍德说。</p> <p>Q_{en}-Q_{zh}: Why did Oliver Wood shake Harry awake? / 奥利弗伍德为什么要摇醒哈利?</p>	17

Table 4: Question categories and reading comprehension skills covered in BiPaR. Bold words are answers while underlined words indicate coreference resolution.

sentence. Actually, correctly answering this question requires very complicated reading comprehension skills, such as multi-sentence reasoning, ellipsis/co-reference resolution, etc. We've found 43 examples of this case among the 49 lexical-match samples.

If there is no lexical match between a question and the corresponding passage but we can find a semantically matched sentence to the question from the passage, we regard this case as *paraphrasing*. Such questions account for 27.0% of the questions. Interestingly, we find questions in BiPaR that are not found in other datasets. We refer to these questions as *summary* questions, which account for 24% of the sampled questions. In order to answer these questions, MRC models need to read the entire passage to detect summative statements. Examples of the summary questions are:

- 1) *What is the situation of the Old Majesty?*
- 2) *What features does Oboi's bedroom show?*

In addition, we also analyze the reading comprehension skills required to answer questions. We find that *coreference resolution*, *multi-sentence reasoning* and *implicit causality under-*

standing frequently appear in answering questions in BiPaR. What deserves our special attention here is the implicit causality, which rarely appears in other datasets. For some questions, it is crucial to understand causality that is not represented by explicit expressions such as “why”, “because”, and “the reason for”. As demonstrated in the last example in Table 4, to correctly answer the question, we must understand the implicit causality: *Quidditch practice* → *Harry, however, was shaken awake several hours earlier than he would have liked by Oliver Wood, Captain of the Gryffindor Quidditch team.*

5 MRC Task Formulation on BiPaR

With aligned passage-question-answer triples, we can define three MRC tasks (monolingual, multilingual and cross-lingual) with seven different forms on this dataset, as demonstrated in Figure 1. Since our goal is to provide benchmark results on this new dataset, we either directly train state-of-the-art MRC models on these tasks or use a straightforward way to adapt existing approaches to the defined tasks in this paper. We leave new approaches, especially those for

multilingual and cross-lingual MRC to our future work.

Monolingual MRC: (P_{en}, Q_{en}, A_{en}) or (P_{zh}, Q_{zh}, A_{zh}) . With these two monolingual MRC forms, we can investigate the performance variation of the same MRC model trained on two different languages with equivalent training instances. In our experiments, we directly train off-the-shelf MRC models on the two monolingual tasks to evaluate their performance on Chinese and English.

Multilingual MRC: $(P_{en}, Q_{en}, A_{en}, P_{zh}, Q_{zh}, A_{zh})$. Similar to multilingual neural machine translation (Johnson et al., 2017), we can build a single MRC model to handle MRC of multiple languages on BiPaR. In our benchmark test, we directly mix training instances of the two languages into a single training set. Correspondingly, the two vocabularies are also combined into one vocabulary for both languages. After that, we train MRC models on this language-mixed dataset to endow them with the multilingual comprehension capacity.

Cross-lingual MRC: The first two forms of cross-lingual MRC are (P_{en}, Q_{zh}, A_{en}) or (P_{zh}, Q_{en}, A_{zh}) , in which we use questions in one language to extract answers from passages written in another language. The other two forms are $(P_{en}, P_{zh}, Q_{zh}, A_{zh}, A_{en})$ or $(P_{zh}, P_{en}, Q_{en}, A_{en}, A_{zh})$, in which we use questions written in one language to extract answers from passages written in multiple languages. For the first two forms of cross-lingual MRC, we use Google Translate³ to translate questions into the language of passages, and then treat them as a monolingual MRC task. For the second two forms of cross-lingual MRC, such as $(P_{zh}, P_{en}, Q_{en}, A_{en}, A_{zh})$, we first obtain A_{en} through a monolingual MRC model, then use the word alignment tool fast_align⁴ to obtain the aligned A_{zh} from P_{zh} . Alternative approaches that do not rely on machine translation or word alignments for the cross-lingual MRC tasks are to directly build cross-lingual MRC models on language-mixed training instances constructed from BiPaR or to explore multi-task learning on multiple languages (Dong et al., 2015).

³<https://translate.google.com/>

⁴https://github.com/clab/fast_align

6 Experiments

We carried out experiments with state-of-the-art MRC models on BiPaR to provide machine results for these 7 MRC tasks defined above. We also provide human performance on the monolingual tasks and demonstrate the performance trajectory of human and machine in answering BiPaR questions.

6.1 Evaluation Metric

Like evaluations on other extraction-based datasets, we used EM and F1 to evaluate model accuracy on BiPaR. Particularly, we used the evaluation program of SQuAD1.1 for the English dataset in BiPaR, and the evaluation program⁵ of CMRC2018 for the Chinese dataset in BiPaR.

6.2 Human Performance Evaluation

In order to assess human performance on BiPaR, we hired three other bilingual crowdsourced workers to independently answer questions (both Chinese and English) on the test set which contains three answers per question as described in Section 3.3. We then calculated the average results of the three human workers as the final human performance on this dataset, which are shown in Table 5.

6.3 Baseline models

We adapted the following state-of-the-art models to the dataset and MRC tasks as described in Section 4.

DrQA⁶: DrQA (Chen et al., 2017) is a simple but effective neural network model for reading comprehension.

BERT⁷: BERT (Devlin et al., 2019) is a strong method for pre-training language representations, which obtains the state-of-the-art results on many reading comprehension datasets. We used the multilingual model of BERT trained on multiple languages for the evaluation of our multilingual MRC task.

6.4 Experimental Setup

All the baselines were tested using their default hyper-parameters except BERT. We only changed the batch size to 8 for BERT_base and 6 for BERT_large due to the memory limit of our

⁵<https://github.com/ymcui/cmrc2018>

⁶<https://github.com/hitvoice/DrQA>

⁷<https://github.com/huggingface/pytorch-pretrained-BERT>

	Monolingual MRC		Multilingual MRC		Cross-lingual MRC			
	1	2	3	4	5	6	7	
Development set								
DrQA	29.87/43.47	36.60/52.90	31.47/44.93	36.68/54.03	27.80/38.94	28.47/43.65	8.07/19.80	6.27/19.79
BERT_base	41.67/56.23	52.53/67.65	42.33/55.49	49.00/63.99	36.27/49.98	41.93/55.66	8.00/24.37	8.60/22.51
BERT_large	44.47/58.94	-	-	-	40.40/53.28	-	7.87/24.60	-
Test set								
DrQA	27.00/39.29	37.40/53.11	28.00/42.49	36.60/53.34	21.93/34.45	27.53/41.08	7.00/18.63	4.07/16.64
BERT_base	41.40/55.03	48.87/64.09	38.33/51.20	49.00/64.06	32.80/46.36	39.87/53.10	5.73/21.08	7.67/20.69
BERT_large	42.53/56.48	-	-	-	37.53/51.51	-	5.60/22.29	-
Human	80.50/91.93	81.50/92.12						

Table 5: Results (EM/F1 score) of models and humans on the development and the test data of BiPaR. 1-7 indicate the seven different MRC tasks on BiPaR: (P_{en}, Q_{en}, A_{en}) , (P_{zh}, Q_{zh}, A_{zh}) , $(P_{en}, Q_{en}, A_{en}, P_{zh}, Q_{zh}, A_{zh})$, (P_{en}, Q_{zh}, A_{en}) , (P_{zh}, Q_{en}, A_{zh}) , $(P_{en}, P_{zh}, Q_{zh}, A_{zh}, A_{en})$, $(P_{zh}, P_{en}, Q_{en}, A_{en}, A_{zh})$. For the 6th and 7th task, we mainly explored the word-alignment method. Hence, EM and F1 scores were evaluated on A_{en} or A_{zh} .

GPUs⁸. We used spaCy⁹ to tokenize sentences and generate part-of-speech and named entity tags that were used to train the DrQA model. We downloaded Chinese models¹⁰ for spaCy to preprocess Chinese datasets. The 300-dimensional Glove word embeddings trained from 840B Web crawl data (Pennington et al., 2014) were used as our pre-trained English word embeddings while the 300-dimensional SGNS word embeddings trained from mixed-sources data (Li et al., 2018) as our pre-trained Chinese word embeddings.

6.5 Evaluation Results

Table 5 presents the results of the models on the development and the test data. As BERT_large is currently not available for Chinese, results on the tasks that need to use BERT_large on Chinese are not provided.

BERT vs. Human: In the monolingual MRC task, the English BERT_large model achieves an F1 score of 56.5%, which is 35.4 points behind human performance (91.9%), and the Chinese BERT_base model achieves an F1 score of 64.1%, which is 28.0 points behind human performance (92.1%), indicating that these tasks are difficult to accomplish with current state-of-the-art models. We also tested the English BERT_large model on a subset of the SQuAD data, which contains the same number of training instances as BiPaR. The F1 score is 86.5%, much more higher than that on

⁸The original batch sizes used in BERT_base/BERT_large are 12/24.

⁹<https://spacy.io/>

¹⁰https://github.com/howl-anderson/Chinese_models_for_SpaCy

BiPaR. This further suggests that BiPaR provides a very challenging dataset for MRC.

English vs. Chinese: On the monolingual and cross-lingual tasks, Chinese results are almost better than the English results of the same MRC model. This does not mean that Chinese MRC is easier than English. One possible reason for this may be that novels originally written in Chinese contribute to 68.5% passages of BiPaR. In the future, we plan to make the dataset more balanced between the two languages.

Monolingual vs. Multilingual: For the DrQA model, we observe that the multilingual training significantly improves the performance on English comparing with the monolingual training with only the English dataset. However, we do not observe this trend on BERT. This suggests that more should be explored on the multilingual MRC setting. We believe that BiPaR opens a door to new MRC approaches that are devoted to using a single model to handle MRC on multiple languages.

Monolingual vs. Cross-lingual: The two simple strategies via machine translation and word alignments for the cross-lingual MRC perform very bad on the four forms of cross-lingual tasks compared to the monolingual task. For (P_{en}, Q_{zh}, A_{en}) and (P_{zh}, Q_{en}, A_{zh}) , the major problem is the low-quality translations of questions, especially for questions from martial arts novels. For the other two cross-lingual tasks, word alignment errors directly result in wrong answers found in passages

of the other language.

6.6 Analysis of BERT and Human in Answering Different BiPaR Questions

Table 6 presents a fine-grained comparison analysis of BERT_base and human results on the English and Chinese monolingual task in terms of both answer types and question categories defined in Section 4.3. We observe that humans have absolute advantages over machines in all answer types and reasoning phenomena. However, humans exhibit different capabilities on answering different questions. They perform worse on answering paraphrasing questions, questions with description answers and questions requiring reading comprehension skills of multi-sentence reasoning than answering other questions.

The performance of BERT_base on questions with description answers is much worse than other questions, e.g., questions with person/location answers. As described in Section 4.2, descriptive answers are often complete sentences or summary statements, which are very long and difficult for machines to detect.

In terms of the relationships between questions and passages, it is clearly observed that the BERT MRC model achieves better performance on answering lexical-match questions than answering paraphrasing and summary questions. In BiPaR, we have more than 50% questions that are paraphrasing and summary questions as described in Section 4.3. Answering them requires a deep understanding of both questions and passages.

As described in Section 4.3, BiPaR produces questions that involve higher-order reading comprehension skills, such as co-reference resolution, multi-sentence reasoning, and understanding of implicit causality. It can be seen from Table 6 that the BERT model is worse on multi-sentence reasoning and implicit causality than co-reference resolution.

7 Conclusion and Future Work

In this paper, we have presented the BiPaR, a bilingual parallel machine reading comprehension dataset on novels. From bilingual parallel passages of Chinese and English novels, we manually created diversified parallel questions and answers of different types via crowdsourced workers with a multi-layer quality control system. Although BiPaR is an extractive MRC dataset, in-

Type	BERT_base	Human
Answer Type		
Non-description	61.97/63.22	91.72/95.10
Description	44.77/54.34	91.87/94.72
Question Category		
Lexical Match	58.69/69.01	96.09/96.73
Paraphrasing	54.34/49.84	85.49/91.36
Summary	50.73/52.76	90.03/95.39
Coreference res.	52.35/58.97	91.78/98.11
Multi-sentence rea.	44.91/45.81	87.22/95.31
Implicit causality	39.31/57.63	92.54/91.28

Table 6: Fine-grained results in terms of different answer types and question categories on the monolingual task. The left side of the slash is the F1 score on the English data, while the right is on Chinese. All F1 scores are calculated on the 100 questions as described in Section 4.3.

depth analyses demonstrate that the dataset is very challenging for state-of-the-art MRC models (performing far behind human) as reading comprehension skills of co-reference resolution, inter-sentential reasoning are needed to answer BiPaR questions. We further define seven types of MRC tasks supported by BiPaR and build baseline models of monolingual, multilingual and cross-lingual MRC on BiPaR.

BiPaR can be extended in several ways. First, we would like to create more parallel triples by adding more novels to make instances more balanced between the two languages. Second, we want to create questions with non-extractive answers. Third, we are also interested in adding multi-passage questions or questions based on the entire novels to BiPaR.

Acknowledgments

The present research was supported by the National Natural Science Foundation of China (Grant No. 61622209). We would like to thank the anonymous reviewers for their insightful comments.

References

- Rita M Aceves-Pérez, Manuel Montes-y Gómez, Luis Villaseñor-Pineda, and L Alfonso Ureña-López. 2008. Two approaches for multilingual question answering: Merging passages vs. merging answers. *International Journal of Computational Linguistics & Chinese Language Processing, Volume 13, Number 1, March 2008: Special Issue on Cross-*

- Lingual Information Retrieval and Question Answering*, 13(1):27–40.
- Rita Marina Aceves-Pérez, Manuel Montes-y Gómez, and Luis Villaseñor-Pineda. 2007. Enhancing cross-language question answering by combining multiple question translations. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 485–493. Springer.
- Akari Asai, Akiko Eriguchi, Kazuma Hashimoto, and Yoshimasa Tsuruoka. 2018. Multilingual extractive reading comprehension by runtime machine translation. *arXiv preprint arXiv:1809.03275*.
- Danqi Chen. 2018. *Neural Reading Comprehension and Beyond*. Ph.D. thesis, Stanford University.
- 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)*, volume 1, pages 2358–2367.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading Wikipedia to answer open-domain questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics.
- Yiming Cui, Zhipeng Chen, Si Wei, Shijin Wang, Ting Liu, and Guoping Hu. 2017. Attention-over-attention neural networks for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 593–602.
- Yiming Cui, Ting Liu, Zhipeng Chen, Shijin Wang, and Guoping Hu. 2016. Consensus attention-based neural networks for chinese reading comprehension. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1777–1786.
- Yiming Cui, Ting Liu, Li Xiao, Zhipeng Chen, Wentao Ma, Wanxiang Che, Shijin Wang, and Guoping Hu. 2018. A span-extraction dataset for chinese machine reading comprehension. *arXiv preprint arXiv:1810.07366*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. Multi-task learning for multiple language translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1723–1732.
- Wei He, Kai Liu, Jing Liu, Yajuan Lyu, Shiqi Zhao, Xinyan Xiao, Yuan Liu, Yizhong Wang, Hua Wu, Qiaoqiao She, et al. 2018. DuReader: a chinese machine reading comprehension dataset from real-world applications. In *Proceedings of the Workshop on Machine Reading for Question Answering*, pages 37–46.
- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 1*, pages 1693–1701. MIT Press.
- Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611.
- Rudolf Kadlec, Martin Schmid, Ondřej Bajgar, and Jan Kleindienst. 2016. Text understanding with the attention sum reader network. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 908–918.
- Tomas Kocisky, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gabor Melis, and Edward Grefenstette. 2018. The narrativeqa reading comprehension challenge. *Transactions of the Association for Computational Linguistics*, 6:317–328.
- 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*, pages 785–794.
- Shen Li, Zhe Zhao, Renfen Hu, Wensi Li, Tao Liu, and Xiaoyong Du. 2018. Analogical reasoning on chinese morphological and semantic relations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 138–143.

- Frank Lin, Hideki Shima, Mengqiu Wang, and Teruko Mitamura. 2005. CMU javelin system for NTCIR5 CLQA1. In *NTCIR*.
- Xiaodong Liu, Yelong Shen, Kevin Duh, and Jianfeng Gao. 2018. Stochastic answer networks for machine reading comprehension. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1694–1704.
- Bernardo Magnini, Danilo Giampiccolo, Lili Aunimo, Christelle Ayache, Petya Osenova, Anselmo Peñas, Maarten de Rijke, Bogdan Sacaleanu, Diana Santos, and Richard Sutcliffe. 2006. The multilingual question answering track at CLEF. In *quot; In Nicoletta Calzolari; Khalid Choukri; Aldo Gangemi; Bente Maegaard; Joseph Mariani; Jan Odjik; Daniel Tapias (ed) Proceedings of the 5th International Conference on Language Resources and Evaluation (LREC'2006)(Genoa Italy 22-28 May 2006)*.
- 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. *choice*, 2640:660.
- César de Pablo-Sánchez, Ana González-Ledesma, José Luis Martínez-Fernández, José María Guirao, Paloma Martínez, and Antonio Moreno-Sandoval. 2005. Miracle’s 2005 approach to cross-lingual question answering. In *CLEF (Working Notes)*.
- Anselmo Peñas, Pamela Forner, Richard Sutcliffe, Álvaro Rodrigo, Corina Forăscu, Iñaki Alegria, Danilo Giampiccolo, Nicolas Moreau, and Petya Osenova. 2009. Overview of respubliqa 2009: question answering evaluation over european legislation. In *Workshop of the Cross-Language Evaluation Forum for European Languages*, pages 174–196. Springer.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Joaquín Pérez, Guillermo Garrido, A Rodrigo, Lourdes Araujo, and Anselmo Peñas. 2009. Information retrieval baselines for the respubliqa task. In *Working Notes for the CLEF 2009 Workshop, Corfu, Greece*.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2019. CoQA: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*.
- Yelong Shen, Po-Sen Huang, Jianfeng Gao, and Weizhu Chen. 2017. Reasonet: Learning to stop reading in machine comprehension. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1047–1055. ACM.
- Richard FE Sutcliffe, Michael Mulcahy, Igal Gabbay, Aoife O’ Gorman, and Darina Slattery. 2005. Cross-language french-english question answering using the dlt system at clef 2005. In *Workshop of the Cross-Language Evaluation Forum for European Languages*, pages 502–509. Springer.
- Wenhui Wang, Nan Yang, Furu Wei, Baobao Chang, and Ming Zhou. 2017. Gated self-matching networks for reading comprehension and question answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 189–198.
- 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. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1918–1927.
- Caiming Xiong, Victor Zhong, and Richard Socher. 2016. Dynamic coattention networks for question answering. *arXiv preprint arXiv:1611.01604*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380.
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V Le. 2018. Qanet: Combining local convolution with global self-attention for reading comprehension. *arXiv preprint arXiv:1804.09541*.