

ChID: A Large-scale Chinese IDiom Dataset for Cloze Test

Chujie Zheng^{1,2}, Minlie Huang^{1*}, Aixin Sun³

¹Institute for Artificial Intelligence, Dept. of Computer Science, Tsinghua University

¹Beijing National Research Center for Information Science and Technology, China

²Dept. of Physics, Tsinghua University, Beijing 100084, China

³School of Computer Science and Engineering, Nanyang Technological University, Singapore

zcj16@mails.tsinghua.edu.cn, aihuang@tsinghua.edu.cn

axsun@ntu.edu.sg

Abstract

Cloze-style reading comprehension in Chinese is still limited due to the lack of various corpora. In this paper we propose a large-scale Chinese cloze test dataset **ChID**, which studies the comprehension of idiom, a unique language phenomenon in Chinese. In this corpus, the idioms in a passage are replaced by blank symbols and the correct answer needs to be chosen from well-designed candidate idioms. We carefully study how the design of candidate idioms and the representation of idioms affect the performance of state-of-the-art models. Results show that the machine accuracy is substantially worse than that of human, indicating a large space for further research.

1 Introduction

Machine reading comprehension aims to assess the ability to comprehend natural language and answer questions from a given document or passage. As a classical method of assessing language proficiency (Fotos, 1991; Jonz, 1991; Tremblay, 2011), cloze test (Taylor, 1953) has been widely employed due to its simplicity in form. Recently, a number of datasets for cloze test have been proposed for different languages. For instance, CNN/Daily Mail (Hermann et al., 2015) provides a benchmark for machine comprehension of English text, while the People Daily and Children’s Fairy Tale dataset (Cui et al., 2016) and CMRC-2017 (Cui et al., 2018) pioneer explorations in Chinese language.

In this paper we explore idiom comprehension (Wray, 2002; Jackendoff and Jackendoff, 2002; Cacciari and Tabossi, 2014; Jiang et al., 2018) in cloze test. Idiom, which is called “成语” (chengyu) in Chinese, is an interesting linguistic phenomena in Chinese language, and this work

*Corresponding author: Minlie Huang.

Idiom	亡羊补牢
Literal	To mend the fence after sheep are lost
Meaning	Never be late to try

Table 1: An example of metaphor in idiom. The sense of “亡羊补牢” should be inferred figuratively but not represented literally using the meanings of the four constituent characters.

Idioms	侃侃而谈	口若悬河
Meanings	To speak with fervour and assurance	Fluently and eloquently
Common	Speak much and long	
Difference	Describe the demeanor of the speaker	Describe the eloquence

Table 2: An example of near-synonyms in idiom, where idioms share similar meanings but are different in language usage.

is in parallel to several datasets (Hill et al., 2016; Xie et al., 2018) that have considered different language phenomena in English. Compared to other types of words, many idioms are unique for their non-compositionality and metaphorical meaning (see an example in Table 1). This feature requires a good representation of idiom. Meanwhile, the characteristic of near-synonym, i.e., words that have similar but not identical meanings (see an example in Table 2), may challenge a machine to choose an accurate idiom in a given context. Due to the fact that idioms are widely used in daily communication and in various literary genres, it is a new challenge to assess the ability of understanding and representing idioms in Chinese reading comprehension.

To this end, we propose **ChID**, a large-scale Chinese **ID**iom dataset for cloze test. ChID con-

Dataset	Lang.	Extractive	Option	Answer Type	Domain	Size
CNN/Daily Mail	EN	Yes	No	Named entities	News	1.38M
Children’s Book Test	EN	Yes	Yes	Multiple types	Children’s Books	678K
Who-did-What	EN	Yes	Yes	Named entities	News	206K
LAMBADA	EN	Partial	No	Multiple types	Novels	10K
Story Cloze Test	EN	No	Yes	Single sentence	Life Stories	102K
CLOTH	EN	No	Yes	Multiple types	Examinations	99K
People Daily & Children’s Fairy Tale	CN	Yes	No	Nouns, Named entities	News, Children’s Stories	880K
CMRC-2017	CN	Yes	No	Nouns, Named entities	Children’s Stories	359K
ChID (this work)	CN	No	Yes	Chinese idioms	News, Novels, Essays	729K

Table 3: Comparison of ChID with other cloze-style reading comprehension datasets. *Extractive* denotes whether the answer is extracted directly from the given context. *Option* denotes whether candidate choices are provided. In the *Answer Type* column, the answers of all datasets except Story Cloze Test are single words. *Size* denotes the total number of queries or blanks in the dataset.

tains 581K passages and 729K blanks, and covers multiple domains. In ChID, the idioms in a passage were replaced with blank symbols. For each blank, a list of candidate idioms including the golden idiom are provided as choice. As the difficulty level of cloze test depends on candidate choices, we investigate several strategies of selecting candidate idioms. We evaluate several state-of-the-art models on the proposed corpus with different representations of idioms. Results show that machine performs much worse than human, which indicates a large room for further research.

Our contributions are summarized as follows:

- We propose a new dataset, ChID, for cloze-style reading comprehension in Chinese language. ChID contains 581K passages and 729K blanks from three domains (news, novels, and essays).
- We conduct extensive experiments on the design of candidate idioms and the idiom representation methods, and compare state-of-the-art models. Results show that the performance of these models is substantially worse than that of human.
- ChID provides a benchmark to evaluate the ability of understanding idioms, a unique yet common language phenomenon in Chinese. To our knowledge, this is the first work where this linguistic phenomenon is studied in the form of machine reading comprehension.

2 Related Work

Recently, machine reading comprehension has been advanced by many corpora with various task settings. For instance, CNN/Daily Mail (Hermann et al., 2015) collects news articles and uses the cloze test (Taylor, 1953) to assess the ability of reading comprehension in English. RACE (Lai et al., 2017) and CLOTH (Xie et al., 2018) are constructed from questions in examinations designed for secondary and high school students. A number of question-answer datasets (Rajpurkar et al., 2016; Reddy et al., 2018) are also proposed and there are many other large-scale datasets (Nguyen et al., 2016; He et al., 2018). These corpora inspire various neural models (Chen et al., 2016; Cui et al., 2016; Seo et al., 2017; Dhingra et al., 2017; Cui et al., 2017). In Table 3, we present a survey on existing cloze-style reading comprehension datasets.

As the earliest cloze-style dataset for machine reading comprehension, CNN/Daily Mail (Hermann et al., 2015) has a very large scale. It collects news articles paired with a number of bullet points, which summarise key aspects of an article. Based on the fact that these summary points are abstractive and do not simply copy sentences from a news article, the corpus is constructed by transforming these bullet points into cloze-style questions, i.e., replacing one entity with a placeholder. Children’s Book Test (CBT) (Hill et al., 2016) also provides a benchmark for machine reading comprehension,

while the key differences from CNN/Daily Mail include: a list of candidate choices is provided for each query, and more types of words are removed, including named entities, (common) nouns, verbs and prepositions. Who-did-What (Onishi et al., 2016) collects its corpus from news and provides options for questions similar to CBT. Each question is formed from two independent articles: an article is treated as context to be read and a separate article on the same event is used to form the query. LAMBADA (Paperno et al., 2016) removed the last word from a given passage and evaluates the ability of word prediction. By contrast, the Story Cloze Test dataset (Mostafazadeh et al., 2017) evaluates the ability of story understanding and script learning, where the task requires to select or generate a reasonable sentence to complete the story context.

To the best of our knowledge, People Daily (PD) and Children’s Fairy Tale (CFT) (Cui et al., 2016) and CMRC-2017 (Cui et al., 2018) are the only two existing cloze-style datasets for Chinese reading comprehension. Similar to CNN/Daily Mail and CBT, PD & CFT and CMRC-2017 replaced a word (usually a noun or named entity) in the document with a blank placeholder and treated the sentence containing this word as a query. PD collects data from news while CFT and CMRC-2017 are from children’s reading materials.

In most datasets, the answer can be directly found from context. CLOTH (Xie et al., 2018) has a similar setting to ChID, where the answer should be selected from given choices. However, CLOTH is collected from English examinations for secondary/high school students, whose size is limited because documents, blanks, and options are all manually created.

3 Chinese Idioms

Idiom is a common language phenomenon and usually called “成语” (chengyu) in Chinese. Thanks to its conciseness in form and expressiveness in meaning, idiom is widely used in daily communication and in various text genres. The main challenges for machine reading comprehension with idiom lie in: idiom representation which represents the meaning of an idiom, and thorough discrimination among the near-synonyms of an idiom.

3.1 Idiom Representation

Many idioms are non-compositional and have metaphorical meanings (see an example in Table 1), which has also made idiom translation a challenging problem and attracted considerable research attentions (Anastasiou, 2010; Salton et al., 2014; Cap et al., 2015; Shao et al., 2017). The meaning of such idioms is generally different from the literal meanings of the constituent characters. Such idioms are usually originated from ancient cultural stories, but the meaning is reserved along the long history of language use. For instance, “塞翁失马” has a metaphorical meaning, which is derived from this story:

Near China’s northern borders lived an old man who bred many horses. One day, one of his horses, for no reason at all, escaped into the territory of the northern tribes. Everyone commiserated with him. “Perhaps this will soon turn out to be a blessing,” said the old man. After a few months, his horse came back, and brought back a fine horse from the north.

So the idiom “塞翁失马” usually refers to a blessing in disguise. Thus comprehending and representing an idiom may require the access to the corresponding cultural history. In addition, due to the polysemy of a single character, even those compositional idioms are likely to have ambiguity, which also makes idiom representation a challenging problem.

3.2 Near-synonyms

It is common that an idiom has near-synonyms. These idioms may be confused in language use due to their similar but not identical meanings¹ (see an example in Table 2). To discriminate those near-synonyms, machine is required to figure out their subtle differences in usage, which is also challenging.

To verify the near-synonym phenomena, we conducted a user study. Based on the idiom vocabulary we collected (see Section 4.1), we manually evaluated the number of near-synonyms per idiom. We randomly sampled 200 idioms. For each idiom, we picked up the 20 most similar idioms whose embedding similarity score to the input idiom is less than some threshold. According

¹Idioms with identical meanings can be interchangeably used in any context.

Passage & Blanks		可是有一个时期大家#idiom-0#，不大敢露面，只有她一个人倚在阳台上看排队的兵走过。
		However, there was a period when everyone #idiom-0# and was scared to show up. Only she leaned on the balcony and watched the soldiers passing by.
#idiom-0# options	Correct	深居简出 Be unwilling to contact people
	Similar	销声匿迹 To disappear from the scene
		离群索居 To stay away from the crowd and live alone
		安分守己 To know one’s place
Random	一帆风顺 To proceed smoothly without a hitch	
	文不对题 Be irrelevant to the subject	
	万里长征 A long and arduous undertaking	

Figure 1: An example in ChID. Each data contains a given passage with several blanks that replace the original idioms (in this example, there is only one blank). For each blank, several options are provided. Among the list of candidate choices, there is one golden answer, three similar idioms and another three random ones.

$\geq K$ NEARs	# Idioms	Proportion
1	179	89.5%
2	131	65.5%
3	80	40.0%
4	40	23.0%

Table 4: Annotation result of near-synonyms. It shows the number of idioms in the 200 sampled idioms that have at least K near-synonyms, for $K = 1, 2, 3, 4$. Fleiss’ kappa is 0.479, indicating moderate agreement.

to the similarity annotation result of Section 4.3 and Table 6, we set this threshold to 0.85. Then we hired four annotators to label these 4,000 idiom pairs in terms of whether a pair is near-synonyms or not. All the annotators have good command of Chinese.

The evaluation result is shown in Table 4. Note that for each idiom, we rounded down the mean of the numbers of near-synonyms labeled by the four annotators. We estimate that about 90% idioms have at least 1 near-synonym. About 23% of the idioms have 4 or more near-synonyms. Fleiss’ kappa (Fleiss, 1971) for measuring inter-annotator agreement is 0.479, indicating moderate agreement (within [0.4, 0.6]). This evaluation result strongly supports our claim that near-synonyms are very common among Chinese idioms.

4 Dataset Collection

Figure 1 presents an example in ChID. In each sample, idioms in a passage are replaced by blank

Level	Freq.	Num.	Prop.
Very Low	[20, 50)	832	21.6%
Low	[50, 100)	742	19.3%
Medium	[100, 200)	822	21.4%
High	[200, 400)	746	19.4%
Very High	[400, 534]	706	18.3%
Total	[20, 534]	3,848	100.0%

Table 5: Idiom frequency statistics in the whole corpus. The minimum and the maximum are 20 and 534 respectively.

symbols, and each blank is provided with several candidate idioms including the golden idiom. The task is to select the golden answer from the candidate choices given the context. Note that the answer is usually not occurring in the context in our setting, which is different from most existing cloze test corpora.

In the following subsections, we will explain the three steps in data collection: (1) Constructing the idiom vocabulary; (2) Extracting passages within a proper length; (3) Designing candidate choices.

4.1 Vocabulary Construction

We collected the idiom vocabulary from Chinese idioms Daquan², which contains over 23K idiom entries. Since vast majority of idioms consist of 4 characters, we only retained idioms with 4 characters in our vocabulary. In order to facilitate the design of candidate choices, we removed

²<http://www.guoxue.com/chengyu/CYML.htm>

those idioms that do not have a pre-trained embedding using the large-scale open-source corpus provided by Song et al. (2018), where approximately 40% idioms were filtered out. We normalized synonyms with only slight morphological variation. Idioms that share the same explanation and meaning, but only differ in one character or the order of characters, are treated as the same idiom. This can be done with the Chinese idiom dictionary because some idioms are marked with: “又作” (also written as), “犹” (like), “同” (the same as), “见” (also see). Such idioms in the passages are all replaced by their re-normalized ones.

We then counted the frequency of each idiom in the corpus, and removed those idioms that appear less than 20 times. Finally, the idiom vocabulary has 3,848 entries in total, and their frequency statistics on the whole corpus is shown in Table 5. The minimum and the maximum idiom frequencies are 20 and 534 respectively. We simply divide the idiom frequency into five intervals: very low (from 20 to 50), low (from 50 to 100), medium (from 100 to 200), high (from 200 to 400) and very high (higher than 400). The proportions of idioms in the frequency intervals are almost uniformly distributed.

4.2 Passage Extraction

To make the topic and domain more diversified, we collected passages from novel and essay on the Internet, and the news articles provided by Sun et al. (2016)³. Since some documents may be very long, we took a paragraph as the basic unit. Each idiom except those in double quotation marks⁴ is replaced with a blank symbol. A paragraph that is shorter than 100 characters is merged with the next paragraph to ensure that the context are sufficient for answer selection. Those passages that are longer than 600 characters are abandoned.

It is worth noting that if some idiom has a much higher word frequency than others, models may tend to bias answer selection to those more frequent idioms. In order to make frequent and infrequent idioms more balanced, we removed some passages which only contain high frequency idioms.

³The news articles are extracted from the THUCNews dataset for Chinese text classification, which can be downloaded from <http://thuctc.thunlp.org/>.

⁴Because words in a quotation mark are usually entities or other content that can not be inferred from the context.

Similarity	SYN	NEAR	OTHER	κ
[0.85, 0.90)	83.2%	16.8%	0.0%	.642
[0.80, 0.85)	53.6%	42.8%	3.6%	.447
[0.75, 0.80)	29.2%	53.6%	17.2%	.485
[0.70, 0.75)	12.0%	57.2%	30.8%	.496
[0.65, 0.70)	0.4%	52.8%	46.8%	.466
[0.60, 0.65)	0.0%	34.0%	66.0%	.528
[0.55, 0.60)	0.0%	10.4%	89.6%	.657
[0.50, 0.55)	0.0%	6.0%	94.0%	.787

Table 6: Annotation result of embedding similarity. The three labels are: **SYN** (synonym), **NEAR** (near-synonym), **OTHER**. κ is the Fleiss’ kappa value.

4.3 Candidate Choice Selection

The semantic relevance between two idioms can be measured by the cosine similarity of their embeddings (Mikolov et al., 2013), which helps us to design candidate choices. However, idioms that are similar in embedding may or may not be synonyms or near-synonyms. We thus manually evaluated the correlation between embedding similarity and idiom synonymity. We split the embedding similarity from 0.9 to 0.5 into 8 intervals. Within each interval, 200 pairs of idioms are sampled. We used three labels to measure the relevance between two idioms: **SYN** (synonym, the two idioms are identical in meaning and can be interchangeably used), **NEAR** (near-synonym, have close or similar meanings but can not be used interchangeably), **OTHER** (irrelevant or opposite in meaning). We hired five annotators to label these samples.

As shown in Table 6, when the similarity score is larger than 0.75, there is a large proportion of idioms pairs that have the same meaning; when the score is between 0.65 and 0.80, there is a large probability that the two idioms are near-synonyms. For those pairs with high (larger than 0.85) or low (smaller than 0.60) similarity, annotators tend to reach substantial agreement⁵ according to Fleiss’ kappa, while we have moderate agreement between the similarity interval [0.65, 0.85].

The above annotation results inspire us to design proper candidate choices for each blank in a passage. First of all, we excluded those idioms that have a similarity score higher than 0.7 to the golden answer. This avoids to include synonyms of the golden answer in the candidate

⁵Substantial agreement corresponds to kappa within [0.6, 0.8].

	In-domain				Out-of-domain	Total
	Train	Dev	Test	Total	Out	
# Passages	520,711	20,000	20,000	560,711	20,096	580,807
Avg. # tokens per passage	99	99	99	99	127	100
# Distinct idioms covered	3,848	3,458	3,502	3,848	3,626	3,848
Avg. idiom frequency	168.6	7.2	7.1	181.6	8.3	189.6
Total # blanks	648,920	24,822	24,948	698,690	30,023	728,713
Avg. # blanks per passage	1.25	1.24	1.25	1.25	1.49	1.25
Single-blank prop.	80.4%	80.7%	80.8%	80.5%	64.7%	79.9%
Multi-blank prop.	19.6%	19.3%	19.2%	19.5%	35.3%	20.1%

Table 7: ChID dataset statistics. The out-of-domain data have longer passages (127 vs. 99) and more blanks per passage (1.49 vs. 1.25) than the in-domain data.

Level	Freq.	In	Out
Very Low	[20, 50)	3.5%	8.2%
Low	[50, 100)	7.2%	12.0%
Medium	[100, 200)	16.0%	19.7%
High	[200, 400)	28.8%	28.7%
Very High	[400, 534]	44.5%	31.4%
Total	[20, 534]	100.0%	100.0%

Table 8: Comparison on idiom frequency distribution between the in-domain and out-of-domain data.

choice. Then, we picked up top 10 similar idioms among the remaining idioms, and randomly chose three idioms as candidate choice. Note that the three idioms have a large probability of being near-synonyms of the golden answer, which affects the difficulty level of the cloze test to some degree. We further randomly sampled another three idioms from the remaining idioms that do not include the top 10 similar idioms. In this manner, the list of candidate choices consists of three parts: the correct answer, three similar idioms, and three other randomly sampled ones, as shown in Figure 1.

4.4 Corpus Statistics

The detailed statistics of ChID is shown in Table 7. News and novels are treated as in-domain data, which are divided into the training set **Train**, the development set **Dev**, and the test set **Test**. Essays are reserved for out-of-domain test **Out** to assess the generalization ability of cloze test models. The in-domain data cover 3,848 Chinese idioms, while **Dev/Test/Out** respectively cover

3,458/3,502/3,626 idioms.

There are some differences between in-domain and out-of-domain data. Firstly, the average length of passages in the in-domain data is nearly 100 words, while Out-of-domain data have longer passages (127 words). The average number of blanks per passage is also different (1.25 vs. 1.49). Secondly, the idiom distributions are different. As shown in Table 8, compared to the in-domain data, low-frequency idioms occupy a higher proportion of all the idiom occurrences in the out-of-domain data (8.2% vs. 3.5% for very low frequency interval and 12.0% vs. 7.2% for low frequency interval) while the high-frequency idioms occur less frequently (31.4% vs. 44.5%). These differences make the out-of-domain test set more challenging.

5 Experiment

5.1 Models

In order to evaluate how well the state-of-the-art models can comprehend Chinese language with idiom, we tested the following models:

Language Model (LM): We trained a bidirectional LSTM (Hochreiter and Schmidhuber, 1997) to obtain the hidden state at the blank (h_b), and use the hidden state to score candidate choices:

$$\vec{h}_b = \overrightarrow{\text{LSTM}}(w_{1:b}), \quad \overleftarrow{h}_b = \overleftarrow{\text{LSTM}}(w_{b:|p|}) \quad (1)$$

$$h_b = \vec{h}_b \oplus \overleftarrow{h}_b \quad (2)$$

$$\alpha_i = \text{softmax}_i(h_b^T c_i) \quad (3)$$

where $|p|$ denotes the length of passage p , $w_{1:b}, w_{b:|p|}$ denote the words in the given context before or after the blank respectively, \oplus denotes concatenation, and c_i denotes the embedding of

each candidate idiom. Then, the option that has the highest α_i is chosen as the answer.

Attentive Reader (AR) (Hermann et al., 2015): The bidirectional LSTM model is augmented with the attention mechanism (Bahdanau et al., 2015). The hidden state at blank h_b is used as the query to attentively read the context as follows:

$$\mathbf{m}_t = \tanh(\mathbf{W}_{hm}\mathbf{h}_t + \mathbf{W}_{bm}\mathbf{h}_b) \quad (4)$$

$$s_t = \text{softmax}_t(\mathbf{w}_{ms}^T \mathbf{m}_t), \mathbf{r} = \sum_{t=1}^{|\mathbf{p}|} s_t \mathbf{h}_t \quad (5)$$

where \mathbf{W}_{hm} , \mathbf{W}_{bm} , \mathbf{w}_{ms} are all parameters. Then, the attention vector \mathbf{r} and the blank vector h_b are used to score each candidate choice:

$$\mathbf{g} = \tanh(\mathbf{W}_{rg}\mathbf{r} + \mathbf{W}_{bg}\mathbf{h}_b) \quad (6)$$

$$\alpha_i = \text{softmax}_i(\mathbf{g}^T \mathbf{c}_i) \quad (7)$$

where \mathbf{W}_{rg} , \mathbf{W}_{bg} are also parameters.

Stanford Attentive Reader (SAR) (Chen et al., 2016): Compared to AR, SAR applies a bilinear matrix \mathbf{W}_s to compute attention weights instead of using a tanh layer. The weighted contextual vector \mathbf{o} is used for scoring candidates:

$$s_t = \text{softmax}_t(\mathbf{h}_b^T \mathbf{W}_s \mathbf{h}_t), \mathbf{o} = \sum_{t=1}^{|\mathbf{p}|} s_t \mathbf{h}_t \quad (8)$$

$$\alpha_i = \text{softmax}_i(\mathbf{o}^T \mathbf{c}_i) \quad (9)$$

5.2 Implementation Details

All the models were implemented with Tensorflow (Abadi et al., 2016). We employed the Jieba Chinese word segmenter⁶ to tokenize passages. We set the vocabulary size to 100K and used the 200-dimensional word embeddings initialized by Song et al. (2018). Those word embeddings that were not matched in Song et al. (2018) were initialized from a uniform distribution between (-0.1, 0.1). We applied a dropout rate of 0.5 on word embeddings. The number of hidden units of RNN cells were all set to 100. The cross entropy cost function is used to compute the training loss. ADAM (Kingma and Ba, 2015) was used to optimize all the models with the initial learning rate to 0.001 and the gradient was clipped when the norm of the gradient was larger than 5. We set the batch size to 32. The training was stopped when the accuracy on **Dev** did not improve within an epoch.

⁶<https://github.com/fxsjy/jieba>

	Dev	Test	Ran	Sim	Out
Human	-	87.1	97.6	82.2	86.2
κ	-	.794	.953	.791	.769
LM	71.8	71.5	80.7	65.6	61.5
AR	72.7	72.4	82.0	66.2	62.9
SAR	71.7	71.5	80.0	64.9	61.7

Table 9: Performance of human and models. κ indicates Fleiss’ kappa. The overall best results are shown in bold, and AR performs significantly better than LM and SAR (sign test, p -value < 0.05).

5.3 Option Settings

To evaluate how the method of candidate choice design will impact the performance, we prepared two additional test sets: **Ran** and **Sim**, both of which have the same passages with **Test**, but candidate choices are designed differently. In **Ran**, all the candidate choices are sampled from the idioms that are not similar to the golden answer. Instead, in **Sim**, all the candidates are sampled from top 10 similar idioms. Therefore, **Sim** is more challenging than **Ran** as the former has more distracting options. Note that each blank has seven choices including the golden answer.

5.4 Results

To explore the ceiling of model performance, we also conducted **Human Evaluation**. We sampled 200 passages respectively from the aforementioned test sets: **Test**, **Ran**, **Sim** and **Out**. We then hired three annotators to complete the 800 cloze tests. These three annotators are first-year or second-year university students and all have very good command of Chinese language. The average accuracy of the annotators and the corresponding Fleiss’ kappa are reported as the final performance.

The experiment results are shown in Table 9. We analyzed the results from the following perspectives:

Option Setting: The setting of similar options is much harder than that of random options. Firstly, we noted that both human and models achieve worse performance on **Test** than on **Ran**, while the accuracy on **Sim** is even lower than **Test**, which indicates that including more similar candidate idioms makes the task more difficult. Secondly, the inter-annotator agreement on **Ran** (Fleiss’ kappa=0.953) is much higher than those

	Dev	Test	Ran	Sim	Out
Idiom Embedding					
LM	71.8	71.5	80.7	65.6	61.5
AR	72.7	72.4	82.0	66.2	62.9
SAR	71.7	71.5	80.0	64.9	61.7
Average Character Embedding					
LM	63.1	63.0	70.6	57.2	53.2
AR	64.5	63.8	72.5	57.8	53.5
SAR	62.9	62.5	71.4	56.9	52.0
Average Character Embedding + MLP					
LM	68.6	68.4	77.4	62.0	57.8
AR	67.1	66.5	75.6	60.7	56.4
SAR	67.0	66.8	75.6	60.8	55.6

Table 10: Performance comparison using different idiom representations.

on other test sets which include similar options. This implies that similar options also make manual annotation harder.

Human vs. Models: Firstly, human performance is substantially better than model performance on all the test sets. The smallest gap between human and machine is 14.6 (on **Test**) and the largest gap is 23.3 (on **Out**). Secondly, humans perform very closely on **Test** and **Out** (87.1 vs. 86.2), however, the models perform much better on **Test** than on **Out** (72.4 vs. 62.9). This observation implies that human has a strong ability to generalize to out-of-domain data while the models cannot generalize well to **Out** which contains more low-frequency idioms.

Model Comparison: AR outperforms all other models significantly. The reason for this may be due to the fact: AR firstly uses the blank representation (h_b) to make an attentive read of the context (see Eq. 4 and 5), and the blank vector is used again with the attentive vector (r) to score a candidate choice. In this manner, the context is attentively used and the blank vector is used twice.

5.5 Comparison on Idiom Representation

In previous experiments, an idiom was treated as a token, and its representation are obtained through pretraining on a large corpus (Song et al., 2018). In this section, we explored another two methods for idiom representation, and evaluated the performance with different idiom representations. One

method simply uses the average embedding of 4 constituent characters as the representation of an idiom. This method mimics to understand idioms purely based on its literal meanings. The other is to apply an MLP (Multi-Layer Perceptron, Bishop et al., 1995; Fine, 1999) which is fed with the concatenation of 4 character embeddings, and the output vector is used to represent an idiom. This method also applies a composition assumption: the representation of an idiom is a composite function of its constituent words. Note that the input to the MLP is an 800-dimension vector, and the MLP has a hidden layer of 400 units and uses tanh as the activation function. The final output of MLP is a 200-dimension vector.

Table 10 shows the performance comparison using three methods for idiom representation. We can observe remarkable drops from *idiom embedding* to *average character embedding + MLP* and to *average character embedding* for all the models, where all the differences are significant (sign test, p -value < 0.01). The results indicate that the other two idiom representation methods are worse than treating an idiom as an independent semantic unit. This study also implies that idiom representation is a key factor for the success of Chinese reading comprehension with idiom. In other words, a good cloze test model should have not only a proper model structure, but also a good method to represent idioms.

6 Conclusion

In this paper, we propose a large-scale Chinese cloze dataset (ChID) which contains 581K passages and 729K queries from news, novels, and essays, covering 3,848 Chinese idioms. The corpus provides a benchmark to evaluate the ability of Chinese cloze test with idiom. Firstly, we analyze how the embedding similarity correlates with synonymity and near-synonymity of Chinese idiom, and find that the difficulty level of Chinese cloze test with idiom correlates positively with the method of choosing candidate choices. Secondly, we find that idiom representation is a key factor to the success of reading comprehension models in this task due to the common non-compositionality and metaphorical meaning of Chinese idiom. Thirdly, we evaluate three state-of-the-art cloze test models on this corpus, and observe that existing model performance is still much worse than human performance. All these

findings indicate that the corpus may be a proper benchmark for Chinese cloze test and worth further research⁷.

Acknowledgments

This work was jointly supported by the National Science Foundation of China (Grant No.61876096), and the National Key R&D Program of China (Grant No. 2018YFC0830200).

References

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pages 265–283.
- Dimitra Anastasiou. 2010. *Idiom treatment experiments in machine translation*. Cambridge Scholars Publishing.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations*.
- Christopher M Bishop et al. 1995. *Neural networks for pattern recognition*. Oxford university press.
- Cristina Cacciari and Patrizia Tabossi. 2014. *Idioms: Processing, structure, and interpretation*. Psychology Press.
- Fabienne Cap, Manju Nirmal, Marion Weller, and Sabine Schulte Im Walde. 2015. How to account for idiomatic german support verb constructions in statistical machine translation. In *Proceedings of the 11th Workshop on Multiword Expressions*, pages 19–28.
- 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.
- 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)*, volume 1, pages 593–602.
- Yiming Cui, Ting Liu, Zhipeng Chen, Wentao Ma, Shijin Wang, and Guoping Hu. 2018. Dataset for the first evaluation on chinese machine reading comprehension. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*.
- 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.
- Bhuwan Dhingra, Hanxiao Liu, Zhilin Yang, William Cohen, and Ruslan Salakhutdinov. 2017. Gated-attention readers for text comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1832–1846.
- Terrence L Fine. 1999. *Feedforward Neural Network Methodology*. Springer Science & Business Media.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Sandra S Fotos. 1991. The cloze test as an integrative measure of efl proficiency: A substitute for essays on college entrance examinations? *Language learning*, 41(3):313–336.
- 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, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems*, pages 1693–1701.
- Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2016. The goldilocks principle: Reading children’s books with explicit memory representations. In *International Conference on Learning Representations*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Ray Jackendoff and Ray S Jackendoff. 2002. *Foundations of language: Brain, meaning, grammar, evolution*. Oxford University Press, USA.
- Zhiying Jiang, Boliang Zhang, Lifu Huang, and Heng Ji. 2018. Chengyu cloze test. In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 154–158.
- Jon Jonz. 1991. Cloze item types and second language comprehension. *Language testing*, 8(1):1–22.

⁷The corpus **Train** is available at <https://github.com/zhengcjl/ChID-Dataset>.

- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- 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.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *International Conference on Learning Representations Workshop*.
- Nasrin Mostafazadeh, Michael Roth, Annie Louis, Nathanael Chambers, and James Allen. 2017. Ls-dsem 2017 shared task: The story cloze test. In *Proceedings of the 2nd Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics*, pages 46–51.
- 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. *arXiv preprint arXiv:1611.09268*.
- Takeshi Onishi, Hai Wang, Mohit Bansal, Kevin Gimpel, and David McAllester. 2016. Who did what: A large-scale person-centered cloze dataset. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2230–2235.
- Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernandez. 2016. The lambada dataset: Word prediction requiring a broad discourse context. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1525–1534.
- 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. 2018. Coqa: A conversational question answering challenge. *arXiv preprint arXiv:1808.07042*.
- Giancarlo D. Salton, Robert J. Ross, and John D. Kelleher. 2014. Evaluation of a substitution method for idiom transformation in statistical machine translation. In *MWE@EACL*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *International Conference on Learning Representations*.
- Yutong Shao, Rico Sennrich, Bonnie L. Webber, and Federico Fancellu. 2017. Evaluating machine translation performance on chinese idioms with a blacklist method. In *Language Resources and Evaluation Conference*.
- Yan Song, Shuming Shi, Jing Li, and Haisong Zhang. 2018. Directional skip-gram: Explicitly distinguishing left and right context for word embeddings. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, volume 2, pages 175–180.
- M Sun, J Li, Z Guo, Z Yu, Y Zheng, X Si, and Z Liu. 2016. Thuctc: an efficient chinese text classifier. *GitHub Repository*.
- Wilson L Taylor. 1953. “cloze procedure”: A new tool for measuring readability. *Journalism Bulletin*, 30(4):415–433.
- Annie Tremblay. 2011. Proficiency assessment standards in second language acquisition research: “clozing” the gap. *Studies in Second Language Acquisition*, 33(3):339–372.
- Alison Wray. 2002. Formulaic language and the lexicon.
- Qizhe Xie, Guokun Lai, Zihang Dai, and Eduard Hovy. 2018. Large-scale cloze test dataset created by teachers. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2344–2356.