
Teaching Pretrained Models with Commonsense Reasoning: A Preliminary KB-Based Approach

Shiyang Li* Jianshu Chen† and Dian Yu†

Abstract

Recently, pretrained language models (e.g., BERT) have achieved great success on many downstream natural language understanding tasks and exhibit a certain level of commonsense reasoning ability. However, their performance on commonsense tasks is still far from that of humans. As a preliminary attempt, we propose a simple yet effective method to teach pretrained models with commonsense reasoning by leveraging the *structured knowledge* in ConceptNet, the largest commonsense knowledge base (KB). Specifically, the structured knowledge in KB allows us to construct various logical forms, and then generate multiple-choice questions requiring commonsense logical reasoning. Experimental results demonstrate that, when refined on these training examples, the pretrained models consistently improve their performance on tasks that require commonsense reasoning, especially in the few-shot learning setting. Besides, we also perform analysis to understand which logical relations are more relevant to commonsense reasoning.

1 Introduction

Recently, pretrained language models [4, 11, 19] have achieved great successes on various natural language understanding tasks, and they are also believed to master a certain level of commonsense reasoning abilities [9, 10, 17]. Equipping machines with commonsense reasoning ability has been seen as one of the key milestones of artificial general intelligence [3]. However, the commonsense reasoning ability of these state-of-the-art pretrained models is still far away from that of humans [8, 16]. One probable reason is that these models are learned from massive amounts of *unstructured* texts with various language model (LM) objectives (e.g., masked language model [4]). That is, the commonsense reasoning capability is never explicitly taught to the pretrained models, but is implicitly acquired through modeling input texts via LM objectives. In this paper, we focus on how to *explicitly* teach the pretrained models the commonsense reasoning ability.

There are several challenges in explicitly injecting commonsense reasoning capability into pretrained models. First, it is generally hard to exploit direct supervision signals for commonsense reasoning from unstructured texts, and it is also expensive, if ever possible, to create a large amount of human-labeled data for learning the commonsense reasoning ability. Second, the pretrained models do not have explicit symbolic reasoning operations; instead, the reasoning is performed implicitly through the neural network operations such as self-attention, and any knowledge relevant to reasoning is stored in the network weights. Note that the weights are only learned to fit certain input-output pairs, where the inputs to the model are natural language sentences, and the outputs are certain items to predict (e.g., masked tokens, next sentence indicator). That is, any reasoning ability has to be acquired implicitly by processing unstructured input texts during pretraining, and it is more difficult to directly supervise the reasoning path for a pretrained model.

*Department of Computer Science, University of California, Santa Barbara, CA 93106. Email: shiyangli@ucsb.edu. The work was done during an internship at Tencent AI Lab, Bellevue, WA.

†Tencent AI Lab, Bellevue, WA, USA. Emails: jianshuchen@tencent.com and yudian@tencent.com.




Figure 1: The generation of logical forms and multiple-choice questions in our proposed approach. The yellow and the red circles in the Venn diagram represent the sets \mathcal{R}_1 and \mathcal{R}_2 , respectively.

To address these challenges, we propose a simple yet effective method to teach pretrained models with explicit commonsense reasoning abilities. The key idea is to exploit the *structured* knowledge in commonsense knowledge bases (e.g., ConceptNet [13]) to generate multiple-choice questions that require commonsense reasoning. Specifically, we sample subgraphs from KB to generate various logical forms and then use text templates to generate natural language questions and candidate answers. As a result, we automatically generate a large-scale multiple-choice question answering dataset with 167 million questions that ask about specific logical relations between different entities/concepts. These questions will be used as the additional training data to further *refine* the pretrained models, which force them to learn the commonsense reasoning ability in order to answer correctly. These training inputs are already in the natural language form, which is consistent with the input of pretrained models. Therefore, it allows the model to *continually* adjust its pretrained weights so that it can master more commonsense reasoning abilities; it naturally combines the power of pretrained weights from unstructured texts and the new information from structured knowledge in KB. Our experimental results show that the proposed approach consistently outperforms the baselines on commonsense reasoning tasks, especially in few-shot learning settings. In addition, we examine which logical relations are more “commonsense” and find that only a few simple ones are most relevant. This work is as a preliminary attempt to integrate structured commonsense knowledge into pretrained models with promising results. As we shall see, the *structured* knowledge in KB allows us to *systematically* construct the logical relations that we want to teach the models. We hope that our work could inspire more research towards combining structured knowledge and pretrained models.

2 The Proposed Approach

The key idea of our method is to generate multiple-choice questions from different subgraphs in KB, and then we use the generated data to further *refine* the pretrained models. The overall idea of the data generation process is shown in Figure 1, which consists of (i) generating different logical forms from a sampled subgraph in KB, (ii) generating multiple-choice questions in natural language form.

2.1 Generating multiple-choice questions as the refinement data

Generating logical forms We first sample a subgraph from KB that is in the following form:

$$(A \xrightarrow{R_1} B \xrightarrow{R_2} C) \quad (1)$$

where A , B , and C are three different entities in the KB, and R_1 and R_2 represent two different relations in the KB. For each of the above subgraph, we will construct a multiple-choice question regarding the entity B in the following manner. First, introduce the following two sets: $\mathcal{R}_1 = \{X \in \Omega : A \xrightarrow{R_1} X\}$, $\mathcal{R}_2 = \{X \in \Omega : X \xrightarrow{R_2} C\}$, where Ω denotes the entire entity set. Note that the set \mathcal{R}_1 represents the set of all (tail) entities that have relation R_1 with A , and \mathcal{R}_2 represents the set of all (head) entities that have relation R_2 with entity C . We use the two circles in the Venn diagram (Fig. 1) to represent these two sets, respectively. Note from Fig. 1 that the entire space could be partitioned into four subsets, denoted as: $\mathcal{S}_1 = \mathcal{R}_1 \cap \mathcal{R}_2^c$, $\mathcal{S}_2 = \mathcal{R}_1 \cap \mathcal{R}_2$, $\mathcal{S}_3 = \mathcal{R}_1^c \cap \mathcal{R}_2$, $\mathcal{S}_4 = \mathcal{R}_1^c \cap \mathcal{R}_2^c$. Each subset represents a certain logical relation. For example, the subset $\mathcal{S}_2 = \mathcal{R}_1 \cap \mathcal{R}_2$ means all the entities that have relation R_1 with A and have relation R_2 with C . Using these four subsets, we could compose questions that ask about all different logical relations from the subgraph in (1). To see this, note that we could compose a set by either choosing or not choosing each subset \mathcal{S}_i , which leads to a total of $2^4 = 16$ subsets. Among them, two trivial cases are excluded: the all-chosen case (full set) and the all-not-chosen set (empty set). Therefore, there are a total of 14 different logical relations about (1) that we could ask (see Appendix B for all the 14 logical forms). To have a more concrete

example, consider the composed subset $\mathcal{S}_1 \cup \mathcal{S}_3$, then we are examining the logical relation:

$$\left((A \xrightarrow{R_1} ?) \vee (? \xrightarrow{R_2} C) \right) \wedge \neg \left((A \xrightarrow{R_1} ?) \wedge (? \xrightarrow{R_2} C) \right) \quad (2)$$

where \wedge and \vee denotes logical AND and logical OR, respectively, and \neg denotes logical negation (NOT). This approach allows us to *systematically* generate all different types of logical relations pertaining to each sampled subgraph from the KB, which even covers questions about a single relation. For example, the logical form corresponding to $\mathcal{S}_1 \cup \mathcal{S}_2$ is “ $A \xrightarrow{R_1} ?$ ”, and the logical form corresponding to $\mathcal{S}_2 \cup \mathcal{S}_3$ is given by “ $? \xrightarrow{R_2} C$ ”, which ask the tail entity and head entity, respectively.

Generating multiple-choice questions Now that once we have a logical form in the form of (2), we can generate natural language questions that ask about this particular logical relation. We achieve this by using text templates. Specifically, we first create two different types of mapping, namely, *affirmative mapping* and *negative mapping*. The affirmative mapping is used to generate sentences with affirmative questions, while the negative mapping is used for generating negative ones. Consider the following specific example of a logical form (also shown in Figure 1):

$$(\text{alone} \xrightarrow{\text{Antonym}} ?) \wedge (? \xrightarrow{\text{CapableOf}} \text{sing in church})$$

where the correct answer for the missing entity is `people`. In the above logical form, the relation `CapableOf` will be mapped into “is capable of” using affirmative mapping. On the other hand, when there is a negation \neg before the relation `CapableOf`, it will be mapped into “is not capable of” using a negative mapping. These obtained strings from relations will be put together with the head entities and the tail entities to generate sentences as natural as possible by using a set of simple heuristic rules. For example, the above logical relation will be mapped into the following natural language sentence: “*which of the following is an antonym of alone and meanwhile is capable of sing in church?*” In Appendix B, we give examples of the possibly generated questions for all the 14 logical relations.

Generating candidate answers The correct answer is obtained from the particular logical form that we want to examine. For example, if we want to generate a question regarding the logical form (2), the set of correct answer is given by $\mathcal{S}_1 \cup \mathcal{S}_3$. On the other hand, for the wrong candidate answers, we will examine three different sampling strategies. The first approach is to *random sample* from the all the other entities in KB [15]. The second one is the *nearest sampling*, which chooses the entity from $\{X \in \Omega : A \xrightarrow{R_1} X, \forall R \neq R_1\}$. The third sampling method is *uniform sampling*: it firstly chooses wrong subset uniformly from $\mathcal{S}_1, \dots, \mathcal{S}_4$ and then samples an entity from the selected subset.

2.2 Refinement: teaching the pretrained models with commonsense reasoning

To teach the pretrained models with commonsense reasoning, we further train the pretrained models on the generated multiple-choice questions to predict the correct answer, which becomes a multi-class classification problem. Afterwards, the model is finetuned on different downstream tasks. We name this step as *refinement* to distinguish it from the *pretraining* and the *finetuning* stages.

3 Experiments

In this section, we examine the performance of the proposed method on different tasks and perform analysis on which logical relations are more “commonsense”. First, we briefly describe the experimental setting, and more details could be found in Appendix A. We first preprocess ConceptNet and keep 3,098,816 English-only triples. Then, we perform search on these triples and obtain a total of 167,395,947 subgraphs that are in the form of (1). These subgraphs would lead to over 167 million multiple-choice questions for further refining the pretrained models. We use the *uniform sampling* method to generate the wrong candidate answers unless otherwise stated. To evaluate the performance, we finetune the refined models on three downstream tasks that require strong commonsense reasoning: *CommonsenseQA* [16], *CosmosQA* [5], and *DREAM* [14] (see Appendix A for the descriptions).

Few-Shot Learning Performance In Table 1, we show the few-shot learning performance of our proposed method on CommonsenseQA. We consider three different types of pretrained models: BERT, GPT, and XLNet. We refine these models on our generated multiple-choice questions and then finetune them on CommonsenseQA. We compare the results to the corresponding models without the refinement process (i.e., directly finetuning on CommonsenseQA). Our method has significantly

better few-shot learning performance with as large as 18% absolute improvement, meaning that the refinement process effectively teaches a pretrained model commonsense reasoning even with a few finetuning samples. With full finetuning data, our method also achieves 2% gain. The above results are obtained using the base models of BERT/XLNet. Additional experimental results in Appendix C show that the same performance gain could carry over to their corresponding large models.

Table 1: The few-shot learning performance in accuracy (%) on the CommonsenseQA development set. Shot percentages are listed in the parentheses. “model + refine” denotes our method. All results are averaged over five independent runs, with standard deviations listed inside the parentheses.

Shot	100 (1.0%)	200 (2.1%)	400 (4.1%)	800 (8.2%)	1600 (16.4%)	3200 (32.9%)	9741 (100%)
BERT [4]	34.94(1.97)	38.41(2.16)	41.73(2.02)	45.44(1.16)	47.53(1.15)	51.84(0.62)	57.33(1.03)
BERT + refine	42.54(1.27)	44.93 (1.69)	47.03 (0.27)	50.58 (0.64)	53.43 (0.85)	54.86 (0.75)	59.28 (0.43)
GPT [11]	27.90(1.30)	28.34(1.39)	30.20(1.99)	33.96(2.53)	38.54(1.55)	45.45(0.65)	50.75(1.08)
GPT + refine	37.69(0.27)	38.56(0.37)	40.49(0.50)	42.49(0.44)	44.24(0.58)	46.50(0.35)	51.52(0.62)
XLNet [19]	25.04(0.67)	27.44(1.08)	29.80(2.17)	34.27(0.89)	38.67(1.30)	47.14(1.25)	57.25(1.14)
XLNet + refine	43.60 (0.15)	43.67(0.24)	45.81(0.23)	47.24(0.47)	50.60(0.53)	53.41(0.32)	59.31 (0.44)

Candidate answer sampling strategies In Table 2, we show the results of different strategies to sample the wrong candidate answers on different datasets. We find that our method is relatively insensitive to different sampling strategies, and the performance varies slightly over different datasets.

Table 2: Performance of different strategies to sample wrong candidate answers.

Dataset	Candidate answer selection	DREAM-dev	DREAM-test	CommonsenseQA-dev	CosmosQA-dev
BERT		62.06(0.75)	61.98(0.79)	57.33(1.03)	58.34(0.82)
BERT + refine	random sampling	63.49 (0.35)	62.89 (0.36)	58.51(0.87)	58.66(0.26)
BERT + refine	nearest sampling	63.02(0.23)	62.56(0.94)	58.97 (0.76)	59.14 (0.89)

Which logical relations are more “commonsense”? To partially answer this question, we refine the pretrained BERT-base model on different subsets of logical relations from all the 14 logical forms in Appendix B and report the results in Table 3. We observe that relatively simple logical relations (#1, #2, #5) (i.e., simple logical AND and single relation reasoning) are more relevant to commonsense; refining on just three of them achieves almost full performance (i.e., BERT + refine (all)). On the other hand, the logical forms (#4, #7, #9), which require more logical compositions and negations, are less commonsense; refining on them does not improve much over the baseline BERT model. This is consistent with our intuition that commonsense should be something relatively straightforward.

Table 3: The relevance of different logical relations to commonsense. “BERT + refine (1,2,5)” means the pretrained BERT model is refined on the logical forms #1, #2, and #5 defined in Appendix B.

Method	BERT	BERT + refine (1,2,5)	BERT + refine (2,4,5)	BERT + refine (4,7,9)	BERT + refine (all)
CommonsenseQA	57.33(1.03)	59.10(0.43)	58.18(0.41)	56.42(0.82)	59.28(0.43)
CosmosQA	58.34(0.82)	59.28(0.81)	59.76(0.75)	58.86(0.74)	58.91(0.44)

4 Related Work

Structured knowledge [2, 13] has been explored for question answering and reading comprehension. Most existing methods [1, 6, 7, 12, 14, 18, 21] only exploit triples relevant to questions. [15] fine-tune models on questions constructed by predicting the head or tail mention in a triple. Very recently, [20] propose to align triples to Wikipedia sentences to form natural language questions.

5 Conclusion

In this paper, we propose a simple yet effective method to use structured knowledge (i.e., ConceptNet) to enhance the commonsense reasoning abilities of pretrained language models. The structured knowledge in KB allows us to construct various logical forms, and then generate multiple-choice questions that require commonsense logical reasoning. Experimental results demonstrate that, when refined on these training examples, these models consistently improve their performance on three datasets that require commonsense knowledge, especially in the few-shot learning setting. In the future, we are interested in designing methods to generate more diverse natural language questions instead of relying on patterns and evaluating on other recent models like RoBERTa.

References

- [1] B. Bi, C. Wu, M. Yan, W. Wang, J. Xia, and C. Li. Incorporating external knowledge into machine reading for generative question answering. In *Proceedings of the EMNLP*, 2019.
- [2] K. Bollacker, C. Evans, P. Paritosh, T. Sturge, and J. Taylor. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the ACM SIGMOD*, 2008.
- [3] E. Davis and G. Marcus. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Commun. ACM*, 58(9):92–103, 2015.
- [4] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [5] L. Huang, R. L. Bras, C. Bhagavatula, and Y. Choi. Cosmos qa: Machine reading comprehension with contextual commonsense reasoning. In *Proceedings of the EMNLP*, 2019.
- [6] L. Laugier, A. Wang, C.-S. Foo, T. Guenais, and V. Chandrasekhar. Encoding knowledge graph with graph cnn for question answering. *Proceedings of the ICLR*, 2019.
- [7] Y. Li and P. Clark. Answering elementary science questions by constructing coherent scenes using background knowledge. In *Proceedings of the EMNLP*, 2015.
- [8] B. Y. Lin, X. Chen, J. Chen, and X. Ren. Kagnet: Knowledge-aware graph networks for commonsense reasoning. *arXiv preprint arXiv:1909.02151*, 2019.
- [9] Y. Liu, F. He, H. Zhang, G. Rao, Z. Feng, and Y. Zhou. How well do machines perform on iq tests: a comparison study on a large-scale dataset. *Proceedings of the IJCAI*, 2019.
- [10] S. Ostermann, M. Roth, A. Modi, S. Thater, and M. Pinkal. SemEval-2018 Task 11: Machine comprehension using commonsense knowledge. In *Proceedings of the SemEval*, 2018.
- [11] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever. Improving language understanding by generative pre-training. In *Preprint*, 2018.
- [12] M. Sachan, A. Dubey, and E. P. Xing. Science question answering using instructional materials. In *Proceedings of the ACL*, 2016.
- [13] R. Speer, J. Chin, and C. Havasi. ConceptNet 5.5: An Open Multilingual Graph of General Knowledge. In *Proceedings of the AAAI*, 2017.
- [14] K. Sun, D. Yu, J. Chen, D. Yu, Y. Choi, and C. Cardie. DREAM: A challenge data set and models for dialogue-based reading comprehension. *Transactions of the Association for Computational Linguistics*, 2019.
- [15] K. Sun, D. Yu, D. Yu, and C. Cardie. Probing prior knowledge needed in challenging chinese machine reading comprehension. *arXiv preprint arXiv:1904.09679*, 2019.
- [16] A. Talmor, J. Herzig, N. Lourie, and J. Berant. Commonsenseqa: A question answering challenge targeting commonsense knowledge. *arXiv preprint arXiv:1811.00937*, 2018.
- [17] T. H. Trinh and Q. V. Le. A simple method for commonsense reasoning. *arXiv preprint arXiv:1806.02847*, 2018.
- [18] L. Wang, M. Sun, W. Zhao, K. Shen, and J. Liu. Yuanfudao at SemEval-2018 Task 11: Three-way attention and relational knowledge for commonsense machine comprehension. In *Proceedings of the SemEval*, 2018.
- [19] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. Salakhutdinov, and Q. V. Le. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*, 2019.
- [20] Z.-X. Ye, Q. Chen, W. Wang, and Z.-H. Ling. Align, mask and select: A simple method for incorporating commonsense knowledge into language representation models. *arXiv*, 2019.
- [21] W. Zhong, D. Tang, N. Duan, M. Zhou, J. Wang, and J. Yin. Improving question answering by commonsense-based pre-training. *arXiv preprint arXiv:1809.03568*, 2018.

Supplementary Material

A Experimental Details

In this section, we describe more details of the experimental settings and the choice of hyper-parameters.

A.1 Experiment details of the refinement process

Handling invalid logical forms. We find that some subgraphs (1) sampled from KB could not generate all the 14 logical forms in Appendix B. For example, if \mathcal{S}_1 is an empty set for a specific subgraph, logical form #0 is invalid. In our implementation, we create a specific 14-dimension 0/1-mask vector for each subgraph to indicate which logical forms are valid for sampling.

Efficiency considerations. In our implementation, we use the `torch.utils.data.Dataset` class in PyTorch to generate the training data for the refinement process on-the-fly. We observe that calculating \mathcal{S}_4 is relatively time-consuming because we have to remove all the elements in \mathcal{S}_1 , \mathcal{S}_2 , and \mathcal{S}_3 from the total set for each sampled subgraph. This can be a bottleneck for the dataloader and will finally reduce the overall GPU utilization. To address this issue, we approximate \mathcal{S}_4 by the total set in our experiments, which is an efficient and relatively accurate approximation. Note that since the number of elements in \mathcal{S}_1 , \mathcal{S}_2 , and \mathcal{S}_3 is much smaller than that in the total set, the chance of sampling an element from \mathcal{S}_1 , \mathcal{S}_2 , and \mathcal{S}_3 is extremely small. Therefore, this could be an efficient and good approximation to sampling from \mathcal{S}_4 .

Hyper-parameters for the refinement process. When we refine BERT, GPT, and XLNet, we only train our models for one epoch. This is because we find that training too many iterations on our generated multiple-choice question answering dataset may make the model forget the pretrained language modeling capability and eventually hurt performance. We set the maximum sequence length to be 40 during refinement as it covers most of the input texts for all three pretrained models, and we set the optimizers and the learning rates to be the same as their default values. The learning rates are set to be 2×10^{-5} , 6.25×10^{-5} , and 2×10^{-5} for BERT, GPT, and XLNet, respectively. We do not tune their hyper-parameters (e.g., learning rate) due to limited resources. Note that for GPT, language model coefficient is set to be 0 during refining since we argue that the texts in our template datasets may not be as natural as the ones used for pretraining.

Experimental setting for Table 3. For BERT + refine (all), we sample over all valid logical forms according to a uniform distribution. For BERT + refine (1,2,5), BERT + refine (2,4,5), and BERT + refine (4,7,9), logical forms are uniformly sampled over (#1, #2, #5), (#2, #4, #5), and (#4, #7, #9), respectively. The training procedures follow the same hyper-parameters described above. For the finetuning process in all experiments, we train a model five times and report their mean values and standard deviations.

A.2 Description of the downstream tasks

CommonsenseQA dataset consists of 12,247 questions with one correct answer and four wrong answers. This dataset has two kinds of splits, namely token split and random split. Our experiments are conducted on the official random split. For few-shot learning experiments, we allow our models to train more epochs to make sure that they converge. Specifically, for $shot \in \{100, 200, 400, 800, 1600, 3200\}$, we train our models with $epoch \in \{100, 50, 25, 12, 8, 8\}$, respectively and keep other settings fixed. For training on the whole dataset, we follow similar settings of officially released code.³ For a fair comparison between baselines and our refining methods, we keep their epochs, batch sizes, and other settings the same. The only differences are parameters where baselines utilize officially pretrained models, and ours use checkpoints during the proposed refining processes.

Cosmos QA is a large-scale dataset of 35.6K problems that require commonsense-based reading comprehension, formulated as multiple-choice questions. It focuses on reading between the lines over

³<https://github.com/jonathanherzig/commonsenseqa/tree/master/bert>.

a diverse collection of people’s everyday narratives, asking questions concerning on the likely causes or effects of events that require reasoning beyond the exact text spans in the context.⁴ Therefore, it is an appropriate dataset for testing commonsense reasoning of models. We finetune baselines and our proposed methods for four epochs with learning rate 2e-5 and batch size of 36. We evaluate models on the development set in every epoch and report the best performance for each experiment.

DREAM is a multiple-choice dialogue-based machine reading comprehension examination dataset. It focuses on in-depth multi-turn multi-party dialogue understanding. Answering 34% of these questions needs commonsense reasoning. We adapt the officially released BERT-based source code on DREAM and choose the same setting as the repository.⁵ Similar to *CosmosQA*, we evaluate development set in every epoch and report the best performance and its corresponding test set accuracy for each experiment.

B All Logical Forms with Example Multiple-Choice Questions

In this appendix, we show all the 14 logical relations that could be sampled from a particular triple pair, and the examples for the corresponding generated multiple-choice questions. Specifically, we consider the following example of triple pair:

$$(\text{arise} \xrightarrow{\text{Antonym}} \text{sit}, \text{sit} \xrightarrow{\text{RelatedTo}} \text{sit up})$$

Then, all the 14 logical forms and the corresponding example questions are given below, where the correct answer is highlighted in red and bolded:

- logical form #0: \mathcal{S}_1

$$(A \xrightarrow{R_1} ?) \wedge \neg (? \xrightarrow{R_2} C) \quad (3)$$

Q: *which of the following is an antonym of arise and meanwhile is not related to sit up ?*

A: **set**

B: *fancifying*

C: *storing space shuttle*

- logical form #1: \mathcal{S}_2

$$(A \xrightarrow{R_1} ?) \wedge (? \xrightarrow{R_2} C) \quad (4)$$

Q: *which of the following is an antonym of arise and meanwhile is related to sit up ?*

A: **sit**

B: *sitting up*

C: *stand up*

- logical form #2: $\mathcal{S}_1 \cup \mathcal{S}_2$

$$(A \xrightarrow{R_1} ?) \quad (5)$$

Q: *which of the following is an antonym of arise ?*

A: *promegapoietin*

B: *sleigher*

C: **set**

- logical form #3: \mathcal{S}_3

$$\neg (A \xrightarrow{R_1} ?) \wedge (? \xrightarrow{R_2} C) \quad (6)$$

⁴<https://leaderboard.allenai.org/cosmosqa/submissions/public>.

⁵https://github.com/nlpdata/mrc_bert_baseline.

Q: which of the following is not an antonym of arise and meanwhile is related to sit up ?
 A: craftist
 B: queer anarchism
 C: **stand up**

- logical form #4: $\mathcal{S}_1 \cup \mathcal{S}_3$

$$\left((A \xrightarrow{R_1} ?) \vee (? \xrightarrow{R_2} C) \right) \wedge \neg \left((A \xrightarrow{R_1} ?) \wedge (? \xrightarrow{R_2} C) \right) \quad (7)$$

Q: which of the following is an antonym of arise or is related to sit up, but not both of them ?
 A: sit down
 B: make refreshing dessert
 C: **lower**

- logical form #5: $\mathcal{S}_2 \cup \mathcal{S}_3$

$$(? \xrightarrow{R_2} C) \quad (8)$$

Q: which of the following is related to sit up ?
 A: lay
 B: **sitting up**
 C: descend

- logical form #6: $\mathcal{S}_1 \cup \mathcal{S}_2 \cup \mathcal{S}_3$

$$(A \xrightarrow{R_1} ?) \vee (? \xrightarrow{R_2} C) \quad (9)$$

Q: which of the following is an antonym of arise or is related to sit up ?
 A: marksberrys
 B: **sit down**
 C: previsive

- logical form #7: \mathcal{S}_4

$$\neg (A \xrightarrow{R_1} ?) \wedge \neg (? \xrightarrow{R_2} C) \quad (10)$$

Q: which of the following is not an antonym of arise and is not related to sit up ?
 A: crunch
 B: **millikin**
 C: sit

- logical form #8: $\mathcal{S}_1 \cup \mathcal{S}_4$

$$\neg (? \xrightarrow{R_2} C) \quad (11)$$

Q: which of the following is not related to sit up ?
 A: sit down
 B: **simpliciter**
 C: crunch

- logical form #9: $\mathcal{S}_2 \cup \mathcal{S}_4$

$$\left((A \xrightarrow{R_1} ?) \wedge (? \xrightarrow{R_2} C) \right) \vee \left(\neg (A \xrightarrow{R_1} ?) \wedge \neg (? \xrightarrow{R_2} C) \right) \quad (12)$$

Q: which of the following is an antonym of arise and is related to sit up, or neither of them ?
 A: fall down
 B: **cremators**
 C: lower

- logical form #10: $\mathcal{S}_1 \cup \mathcal{S}_2 \cup \mathcal{S}_4$

$$(A \xrightarrow{R_1} ?) \vee \neg(? \xrightarrow{R_2} C) \quad (13)$$

Q: *which of the following is an antonym of arise or is not related to sit up ?*

A: *sitting up*

B: ***sit down***

C: *stand up*

- logical form #11: $\mathcal{S}_3 \cup \mathcal{S}_4$

$$\neg(A \xrightarrow{R_1} ?) \quad (14)$$

Q: *which of the following is not an antonym of arise ?*

A: *lay down*

B: ***free criminals***

C: *abed*

- logical form #12: $\mathcal{S}_1 \cup \mathcal{S}_3 \cup \mathcal{S}_4$

$$\neg(A \xrightarrow{R_1} ?) \vee \neg(? \xrightarrow{R_2} C) \quad (15)$$

Q: *which of the following is not an antonym of arise or is not related to sit up ?*

A: *sit down*

B: *sit down*

C: ***snub line***

- logical form #13: $\mathcal{S}_2 \cup \mathcal{S}_3 \cup \mathcal{S}_4$

$$\neg(A \xrightarrow{R_1} ?) \vee (? \xrightarrow{R_2} C) \quad (16)$$

Q: *which of the following is not an antonym of arise or is related to sit up ?*

A: *lower*

B: *fall*

C: ***sit down***

C Additional Experiments

In this section, we list additional experiments with the BERT-large model. All results are averaged over five independent runs, with standard deviations listed inside the parentheses.

Table 4: Performance of our proposed method on BERT-large vs. on BERT-base, evaluated on the DREAM dataset. The performance metric is prediction accuracy (%). The standard deviations are listed inside the parenthesis.

Data	Candidate answer selection	DREAM-dev	DREAM-test
BERT-base	random sampling	62.06(0.75)	61.98(0.79)
BERT-base + refine	random sampling	63.49(0.35)	62.89(0.36)
BERT-large	random sampling	65.62(0.62)	66.07(1.00)
BERT-large + refine	random sampling	67.11(0.48)	67.36(0.63)