

GO BEYOND PLAIN FINE-TUNING: IMPROVING PRETRAINED MODELS FOR SOCIAL COMMONSENSE

Ting-Yun Chang^{1*}, Yang Liu², Karthik Gopalakrishnan², Behnam Hedayatnia²,
Pei Zhou³, Dilek Hakkani-Tür²

¹ Academia Sinica, Taiwan; ² Amazon, Alexa AI, USA; ³ University of Southern California

ABSTRACT

Pretrained language models have demonstrated outstanding performance in many NLP tasks recently. However, their social intelligence, which requires commonsense reasoning about the current situation and mental states of others, is still developing. Towards improving language models’ social intelligence, in this study we focus on the Social IQA dataset, a task requiring social and emotional commonsense reasoning. Building on top of the pretrained RoBERTa and GPT2 models, we propose several architecture variations and extensions, as well as leveraging external commonsense corpora, to optimize the model for Social IQA. Our proposed system achieves competitive results as those top-ranking models on the leaderboard. This work demonstrates the strengths of pretrained language models, and provides viable ways to improve their performance for a particular task.

Index Terms— Commonsense Reasoning, Social IQA, Pretrained Language Models

1. INTRODUCTION

Recently, there has been emerging research interest in developing machine common sense. Several commonsense datasets have been proposed [1, 2, 3, 4, 5]. In particular, Social IQA was introduced in [4], which is a multiple-choice QA dataset probing machine’s emotional and social commonsense reasoning in a variety of everyday situations. Afterward, different efforts have been put forth for this challenging task. One line of work introduces additional knowledge bases, providing the model with more information underlying the given text [6, 7, 8, 9]. Another line generalizes existing well-performed models, focusing on zero-shot setting [6, 10]. Generally, large pretrained language models are leveraged in previous work, such as BERT [11], RoBERTa [12], GPT and GPT2 [13, 14], and T5 [15]. For example, [16] recently reformulates Social IQA as a generation task and achieves remarkable performance by fine-tuning T5 [15].

A standard practice in NLP recently is to simply fine-tune the large pre-trained language models on the particular downstream task. This is indeed a very strong baseline for many

tasks, some even reaching human-level performance. It is also the backbone of most of the top-performing methods for the Social IQA task. In this study, we start from such a practice, and aim to improve it in the context of Social IQA from several directions, which we believe are applicable to other language understanding tasks as well.

(i) How can we improve the model’s stability? Training instability is a common issue when fine-tuning large pretrained language models on downstream tasks [11, 17]. We propose to use multi-task learning with both generation loss and classification loss. We also introduce an ensemble method for incorporating different pretrained language models, such as GPT2 and RoBERTa, making the final result less sensitive to an individual model.

(ii) How can we optimize the model for the multi-choice task? Rather than treating each answer candidate independently, we propose to look at all the choices together, and thus allowing the model to compare them when making predictions. In addition, we introduce inter-segment attention to model the relationships between the given context, question, and answer options for the Social IQA multi-choice task.

(iii) How can we utilize other datasets? There are other multiple-choice datasets related to common sense, with different focuses from Social IQA. We adopt a two-phase approach, where we first fine-tune the models on other tasks, and then on Social IQA. This way the model learns general common sense first and then is optimized for the specific Social IQA task.

We conduct thorough experiments on the Social IQA task, showing that while the model size and pretraining corpora have a great influence on the final result, sophisticated architecture and similar-domain knowledge usage in the fine-tuning step are very beneficial. Our experimental results demonstrate the contribution and the competitive performance of different methods we proposed. We expect that our analyses may help future work to better understand pretrained models’ capability of social intelligence.

*Work was done while the first and fifth author were interns at Amazon

2. USING PRETRAINED LANGUAGE MODELS AS BASELINE

In Social IQA [4], given a context C about an event, and a corresponding question Q , the goal is to select the correct answer from the set $A = (A_1, A_2, A_3)$. Figure 1 shows an example.

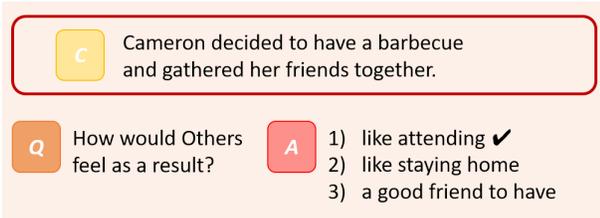


Fig. 1. An example in the Social IQA data.

A typical approach for solving this kind of multiple-choice problem with a pretrained transformer-based language model is done by concatenating C , Q , and A_i with a separator token, and then letting the model output a score via an MLP (multi-layer perceptron) built on top of the final hidden representation of the classifier token $[CLS]$. Finally, scores for each data point are normalized across all (C, Q, A_i) instances with softmax. For example, if we are fine-tuning BERT, the input is formatted as below:

$$[CLS] C [SEP] Q [SEP] A_i [SEP]$$

On the other hand, when using GPT-based models, the $[CLS]$ token is placed at the end of the sequence instead of the beginning, since it is feed-forward. Cross-entropy loss is generally used for model training.

As RoBERTa [12] has shown competitive performance on the Social IQA leaderboard, we use it as the main baseline model. In the following, we propose methods to improve the model’s stability and better represent the multi-choice task.

3. IMPROVING MODEL ROBUSTNESS

3.1. Training Stability

Training instability is a common issue when fine-tuning large pretrained language models on downstream tasks [11, 17]. With a poor hyper-parameter setting, a model achieving state-of-the-art may perform no better than random guessing. A probable explanation is that the model suffers from *catastrophic forgetting* [18, 19].

To address this problem, in addition to cross-entropy loss for multiple-choice (MC) classification, we add the masked language modeling (MLM) loss [11] in a multi-task training setup. Therefore the loss in BERT-based models is the sum of the MLM and MC CE loss, and for GPT-based models, it is the unidirectional LM loss plus MC loss. The idea is that while a new downstream task might be challenging to learn

at the beginning of the fine-tuning process, the generation loss can provide a stable objective that matches the pretraining phase. Therefore, the model could make steady progress and be less sensitive to different hyper-parameters. Note that we only train the (masked) LM loss on inputs with the correct answer A_{y^*} , where y^* is the reference label.

3.2. Combining RoBERTa with GPT2

In our initial experiments, we observed RoBERTa-large outperforms GPT2-medium model (both fine-tuned on the social IQA data set), with $\sim 10\%$ performance gap. With such a large difference, simply combining the two models’ predictions is not an effective way for system combination. Since GPT2 is a unidirectional regressive model and there are some differences in the pretraining corpora used between GPT2 and RoBERTa, we expect these two models may be complementary. To fuse these two models effectively, we do not train them jointly since it is computationally expensive; instead, we propose a two-step approach, which only needs to keep one pretrained model in the memory at each step.

Figure 2 illustrates the two steps. First, we fine-tune the GPT2 model on Social IQA, with the $[CLS]$ token placed at the end of the concatenated sequence of C , Q and A_i , and an MLP classifier stacked upon it. We believe that the final hidden representation of $[CLS]$ condenses some semantic meanings of the entire sequence, and it provides useful information for this multi-choice task after fine-tuning. Hence, we extract this $[CLS]$ token representation and use it as an additional feature when we move on to fine-tune RoBERTa (the second step). Our proposed 2-step ensemble method is especially tailored to large pretrained models, while sharing a similar idea to the traditional Stacked generalization [20].

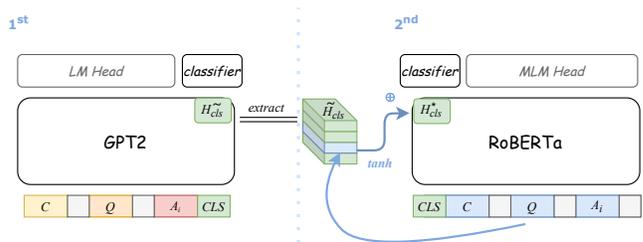


Fig. 2. The pipeline of combining GP2 and RoBERTa. First, we fine-tune GPT2 and then store its representations for all the instances. Secondly, RoBERTa retrieves the corresponding one (H_{CLS}) to use in its fine-tuning step.

Specifically, let H_{CLS} and $H_{CLS}^* \in \mathbb{R}^d$ be the $[CLS]$ hidden representation of GPT2 and RoBERTa respectively. For each training instance, we extract H_{CLS} from the fine-tuned GPT2, and then combine it with H_{CLS}^* while fine-tuning

RoBERTa for the multi-choice task:

$$H_{CLS}^- = [\tanh(H_{CLS}^-) \oplus H_{CLS}^*] \quad (1)$$

$$Score = H_{CLS}^- W_t + b_t \quad (2)$$

where \oplus can be concatenation or mean-pooling¹. Finally, the combined H_{CLS}^- is passed to RoBERTa’s classifier to predict the score. In both steps, we apply the generation loss besides classification loss (Section 3.1), using unidirectional LM for GPT2 and MLM for RoBERTa. Note that H_{CLS}^- serves as external features for RoBERTa and will not be tuned while updating RoBERTa’s parameters.

4. OTHER PROPOSED MODEL VARIATIONS

4.1. Cross-Segment Attention

In the standard self-attention mechanism, each token attends to *all the tokens* in the sequence. Yet in a task like Social IQA, there is a distinct separation of context, question, and answer, and we would like to emphasize the differences between these segments. Previous work has dealt with this property with different architectures, usually using complex cross-segment (or named multiway) attention blocks [2, 21, 22, 23].

In this paper, we apply a simple adaption of previous work [2] that suits BERT well. We keep the $[SEP]$ token in the input and leave the self-attention layer untouched, but stack an inter-segment attention layer using the final hidden representations. This additional attention layer attends between different segments only. This design is based on two considerations:

- We would like to follow the original BERT framework that inherits the huge transformer block pretrained on large corpora, and fine-tune it with the simple task-specific layer built atop [11].
- We want to keep the advantage of seeing global information through self-attention [24], at the same time stress the difference of segments on the semantic level, which is believed to be learned by the higher layers of BERT [25].

Formally, let d be the hidden dimension, l be the sequence length of the input, $H_C, H_Q,$ and $H_{A_i} \in \mathbb{R}^{l \times d}$ be BERT’s final representation (zero-padded to l for parallelism), corresponding to the context, question and the i^{th} answer respec-

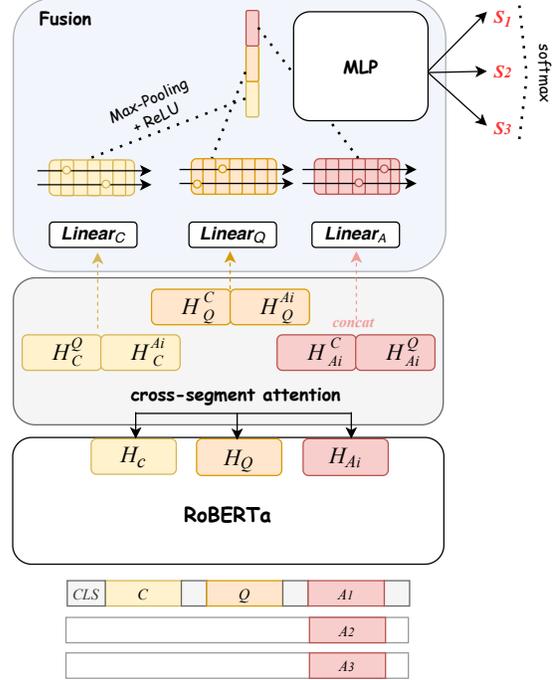


Fig. 3. The Cross-Segment Attention model.

tively. We apply the inter-segment attention:

$$\begin{aligned} \hat{H}_C &= H_C W + \mathbf{1} * b_1^T, \\ \hat{H}_Q &= H_Q W + \mathbf{1} * b_1^T, \\ \hat{H}_{A_i} &= H_{A_i} W + \mathbf{1} * b_1^T, \\ H_C^Q &= \text{Softmax}\left(\frac{\hat{H}_C \hat{H}_Q^T}{\sqrt{d}}\right) H_Q \\ H_{A_i}^{A_i} &= \text{Softmax}\left(\frac{\hat{H}_{A_i} \hat{H}_{A_i}^T}{\sqrt{d}}\right) H_{A_i} \end{aligned} \quad (3)$$

where $W \in \mathbb{R}^{d \times d}, b \in \mathbb{R}^d, \mathbf{1} \in \mathbb{R}^l$ (a vector of all-ones), H_C^Q and $H_{A_i}^{A_i} \in \mathbb{R}^{l \times d}$ are question-attentive and answer-attentive representation (zero-padded) for the context C respectively, and \sqrt{d} is used to smooth the softmax results [24]. Similarly, we obtain context-attentive and answer-attentive representation (H_Q^C and $H_{A_i}^Q$) for question Q , and context-attentive and question-attentive representation for answer A_i ($H_{A_i}^C$ and $H_{A_i}^Q$).

Once we obtain all these representations ($H_C^Q, H_{A_i}^{A_i}, \dots$), we can then apply a *fusion layer* to utilize all of them and output a score for classification. In our experiments, we evaluated different commonly used fusion methods, including element-wise addition, subtraction, concatenation, ReLU and max-pooling over the sequence dimension, and several linear transformations. The final model architecture is decided by our experimental results. Figure 3 illustrates the model using cross segment attentions.

¹Note that in Hugging Face’s implementation, BERT-based models’ final $[CLS]$ representation is normalized by tanh, so we scale H_{CLS}^- by tanh as well before the combination. RoBERTa-large and GPT2-medium have the same hidden dimension $d = 1024$.

4.2. Seeing all Choices

In the baseline model introduced in Section 2, the model assigns a score for each answer independently, without looking at other answer candidates. However, we human often compare all choices when solving multiple-choice problems, and when multiple answers look reasonable, we choose the most plausible one, and we use the method of exclusion frequently. Doing these kinds of reasoning requires comparisons of different choices, instead of independently generating prediction scores for each candidate. Motivated by this, we re-format the input to allow the model to see all the choices as below:

[CLS] C [SEP] Q [SEP] A₂ [SEP] A₃ [UNK] A₁
[CLS] C [SEP] Q [SEP] A₁ [SEP] A₃ [UNK] A₂
[CLS] C [SEP] Q [SEP] A₁ [SEP] A₂ [UNK] A₃

where the answer after the unknown token [UNK] is the target candidate, i.e., the model’s prediction score is for that choice.

5. USING EXTERNAL RESOURCES

The models introduced in the previous section heavily rely on fine-tuning pretrained models on the Social IQA dataset. Given the size of the training data and the complexity of social common sense, it is likely the models are undertrained. We believe it is beneficial for the model to access external information to learn more about language related to social scenarios. There are plenty of common sense resources with different focuses [1, 2, 3, 5, 26, 27, 28], which may be very helpful for our model to learn more broad and general concepts of our world. In [29], we investigated how to incorporate knowledge graph into the pretrained models.

In this study, our focus is on leveraging other multiple-choice datasets of common sense to help the model learn more general concepts:

- HellaSwag [3] tests a model’s ability to choose the most plausible continuation of a story. There are no questions in this data, so each instance is formatted as [CLS] C [SEP] A_i [SEP].
- CosmosQA [2] evaluates machine’s reading comprehension of people’s everyday narratives. It is formatted the same as Social IQA (see Section 2).
- WinoGrande [5] improves the scale and the hardness of WSC [30], focusing on coreference resolution. WinoGrande contains one sentence with a blank and two options to be filled in. Following [5], an example is formatted as ” [CLS] *The trophy doesn’t fit into the brown suitcase because the [SEP] _ is too large. [SEP]*”, where the blank is filled with either A₁ or A₂.

Note that these additional datasets are built based on different sources, and ATOMIC [28], the source of Social IQA, is not

included, therefore using these datasets should not compromise testing on Social IQA’s data.

We adopt 2-stage fine-tuning [31, 32], where we use these additional datasets in the first stage fine-tuning, and then continue to tune the model on Social IQA in the second stage. Specifically, in the first stage, we use a multi-task learning setup as MTDNN [33], where all the additional datasets are jointly trained at equal sample rates. Since all the datasets are multiple-choice tasks, we adopt a unified classifier to assign a score for each choice, instead of one classifier per task. Empirically, we found no significant differences in sharing it or not. We trained our model with N -way cross-entropy loss, where N is the number of choices, with $N_{SocialIQA} = 3$, $N_{CosmosQA} = 4$, $N_{HellaSwag} = 4$, and $N_{WinoGrande} = 2$. Note that we not only utilize the classification labels in these external datasets, but also train our model with MLM loss (Section 3.1) on the diverse text they provided. Fine-tuning on these datasets is expected to improve the model’s general representation for common sense related words and tasks. The idea is similar to [34], where they show the benefits when continuing to pretrain models on data in the task domain.

6. EXPERIMENTS AND RESULTS

6.1. Experimental Setup

We train our models on the 33k Social IQA training instances, running hyper-parameters search for every proposed model respectively, and report their best performance on the dev set. We ran a grid search over the learning rate in $\{5e - 6, 1e - 5, 2.5e - 5, 5e - 5\}$, and the *effective batch size* (number of GPUs \times batch size per GPU \times gradient accumulation steps) in $\{16, 32, 64, 256\}$. Every model was trained on 8 NVIDIA V100 16GB GPUs, using Hugging Face’s `transformers` toolkit².

6.2. Results

Pretraining To evaluate the impact of pretraining, we first fine-tuned several large pre-trained language models using the same framework as in [4], including variants of BERT, RoBERTa, and GPT2. Table 1 shows the results. We can see a consistent trend: the larger the model size and the pre-trained corpus, the better the performance is. Specifically, RoBERTa-large and GPT2-medium have similar model size; however, RoBERTa was pretrained on a larger corpus than GPT2’s (160GB vs. 40GB). Furthermore, our experiment of training a GRU-RNN model from scratch only achieved about 52% accuracy, again showing the necessity of using large-scale pretrained models for such a commonsense task. All our following experiments are based on RoBERTa-large as it performs the best among all these baseline models.

²<https://huggingface.co/transformers/>

Model	Accuracy (%)	Model Size (M)
BERT-base	63.3	110
BERT-large	66.0	340
GPT2-medium	66.2	345
GPT2-large	71.7	1558
RoBERTa-base	69.0	125
RoBERTa-large	78.0	355

Table 1. Results on Dev set using different pretrained models fine-tuned on Social IQA.

MLM We then demonstrate the effectiveness of jointly training the MLM loss (see Section 3.1) in Table 2. As our goal is to improve training stability, we perform two rounds of 18 experiments with different learning rates and gradient accumulation steps, comparing the average performance of RoBERTa-large with and without MLM. While the best and worst performances of the two models are about the same, the average performance of the one with MLM is much better ($\sim 5\%$ improvement). Note here we do not apply other proposed methods and follow the basic framework [4] to evaluate the impact of adding MLM in training.

Model	Best	Worst	Mean, std
RoBERTa-large	78.0	32.9	63.8, 0.17
+ MLM	78.4	33.2	<u>69.4</u> , 0.13

Table 2. Results (accuracy in %) on the Dev set showing the effectiveness of incorporating MLM loss.

Proposed Methods We evaluate our proposed models and report their best performances on the dev set in Table 3. All of the proposed methods lead to better results. Specifically, we experimented with three GPT2-medium models fine-tuned with different hyper-parameters as feature extractors, and found all of them are beneficial. However, when adopting features from GPT2-xl, which is $\sim 5\%$ better than GPT2-medium on the Social IQA dev set, we found no improvement than using RoBERTa only. A possible explanation is that GPT2-medium’s features are more complementary to RoBERTa-large. It is worth noting that when we applied such a GPT2-medium and RoBERTa-large combination to other commonsense datasets, we also found noticeable improvements in preliminary experiments, including CosmosQA, HellaSwag, and WinoGrande.

The last row in Table 3 shows the results when ensembling 5 models with majority vote, including See all Choices, Cross-Segment Attention, RoBERTa + GPT2, Multiple-Choice Datasets, and one system using external concept knowledge graph that has a classification accuracy of 79.2% [29]. This ensemble outperforms all the individual models. We also submitted this method for the test set to Social IQA leaderboard. Our system achieves 78.3%, a competitive result to other state-of-the-art models.

Model	Accuracy (%)
Sap et al. [4]	78.0
See all Choices	78.9
Cross-Segment Attention	79.2
RoBERTa + GPT2	78.6, 79.1, 79.7
Multi-Choice Datasets	79.3
Ensemble	81.1

Table 3. Comparisons of different proposed models. The three values in RoBERTa + GPT2 row correspond to three different GPT2-medium models fine-tuned on Social IQA.

6.3. Analyses

To better understand the behavior of pretrained models on the Social IQA task, we analyzed the RoBERTa-large model (with MLM) from several aspects.

Few-Shot Learning We randomly chose 10% training instances and found the model achieved 72.3% on the dev set, which is better than any other models in Table 1. This may imply that a large pretrained model like RoBERTa does learn some commonsense of social intelligence during pretraining, and it only needs a small amount of training instances to adapt to a specific task.

Question Type Transfer Furthermore, we fine-tuned our model based on different question types of training instances³, where the model is tuned on only one type and evaluated on the others. This setting challenges the model’s ability to do zero-shot question type transferring. The results in Figure 4 show that most of the pairs are much better than random guessing ($\sim 33\%$), even though some types, such as *why*, *need*, *happen*, account for less than 15% training instances.

Zero-Shot Task Transfer The model utilizing additional multiple-choice datasets shows promising results in Table 3. Here, we perform zero-shot transferring between these datasets and Social IQA, showing that the knowledge learned from these tasks indeed is highly transferable to Social IQA (Table 4). Among them, CosmosQA is the most transferable, which makes sense since it has the same data format of (C, Q, A) as Social IQA, despite having much longer context.

Information Masking Previous work [35, 36, 37, 38] on question-answering has shown that some models tend to learn from artificial or superficial patterns of the dataset, and they can still predict the correct answer after import clues (to human) in the premise are masked. Therefore, we challenge our well-trained model, a RoBERTa with MLM fine-tuned on the original training set, by masking context, question, and both, respectively, during inference. We evaluated three different

³Instances in Social IQA can be divided into 6 question types, i.e., *wants*, *reactions*, *descriptions*, *motivations*, *needs*, *effects*. Since the data does not contain such annotations, here we categorize instances by keyword matching. Our categorization has a similar distribution to that in the initial paper.

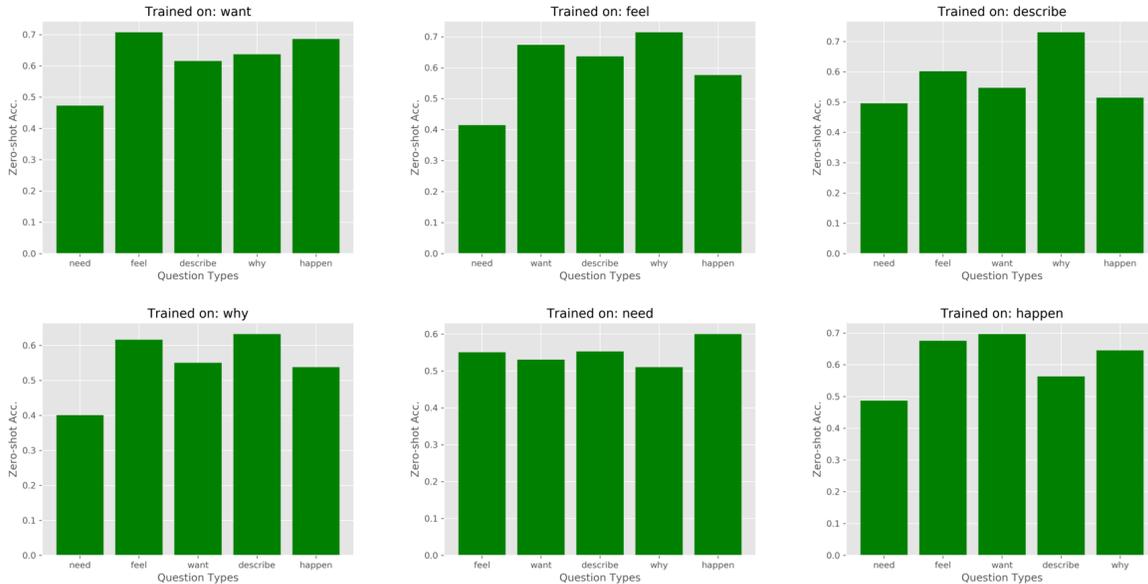


Fig. 4. Zero-shot transferring among 6 question types. The top-left type (*want*) accounts for the most training instances, while the bottom-right type (*happen*) has the least. Each bar corresponds a transfer pair.

Training Task	Social IQA (%)
HellaSwag	52.7
CosmosQA	63.7
WinoGrande	59.2
None	33.3

Table 4. Fine-tuning on three different commonsense tasks and evaluating on Social IQA’s dev set.

Info.	Accuracy (%)		
	Pad	Delete	Filter
Full	78.4		
- Context	51.8	52.2	58.4
- Question	58.4	57.8	64.3
- Both	44.3	40.0	47.7

Table 5. Effect of different information masking methods on a well-trained ROBERTa model.

masking methods:

- Zero-pad the masked segment, so the sequence length remains unchanged.
- Delete the entire masked segment.
- Filter out the keywords in the masked segment, where we define keywords as *noun*, *verb*, *adjective*⁴ and other task-important words like “*before*, *next*, *What*, *How*, *Why*”. For example, “*Tracy didn’t go home that evening and resisted Riley’s attacks.*” becomes “*Tracy didn’t that and Riley’s*”.

Table 5 shows that without either context or question, the model’s performance does drop significantly, and when both are not provided, it degrades even more. The patterns for different masking approaches are the same. Note that **Filter** may not remove all the useful words, but we can already see the significant performance drops. This analysis shows the model is relying on meaningful information to make predictions.

⁴We used the POS tagger from <https://spacy.io/models/en>

7. CONCLUSION

In this paper, we attempt to build a system with commonsense social intelligence. Building on powerful pretrained language models, we propose several variations focusing on different aspects: alleviating the training instability problems by jointly training the model with the MLM and multiple-choice classification loss, considering Social IQA’s properties and designing more sophisticated architectures, utilizing external commonsense resources, and ensembling models with complementary properties. Our proposed methods achieve competitive results on Social IQA leaderboard. Besides, we found that the pretraining step highly dominates the performance, and a large pretrained language model such as ROBERTa can generalize commonsense knowledge well. Though our experiments are focused on Social IQA, we believe our methods and findings may help future work to develop models for similar tasks and machine social intelligence in general.

8. REFERENCES

- [1] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant, “CommonsenseQA: A question answering challenge targeting commonsense knowledge,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019, pp. 4149–4158.
- [2] Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi, “Cosmos qa: Machine reading comprehension with contextual commonsense reasoning,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019, pp. 2391–2401.
- [3] Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi, “HellaSwag: Can a machine really finish your sentence?,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 4791–4800.
- [4] Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi, “Social iqa: Commonsense reasoning about social interactions,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019, pp. 4453–4463.
- [5] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi, “Winogrande: An adversarial winograd schema challenge at scale,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019, pp. 8732–8740.
- [6] Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi, “Unsupervised commonsense question answering with self-talk,” *arXiv preprint arXiv:2004.05483*, 2020.
- [7] Bill Yuchen Lin, Xinyue Chen, Jamin Chen, and Xiang Ren, “Kagnet: Knowledge-aware graph networks for commonsense reasoning,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019, pp. 2822–2832.
- [8] Arindam Mitra, Pratyay Banerjee, Kuntal Pal, Swaroop Mishra, and Chitta Baral, “How additional knowledge can improve natural language commonsense question answering?,” *arXiv preprint arXiv:1909.08855*, 2019.
- [9] Debjit Paul and Anette Frank, “Social commonsense reasoning with multi-head knowledge attention,” *arXiv preprint arXiv:2010.05587*, 2020.
- [10] Antoine Bosselut and Yejin Choi, “Dynamic knowledge graph construction for zero-shot commonsense question answering,” *arXiv preprint arXiv:1911.03876*, 2019.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “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 (NAACL)*, 2019, pp. 4171–4186.
- [12] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov, “Roberta: A robustly optimized bert pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [13] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever, “Improving language understanding by generative pre-training,” 2018.
- [14] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever, “Language models are unsupervised multitask learners,” 2019.
- [15] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu, “Exploring the limits of transfer learning with a unified text-to-text transformer,” *arXiv preprint arXiv:1910.10683*, 2019.
- [16] Daniel Khoshdel, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi, “Unifiedqa: Crossing format boundaries with a single qa system,” *arXiv preprint arXiv:2005.00700*, 2020.
- [17] Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah Smith, “Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping,” *arXiv preprint arXiv:2002.06305*, 2020.
- [18] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al., “Overcoming catastrophic forgetting in neural networks,” *Proceedings of the national academy of sciences*, vol. 114, no. 13, pp. 3521–3526, 2017.
- [19] Xinyang Chen, Sinan Wang, Bo Fu, Mingsheng Long, and Jianmin Wang, “Catastrophic forgetting meets negative transfer: Batch spectral shrinkage for safe transfer learning,” in *Advances in Neural Information Processing Systems*, 2019, pp. 1908–1918.
- [20] David H Wolpert, “Stacked generalization,” *Neural networks*, vol. 5, no. 2, pp. 241–259, 1992.

- [21] Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi, “Bidirectional attention flow for machine comprehension,” in *International Conference on Learning Representations*, 2017.
- [22] Haichao Zhu, Furu Wei, Bing Qin, and Ting Liu, “Hierarchical attention flow for multiple-choice reading comprehension,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- [23] Liang Wang, Meng Sun, Wei Zhao, Kewei Shen, and Jingming Liu, “Yuanfudao at semeval-2018 task 11: Three-way attention and relational knowledge for commonsense machine comprehension,” in *Proceedings of The 12th International Workshop on Semantic Evaluation*, 2018, pp. 758–762.
- [24] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [25] Ian Tenney, Dipanjan Das, and Ellie Pavlick, “BERT rediscovers the classical NLP pipeline,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 4593–4601.
- [26] Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi, “Piqa: Reasoning about physical commonsense in natural language,” *arXiv preprint arXiv:1911.11641*, 2020.
- [27] Hugo Liu and Push Singh, “Conceptnet—a practical commonsense reasoning tool-kit,” *BT technology journal*, vol. 22, no. 4, pp. 211–226, 2004.
- [28] Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A Smith, and Yejin Choi, “Atomic: An atlas of machine commonsense for if-then reasoning,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019, vol. 33, pp. 3027–3035.
- [29] Ting-Yun Chang, Yang Liu, Karthik Gopalakrishnan, Behnam Hedayatnia, Pei Zhou, and Dilek Hakkani-Tur, “Incorporating commonsense knowledge graph in pre-trained models for social commonsense tasks,” in *Proceedings of Deep Learning Inside Out (DeeLIO): The First Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, 2020, pp. 74–79.
- [30] Hector Levesque, Ernest Davis, and Leora Morgenstern, “The winograd schema challenge,” in *Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*, 2012.
- [31] Jason Phang, Thibault Févry, and Samuel R Bowman, “Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks,” *arXiv preprint arXiv:1811.01088*, 2018.
- [32] Yada Pruksachatkun, Jason Phang, Haokun Liu, Phu Mon Htut, Xiaoyi Zhang, Richard Yuanzhe Pang, Clara Vania, Katharina Kann, and Samuel R Bowman, “Intermediate-task transfer learning with pretrained models for natural language understanding: When and why does it work?,” *arXiv preprint arXiv:2005.00628*, 2020.
- [33] Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao, “Multi-task deep neural networks for natural language understanding,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 4487–4496.
- [34] Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A Smith, “Don’t stop pretraining: Adapt language models to domains and tasks,” *arXiv preprint arXiv:2004.10964*, 2020.
- [35] Shi Feng, Eric Wallace, Alvin Grissom II, Pedro Rodriguez, Mohit Iyyer, and Jordan Boyd-Graber, “Pathologies of neural models make interpretation difficult,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.
- [36] Pride Kavumba, Naoya Inoue, Benjamin Heinzerling, Keshav Singh, Paul Reiser, and Kentaro Inui, “When choosing plausible alternatives, clever hans can be clever,” in *Proceedings of the First Workshop on Commonsense Inference in Natural Language Processing*, 2019, pp. 33–42.
- [37] Paul Trichelair, Ali Emami, Adam Trischler, Kaheer Suleman, and Jackie Chi Kit Cheung, “How reasonable are common-sense reasoning tasks: A case-study on the winograd schema challenge and swag,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019, pp. 3373–3378.
- [38] Alexandre Tamborrino, Nicola Pellicano, Baptiste Pannier, Pascal Voitot, and Louise Naudin, “Pre-training is (almost) all you need: An application to commonsense reasoning,” *arXiv preprint arXiv:2004.14074*, 2020.