

RFiD: Towards Rational Fusion-in-Decoder for Open-Domain Question Answering

Cunxiang Wang^{*}, Haofei Yu^{♡*}, Yue Zhang^{♣†}

[♣]School of Engineering, Westlake University, China

[♡]Language Technologies Institute, Carnegie Mellon University, USA
{wangcunxiang, zhangyue}@westlake.edu.cn; haofei@cs.cmu.edu

Abstract

Open-Domain Question Answering (ODQA) systems necessitate a reader model capable of generating answers by simultaneously referring to multiple passages. Although representative models like Fusion-in-Decoder (FiD) have been proposed to address this challenge, these systems can inadvertently rely on spurious features instead of genuine causal relationships between the question and the passages to generate answers. To counter this problem, we introduce the Rational Fusion-in-Decoder (RFiD) model. Our model leverages the encoders of FiD to differentiate between causal relationships and spurious features, subsequently guiding the decoder to generate answers informed by this discernment. Experimental results on two ODQA datasets, Natural Questions (NQ) and TriviaQA (TQ), demonstrate that our model surpasses previous methods, achieving improvements of up to 1.5 and 0.7 in Exact Match scores on NQ, and exhibits an enhanced ability to identify causal relationships.¹

1 Introduction

Open-domain Question Answering (ODQA) has garnered significant attention (Chen et al., 2017; Kwiatkowski et al., 2019; Joshi et al., 2017), leading to the development of various systems designed to retrieve relevant passages (Karpukhin et al., 2020; Bevilacqua et al., 2022) from large databases and generate corresponding answers (Izacard and Grave, 2020b; Lewis et al., 2020). We utilize the Fusion-in-Decoder (FiD) (Izacard and Grave, 2020b) model as our baseline model, a sequence-to-sequence paradigm based on the T5 model (Raffel et al., 2020). Given a question, the FiD model encodes K retrieved passages using K respective T5 encoders, concatenates these K encoder hidden

^{*}Co-first Author

[†]The corresponding author.

¹Our code and data are available at <https://github.com/wangcunxiang/RFiD>

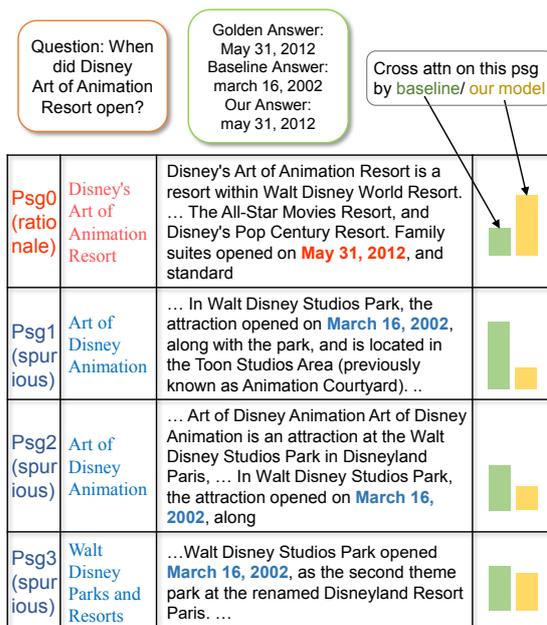


Figure 1: An example from our experiments. The question has only one relevant passage (red Psg0), while the remaining blue ones represent three passages that contain the wrong answer generated by the baseline model.

states, and then feeds the result into a T5 decoder to generate the answer.

FiD treats all passages equally in its encoders, relying exclusively on the cross-attention mechanism to establish correlations between the decoder and encoders. However, the cross-attention mechanism lacks an explicit mechanism for distinguishing differences among passages, which can result in the detection of spurious patterns (Slack et al., 2020; Jo and Bengio, 2017). Consequently, it becomes challenging for the model to identify crucial passages. An example of such spurious patterns observed in our experiment is depicted in Figure 1, where the model confuses "Disney Art of Animation Resort" with "Art of Disney Animation" due to the prevalence of passages about the latter, resulting in an incorrect answer.

To address this issue, we propose a conceptually straightforward strategy by introducing a rationalization process to explicitly determine which retrieved passages contain the answer before conducting answer generation. This process assigns different embeddings to rationale passages and irrelevant passages. These embeddings then guide the cross-attention during the answer generation phase. We dub this new model the **Rational Fusion-in-Decoder (RFiD)**.

We evaluate the effectiveness of our proposed RFiD model through experiments on the Natural Questions (NQ) (Kwiatkowski et al., 2019) and TriviaQA (TQ) (Joshi et al., 2017) datasets. Our results demonstrate that our methods can help models overcome spurious patterns and enhance their reasoning abilities, leading to an improvement of up to 1.5/0.7 Exact Match points on the NQ/TQ datasets respectively. Further analysis reveals that our methods effectively direct models to focus more on correct causal features and less on spurious features. For instance, as seen in the rightmost column of Figure 1, our model has increased its attention on the relevant passage.

To the best of our knowledge, we are the first to incorporate rationalization into ODQA models, thus underscoring the importance of passage rationalization.

2 Related Work

Open Domain Question Answering (ODQA). The prevailing approach to ODQA involves using a retriever to pinpoint relevant passages for a given question from a vast database, such as Wikipedia, and then employing a reader to generate the final answer. This is achieved by integrating the retrieved passages and question with a large pretrained language model. Retrievers commonly use methods ranging from string-matching algorithms like BM25, to dense retrievers such as DPR (Karpukhin et al., 2020), and semi-parametric retrievers like SEAL (Bevilacqua et al., 2022). The reader models fall into two primary categories: Extractive readers, such as DPR-reader (Karpukhin et al., 2020), identify the answer spans from the retrieved passages, while Generative readers, including the Fusion-in-Decoder model (FiD) (Izacard and Grave, 2020b) and the Retrieval-Augmented Generation model (RAG) (Lewis et al., 2020), generate the answer in a sequence-to-sequence manner.

Our work seeks to enhance the reasoning abil-

ity of the FiD reader without modifying the retriever. To this end, KG-FiD (Yu et al., 2022) uses knowledge graphs to rerank and concatenate related passages for improved performance and efficiency. GRAPE (Ju et al., 2022) incorporates knowledge graphs into FiD by integrating the Relation-aware graph neural network into the encoder. R2-D2 (Fajcik et al., 2021) combines a passage reranker, an extractive reader, and two generative readers into a single comprehensive ensemble model. Unlike these approaches, our work does not involve the use of external knowledge or alternate model architectures, instead focusing on spurious patterns and reader rationale capabilities.

Rationale. Recently, spurious patterns have come into the spotlight in NLP (Slack et al., 2020; Jo and Bengio, 2017), demonstrating a significant impact on model generalization (Kaushik et al., 2020, 2021). Various strategies have been implemented to curb spurious features in tasks like sentiment analysis (Lu et al., 2022; Yang et al., 2021), NER (Zeng et al., 2020; Yang et al., 2022), NLI (Wu et al., 2022) and more (Wang and Culotta, 2020).

Our work shares a common goal of overcoming spurious patterns and prioritizing causal features, but it distinguishes itself by using an encoder to identify causal features instead of data augmentation. To the best of our knowledge, we are the first to incorporate rationalization into ODQA.

Asai et al. (2022) also devise multi-task learning methods to train the model to select evidential passages during answer generation, a technique somewhat similar to ours. However, our work differs in two fundamental ways: 1. We strive to guide the decoder with a learnable embedding, which they do not. This approach results in superior performance with an accuracy of 50.7 vs 49.8 on NQ and 69.6 vs 67.8 on TQ. 2. We analyze the rationale ability of our RFiD model, explaining the performance gain and aligning with our motivation, which we consider a significant contribution of this paper.

LLMs in ODQA Initial attempts to employ Pretrained Language Models (PLMs) to directly answer open-domain questions without retrieval reported inferior performance compared to DPR+FiD (Yu et al., 2023; Wang et al., 2021; Rosset et al., 2021). However, with the advent of Large Language Models (LLMs) like ChatGPT and others, the promise of directly answering open questions

based solely on internal parameters became increasingly feasible (Shi et al., 2023).

A study by Wang et al. (2023) manually evaluated the performance of LLMs, including ChatGPT-(3.5/4), GPT-3.5, and Bing Chat, alongside DPR+FiD on Natural Questions (NQ) and TriviaQA (TQ) test sets. The results revealed that while FiD surpassed ChatGPT-3.5 and GPT-3.5 on NQ and GPT-3.5 on TQ, the combination of DPR+FiD still showcased considerable potential in the era of LLMs.

3 Method

In this section, we explain the baseline Fusion-in-Decoder (FiD) model and our Rational-FiD (RFiD) model. The RFiD model uses a passage-level classifier on top of each FiD encoder to determine whether the corresponding passage is a rationale for the question. It guides the decoder with a rationale embedding concatenated to the encoder hidden states, as well as a multi-task framework to blend FiD training with rationale training.

3.1 Fusion-in-Decoder for ODQA

The overall input to the reader is a question and K retrieved passages. We feed the text sequence concatenated by the question and one passage to each encoder, and the concatenation detail is in Appendix A.1. Formally, for the pair of the question and the k_{th} passage, the input textual sequence X_k is as $x_{k,1}, \dots, x_{k,i}, \dots, x_{k,L}$, where $x_{k,i}$ represents the i^{th} token and L is the maximum tokens length. We denote the target answer as Y , which is also a textual sequence. Therefore, multi-passage QA can be defined as learning the conditional probability $p(Y|X_1, \dots, X_K; \theta)$, where θ denotes the model parameters. Such a model factorizes the conditional probability into $p(y_i|y_1, \dots, y_{i-1}, X_1, \dots, X_K)$ and is trained in an auto-regressive way. We denote the FiD training loss as \mathcal{L}_{FiD} for further usage.

The Fusion-in-Decoder (FiD) (Izacard and Grave, 2020b) model has been used as a standard baseline for calculating the above probability and find the most probable answer given K question-passage sequences. FiD has a multi-encoder architecture, with shared parameters.

3.2 Passage Rationale

We define a passage as a rational passage to a question if the passage contain at least one answer span

from all golden answers, or it is a spurious passage. This is inspired by (Karpukhin et al., 2020) who use a similar method to define positive or negative passages for training the retriever. We ask the encoders of FiD to distinguish rational and spurious and guide the decoder with the results.

Formally, we denote $\mathbf{H}_k \in \mathbb{R}^{L \times d}$ as the output encoder hidden states of the k_{th} encoder, where L is the maximum tokens length and d is the dimension of hidden states. Therefore, the annotation for the input of fusion decoder can be defined as $[\mathbf{H}_1; \dots, \mathbf{H}_k; \dots, \mathbf{H}_K]$. For the k_{th} encoder and its hidden states \mathbf{H}_k , we apply a binary classifier on the top of the first token’s hidden states $\mathbf{H}_{k,1}$ to distinguish whether the passage is a rationale passage to the question. The binary classification result of the k_{th} encoder is

$$\hat{b}_k = Classifier(\mathbf{H}_{k,1}) \in \mathbb{R}^2 \quad (1)$$

The training loss used for this passage rationale task is can be defined using the Cross Entropy loss

$$\mathcal{L}_{ratn} = -(b \log(\hat{b}) + (1 - b) \log(1 - \hat{b})) \quad (2)$$

where b is the rational/spurious label while $\hat{b} \in \mathbb{R}^2$ is the classification output .

Guiding Decoder. After obtaining the output, we guide the decoder with the result by appending additional embeddings to the end of the encoder hidden states and feeding the new encoder hidden states to the decoder.

the prediction label of classification is

$$pred_k = \arg \max(\hat{b}_k) \in \{0, 1\}$$

In particular, we use two updatable embeddings

$$\mathbf{E}_{\{0,1\}}^{ratn} \in \mathbb{R}^{2 \times d}$$

to represent the passage rationale information, where d is the dimension of encoder hidden states.

The rationale embedding for the k_{th} encoder is

$$\mathbf{H}_{k,ratn} = \mathbf{E}_{pred_k}^{ratn} \in \mathbb{R}^d \quad (3)$$

So, the modified encoder hidden states of the k_{th} encoder with rationale embedding is

$$\mathbf{H}_k = [\mathbf{H}_{k,1}; \dots; \mathbf{H}_{k,j}; \dots; \mathbf{H}_{k,L}; \mathbf{H}_{k,ratn}] \in \mathbb{R}^{(L+1) \times d} \quad (4)$$

where L is the maximum tokens length and $\mathbf{H}_{k,j}$ is the hidden state of the j_{th} token.

	train	dev	test
Natural Questions	79168	8757	3610
TriviaQA	78785	8837	11313

Table 1: Data details of two datasets.

	#para	NQ		TQ	
		dev	test	dev	test
RAG	626M	-	44.5	-	56.1
FiD-base	440M	-	48.2	-	65.0
FiD-large	990M	-	51.4	-	67.6
KG-FiD-base	443M	-	49.6	-	66.7
KG-FiD-large (Yu et al., 2022)	994M	-	53.4	-	69.8
GRAPE-base	454M	-	48.7	-	66.2
GRAPE-large (Ju et al., 2022)	1.01B	-	53.5	-	69.8
Our Implementations					
FiD-base	440M	49.3	50.2	68.6	69.0
RFiD-base	440M	50.0	50.7	69.6	69.6
FiD-large	990M	51.6	52.8	71.6	71.9
RFiD-large	990M	52.5	54.3	72.7	72.6

Table 2: Exact Match scores of different models on the dev/test set of Natural Questions and TriviaQA. The upper part presents results of related reader models, while the lower part presents our implemented FiD and RFiD models. Note that direct comparison between the results of previous work and ours may be affected by differences in retrieved passages.

Multi-task Learning. We propose a multi-learning framework to train the rationale classifier and the sequence-to-sequence architecture of FiD at the same time. We define the overall training loss as the sum of the binary classifier training loss \mathcal{L}_{ratn} and the FiD training loss \mathcal{L}_{FiD} :

$$\mathcal{L}_{total} = \mathcal{L}_{ratn} + \mathcal{L}_{FiD} \quad (5)$$

Details. We conduct experiments with FiD-base/large and RFiD-base/large on Natural Questions (NQ) (Kwiatkowski et al., 2019) and TriviaQA (TQ) (Joshi et al., 2017). Their statistics are shown in the Appendix Table 1. To avoid the retrieval bias, we follow Izacard and Grave (2020b,a) and adopt fixed DPR retrievers to obtaining 100 passages for each question and fix the passages in the following experiments.

4 Experiment

4.1 Main Results

Table 2 presents our main results. Our RFiD model outperforms the baseline FiD model on both

	NQ		TQ	
	EM	$r_{pos/neg}$	EM	$r_{pos/neg}$
FiD-base	50.3	3.71	69.0	2.14
RFiD-base	50.7	4.31	69.6	2.32
FiD-large	52.8	3.82	71.9	2.17
RFiD-large	54.3	4.41	72.6	2.52
RFiD-large w/o guiding decoder	53.4	4.02	72.2	2.26

Table 3: The right column shows the ratio of cross-attention scores from the decoder for positive passages versus negative passages. The experiment was conducted on the NQ-test set.

the Natural Questions (NQ) and TriviaQA (TQ) datasets.

RFiD-large achieved an exact match score of 54.3 on the NQ test set, surpassing the FiD-large baseline score of 52.8 by 1.5 points. This represents a performance increase of roughly 2.8%. On the TQ test set, RFiD-large scored 72.6, which is 0.7 points higher than the FiD-large score of 71.9, representing an improvement of approximately 0.9%. When comparing the base models, RFiD-base scored 50.7 on the NQ test set, which is 0.5 points higher than the FiD-base score of 50.2, corresponding to an approximate improvement of 1.0%. On the TQ test set, RFiD-base scored 69.6, 0.6 points higher than the FiD-base score of 69.0, reflecting an improvement of around 0.9%.

These consistent improvements across both base and large models in two different datasets highlight the robustness of our RFiD model in various contexts. Furthermore, these results support the hypothesis that incorporating rationale embeddings in the Fusion-in-Decoder architecture indeed benefits the model’s reasoning ability and overall performance. Additionally, it’s worth mentioning that our RFiD model’s performance increase is achieved with a negligible increase in parameters. This is demonstrated in the ‘#para’ column in Table 2, further validating the efficiency and practicality of our proposed approach.

In summary, our RFiD model effectively enhances the rationale ability of the Fusion-in-Decoder model, leading to improved performance on open-domain question answering tasks. Our model outperforms both the baseline and other state-of-the-art models on the Natural Questions and TriviaQA datasets, demonstrating the power of our simple yet effective approach.

4.2 Cross Attention Analysis

To evaluate the ability of our RFiD models to distinguish between positive and negative passages, we conducted a cross-attention analysis of the decoder. The principle here is straightforward: better performance would be indicated by more cross-attention focused on positive passages and less cross-attention directed towards negative passages.

The calculations for this analysis are based on a set of equations. To start with, we define the average cross-attention scores on positive passages as follows:

$$C\bar{A}_{pos} = \frac{1}{N_q} \sum_q \left(\frac{1}{N_{pos}^q} \sum_{p \in P_{pos}^q} CA_{\{q;p\}} \right) \quad (6)$$

where N_q is the number of questions, P_{pos}^q is the set of positive passages on q and N_{pos}^q is its amount, and $CA_{\{q;p\}}$ is the overall cross attention of decoder on the passage p when the question is q , which can be calculated as

$$CA_{\{q;p\}} = \sum_l^{N_{ly}} \left(\sum_{j=1}^L ca_{\{l;j\}} \right) \quad (7)$$

where N_{ly} is the number of layers, L is the maximum token length and $ca_{\{l;j\}}$ is the cross attention score of the l_{th} layer of decoder on the j_{th} token.

Thus, the ratio $r_{pos/neg}$ of the average cross attention scores of positives passages over the scores of negative passages is

$$r_{pos/neg} = \frac{C\bar{A}_{pos}}{C\bar{A}_{neg}} \quad (8)$$

The results shown in Table 3 reveal that RFiD models have higher $r_{pos/neg}$ values compared to FiD models, suggesting that RFiD focuses more on positive passages and less on negative passages. For example, the $r_{pos/neg}$ of RFiD-large is 4.41/2.52 on NQ/TQ, which is 0.39/0.35 higher than the FiD-large. Similarly, for the base model, the improvements are 0.60/0.18 on NQ/TQ. The improved ability to identify relevant passages contributes to the overall performance increase seen in our experimental results.

4.3 Analysis Without Guiding the Decoder

We also conduct ablation experiment of not guiding the decoder, which means in the Equation 4, the $\mathbf{H}_k = [\mathbf{H}_{k,1}; \dots; \mathbf{H}_{k,L}]$.

The results are displayed in the last row of Table 3. ‘RFiD-large w/o guiding decoder’ achieves

53.4 and 72.2 EM on the NQ-test and TQ-test, respectively, outperforming the baseline FiD-large by 0.6 and 0.3 EM. The $r_{pos/neg}$ values are 4.02 and 2.26, respectively, which are also higher than the baseline. The results suggest that even without explicitly guiding the decoder, encouraging encoders to discern rationales can still improve performance and the model’s ability to identify rationales. This may be because the encoders implicitly encode the rationale information into the hidden states and feed it to the decoder.

4.4 Case Study

As depicted in Figure 1, the baseline FiD identifies the incorrect answer "March 16, 2002" by referring to spurious passages (shown in blue). It confuses "Disney Art of Animation Resort" with "Art of Disney Animation". In the top 100 retrieved passages, there are many passages about the latter but only one rationale passage (shown in red) about the former. However, our RFiD model can distinguish the only rationale passage from all the spurious ones and mark it for the decoder via an explicit embedding. This enables the decoder to focus more on the rationale, leading to the correct answer.

5 Conclusion

In this study, we sought to rationalize the reader component of open question answering by introducing an explicit mechanism to identify potential answer-containing passages among the retrieved candidates. Experimental results show that our RFiD model effectively improves the rationale ability of the Fusion-in-Decoder model, leading to enhanced performance on ODQA tasks. Additionally, it outperforms the baseline and other concurrent models on both the Natural Questions and TriviaQA datasets, with only a minimal increase in the number of parameters.

Acknowledgement

We thank for Linyi Yang, Sirui Cheng for their generous help and discussion. This publication has emanated from research conducted with the financial support of the Pioneer and ‘Leading Goose’ R&D Program of Zhejiang under Grant Number 2022SDXHDX0003.

Limitations

The method of identifying rational and spurious passages sometimes makes mistakes when

1. one passage actually contains one golden answer but the content slightly differ the golden answer span, for example, ‘Messi’ vs ‘Lionel Messi’ vs ‘Lionel Andrés Messi’ vs ‘Lionel Andres Messi’;

2. one passage actually does not relate the answer but the answer span is too common and appears in the passage, such as ‘2’.

We just use only seed=0 for the experiments.

Ethics Statement

There are no known potential risks.

References

- Akari Asai, Matt Gardner, and Hannaneh Hajishirzi. 2022. [Evidentiality-guided generation for knowledge-intensive NLP tasks](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2226–2243, Seattle, United States. Association for Computational Linguistics.
- Michele Bevilacqua, Giuseppe Ottaviano, Patrick Lewis, Wen tau Yih, Sebastian Riedel, and Fabio Petroni. 2022. [Autoregressive search engines: Generating substrings as document identifiers](#). In *arXiv pre-print 2204.10628*.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading Wikipedia to answer open-domain questions. In *Association for Computational Linguistics (ACL)*.
- Martin Fajcik, Martin Docekal, Karel Ondrej, and Pavel Smrz. 2021. [R2-D2: A modular baseline for open-domain question answering](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 854–870, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Gautier Izacard and Edouard Grave. 2020a. [Distilling knowledge from reader to retriever for question answering](#).
- Gautier Izacard and Edouard Grave. 2020b. [Leveraging passage retrieval with generative models for open domain question answering](#). *CoRR*, abs/2007.01282.
- Jason Jo and Yoshua Bengio. 2017. Measuring the tendency of cnns to learn surface statistical regularities. *ArXiv*, abs/1711.11561.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, Vancouver, Canada. Association for Computational Linguistics.
- Mingxuan Ju, Wenhao Yu, Tong Zhao, Chuxu Zhang, and Yanfang Ye. 2022. [Grape: Knowledge graph enhanced passage reader for open-domain question answering](#). In *Findings of Empirical Methods in Natural Language Processing*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Divyansh Kaushik, Eduard Hovy, and Zachary C Lipton. 2020. Learning the difference that makes a difference with counterfactually augmented data. *International Conference on Learning Representations (ICLR)*.
- Divyansh Kaushik, Amrith Setlur, Eduard Hovy, and Zachary C Lipton. 2021. Explaining the efficacy of counterfactually augmented data. *International Conference on Learning Representations (ICLR)*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association of Computational Linguistics*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Jinghui Lu, Linyi Yang, Brian Mac Namee, and Yue Zhang. 2022. A rationale-centric framework for human-in-the-loop machine learning. *arXiv preprint arXiv:2203.12918*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Corbin L Rosset, Chenyan Xiong, Minh Phan, Xia Song, Paul N. Bennett, and saurabh tiwary. 2021. [Pretrain knowledge-aware language models](#).
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen tau Yih. 2023. [Replug: Retrieval-augmented black-box language models](#).

- Dylan Slack, Sophie Hilgard, Emily Jia, Sameer Singh, and Himabindu Lakkaraju. 2020. [Fooling lime and shap: Adversarial attacks on post hoc explanation methods](#). In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, AIES '20*, page 180–186, New York, NY, USA. Association for Computing Machinery.
- Cunxiang Wang, Sirui Cheng, Zhikun Xu, Bowen Ding, Yidong Wang, and Yue Zhang. 2023. [Evaluating open question answering evaluation](#).
- Cunxiang Wang, Pai Liu, and Yue Zhang. 2021. [Can generative pre-trained language models serve as knowledge bases for closed-book QA?](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3241–3251, Online. Association for Computational Linguistics.
- Zhao Wang and Aron Culotta. 2020. Robustness to spurious correlations in text classification via automatically generated counterfactuals. In *AAAI Conference on Artificial Intelligence*.
- Yuxiang Wu, Matt Gardner, Pontus Stenetorp, and Pradeep Dasigi. 2022. [Generating data to mitigate spurious correlations in natural language inference datasets](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2660–2676, Dublin, Ireland. Association for Computational Linguistics.
- Linyi Yang, Jiazheng Li, Padraig Cunningham, Yue Zhang, Barry Smyth, and Ruihai Dong. 2021. [Exploring the efficacy of automatically generated counterfactuals for sentiment analysis](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 306–316, Online. Association for Computational Linguistics.
- Linyi Yang, Lifan Yuan, Leyang Cui, Wenyang Gao, and Yue Zhang. 2022. Factmix: Using a few labeled in-domain examples to generalize to cross-domain named entity recognition. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5360–5371.
- Donghan Yu, Chenguang Zhu, Yuwei Fang, Wenhao Yu, Shuohang Wang, Yichong Xu, Xiang Ren, Yiming Yang, and Michael Zeng. 2022. [KG-FiD: Infusing knowledge graph in fusion-in-decoder for open-domain question answering](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4961–4974, Dublin, Ireland. Association for Computational Linguistics.
- Wenhao Yu, Dan Iter, Shuohang Wang, Yichong Xu, Mingxuan Ju, Soumya Sanyal, Chenguang Zhu, Michael Zeng, and Meng Jiang. 2023. Generate rather than retrieve: Large language models are strong context generators. In *International Conference for Learning Representation (ICLR)*.
- Xiangji Zeng, Yunliang Li, Yuchen Zhai, and Yin Zhang. 2020. [Counterfactual generator: A weakly-supervised method for named entity recognition](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7270–7280, Online. Association for Computational Linguistics.

A Experiments

A.1 Data Process

Following [Izcard and Grave \(2020b\)](#), we concatenate the question and the passage in the form of “Question : <question> ; Title : <title> ; Context : <context> ”.

A.2 Implementation Details

We conduct our experiments on 2 A100-80G-SXM GPUs.

In training, the optimizer is AdamW and the learning rate is $1e-4$ with weight decay rate as 0.01; the batch size is 64 and the total training step is 320k.

In evaluating, the eval step is 10k and the best-dev checkpoint is used for the test.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Just Limitations; no section number
- A2. Did you discuss any potential risks of your work?
Just Ethics Statement; no section number
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract Section and Section 1
- A4. Have you used AI writing assistants when working on this paper?
grammarly and chatgpt, check grammar mistakes and polish the writing.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
Section 4
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Section 4
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Not applicable. Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 4

C Did you run computational experiments?

Section 4

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 4

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section 4

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 4

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Not applicable. Left blank.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.