

Long Context Question Answering via Supervised Contrastive Learning

Avi Caciularu^{♣*} Ido Dagan[♣] Jacob Goldberger[♠] Arman Cohan^{◇,♡}

[♣]Computer Science Department, Bar-Ilan University, Ramat-Gan, Israel

[♠]Faculty of Engineering, Bar-Ilan University, Ramat-Gan, Israel

[◇]Allen Institute for AI, Seattle, WA

[♡]Paul G. Allen School of Computer Science, University of Washington, Seattle, WA

avi.c33@gmail.com, armanc@allenai.org

dagan@cs.biu.ac.il, jacob.goldberger@biu.ac.il

Abstract

Long-context question answering (QA) tasks require reasoning over a long document or multiple documents. Addressing these tasks often benefits from identifying a set of evidence spans (e.g., sentences), which provide supporting evidence for answering the question. In this work, we propose a novel method for equipping long-context QA models with an additional sequence-level objective for better identification of the supporting evidence. We achieve this via an additional contrastive supervision signal in finetuning, where the model is encouraged to explicitly discriminate supporting evidence sentences from negative ones by maximizing question-evidence similarity. The proposed additional loss exhibits consistent improvements on three different strong long-context transformer models, across two challenging question answering benchmarks – HotpotQA and QAsper.¹

1 Introduction

Answering questions that require reasoning over a long sequence, such over long documents or multiple documents, is a challenging task (Pang et al., 2021). Research in this domain mostly includes tasks that involve multiple text segments, over benchmarks like HotpotQA (Yang et al., 2018) and QAsper (Dasigi et al., 2021). HotpotQA is a multi-hop QA benchmark over multiple paragraphs from Wikipedia, while QAsper involves reading comprehension from long academic papers, where relevant information on a question could be spread across the paper.

Given the task complexity (Choi et al., 2017), benchmarks often provide an additional set of evidence spans, such as sentences or paragraphs, for a given question answer pair. This breaks down the long-context QA task, adding a preliminary

evidence span detection, which is crucial for successfully finding the correct answer, and also potentially helps in model interpretability. In this work, we propose a method for improving long-context QA via leveraging such evidence spans, by maximizing their similarity with the question.

Since identifying the evidences provides relevant information for answering the question, prior work showed that jointly training models to perform evidence span extraction in addition to answer generation is important for achieving high performance (Yang et al., 2018; Dasigi et al., 2021). To jointly perform evidence extraction and question answering, models utilize sentence representations (marker tokens) in the input; the final layer representation corresponding to these markers is then passed through a classification layer and is optimized using the cross-entropy loss in conjunction with the answer extraction/generation loss. We conjecture and demonstrate (see Table 1) that this objective does not sufficiently capture relationships between the question and the candidate evidence spans. Thus, we propose a complementary objective for enforcing question-evidence *similarity* in the model representation (see Fig. 1). Further, we show that optimizing the question-evidence similarity under a certain subspace by using linear projection of the raw representations may softly impose information encoding about their relatedness. Since questions may be partitioned into several types (e.g., yes/no, generative, non answerable, etc.) in the common QA settings, we also investigate learning separate projections per question type.

Driven by the intuition that questions should be similar to their supporting evidences, under some specific geometric sub-space, we propose a novel supervised contrastive learning objective for the finetuning stage, aiming to maximize similarity of question-evidence representations. Contrastive learning has been recently applied to a variety of

* Work partly done as an intern at AI2.

¹Our code is available at <https://github.com/aviclulong-context-qa-contrast>.

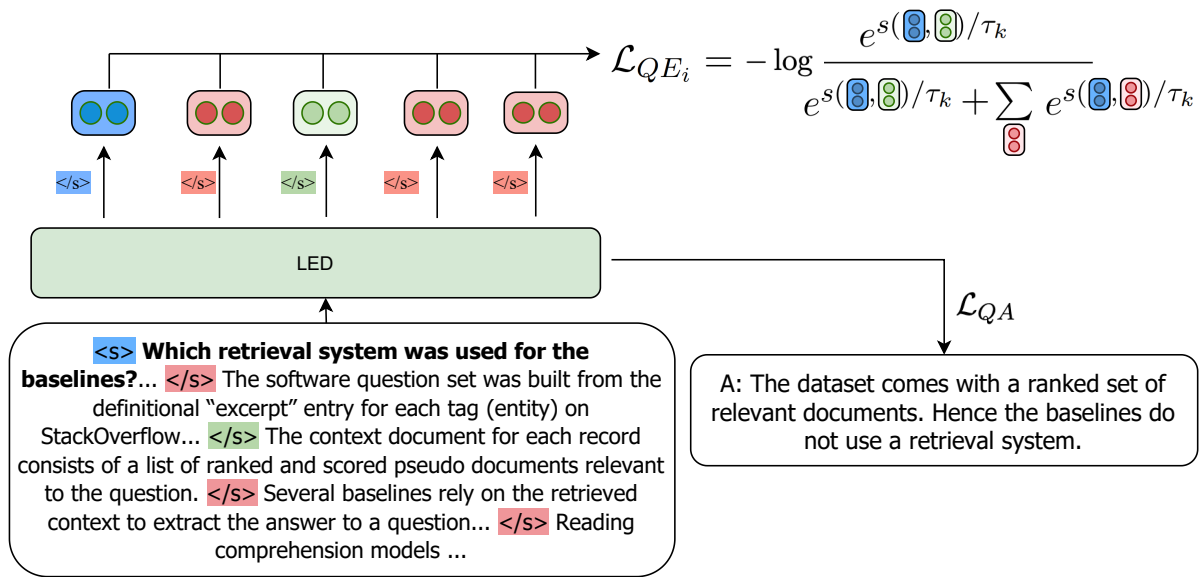


Figure 1: Demonstration of our method over an example instance from QAsper (Dasigi et al., 2021). A long sequence is fed to a model, producing representations for the marker special tokens. Then, these vectors are used to compute our contrastive objective. The token colored in blue represents the question, and the tokens colored in green (red) represent positive (negative) evidence sentences. The goal is to maximize the similarity between the blue vector and the green vectors.

deep learning models in computer vision (Chen et al., 2020; Chen and He, 2021) and NLP (Gao et al., 2021; Gunel et al., 2021). Unlike prior NLP related methods, our proposed loss term is model-agnostic and operates in the *sequence-level* in a supervised manner. This objective targets question and evidence representations within input sequences, and unlike prior work, it is not based on individually encoding sentences or paragraphs. We show that our additional contrastive supervision provides consistent improvements on three different models and two datasets, demonstrating its effectiveness and versatility.

2 Setup

In this section, we define and elaborate the common setup and notations for long-context QA.

Assume that we are given a question q_i and a context $\mathcal{S}_i = \langle s_1, \dots, s_M \rangle$ consisting of M sentences (s_j can be also a document/paragraph/passage, depending on the dataset). From \mathcal{S}_i , the task is to identify the correct answer a_i and a set of N evidence spans $\mathcal{S}_i^+ = \{s_{i_1}, \dots, s_{i_N}\}$ where i_j are indices of the sentences that are the supporting evidence for answering the question q_i .

As common in the input setup of long-context transformer models (Beltagy et al., 2020; Zaheer et al., 2020; Caciularu et al., 2021), which are the

current state-of-the-art models for solving long-context QA tasks, the question and context sentences are concatenated in a single long sequence with special sentence tokens specifying sentence boundaries. Then the input is passed to the long-context transformer, which is trained to jointly identify the evidence sentences and extract or generate the answer.

Concretely, for each example, we prepare the following concatenated input sequence:

$$\langle s \rangle, q_i, \langle /s \rangle, s_1, \langle /s \rangle, s_2, \dots, \langle /s \rangle, s_M$$

where “,” is the string concatenation operation, q_i and s_j are sequences of tokens corresponding to the question and the j th sentence in the input context, and $\langle s \rangle$ and $\langle /s \rangle$ are special tokens representing the question and a context sentence, respectively (See Fig. 1 for an example).

Then, a QA loss, which we denote by \mathcal{L}_{QA} is applied over the contextualized representation of each sentence token, and is optimized using supervision. \mathcal{L}_{QA} depends on the dataset and can take the form of a multi-task objective, representing multiple tasks in the context of QA (Dasigi et al., 2021) (particularly evidence extraction and answer generation).

3 Question-Evidence Contrastive Loss

In this section, we elaborate on our novel proposed contrastive loss term, combining question-evidence similarity maximization (§3.1), and question-type aware projections (§3.2).

3.1 Question-Evidence Similarity

To encourage the long-context transformer model to capture relationships between the question and evidence sentences, we introduce an additional sequence-level loss that compares and contrasts the question with context sentences.

The additional proposed loss \mathcal{L}_{QE} is based on the InfoNCE loss (Oord et al., 2018), and is applied over instances consisting of triplet of vectors representing the question, an evidence sentence and distractor sentences. The loss encourages the question and evidence representations to become closer to each other, while pushing the question and distracting sentences away.

Formally, the contrastive loss is defined as the sum of negative log-likelihood losses over all input examples, where each loss term discriminates the positive units from negative ones:

$$\mathcal{L}_{QE_i} = -\log \sum_{s^+ \in \mathcal{S}_i^+} \left(\frac{e^{sim(s^+, q_i)/\tau_k}}{\sum_{s \in \mathcal{S}_i} e^{sim(s, q_i)/\tau_k}} \right), \quad (1)$$

where s (q_i) is the sentence (question) marker vector representation (see $\langle s \rangle$ and $\langle /s \rangle$ in Fig. 1), τ_k is the configurable temperature hyperparameter, and $sim(\cdot)$ is a similarity metric. \mathcal{L}_{QE_i} serves a single example, and the final aggregated loss \mathcal{L}_{QE} is obtained by averaging over all the examples.

We incorporate \mathcal{L}_{QE} into the main QA span extraction/generation loss \mathcal{L}_{QA} using the augmented loss:

$$\mathcal{L} = (1 - \lambda) \cdot \mathcal{L}_{QA} + \lambda \cdot \mathcal{L}_{QE},$$

where λ is a weighting hyperparameter.

The underlying $sim(\cdot)$ can take the form of a non-parametric similarity function, e.g., the dot product ($sim(s, q) = s^\top q$) or the cosine similarity ($sim(s, q) = \frac{s^\top q}{\|s\| \|q\|}$). However, we show empirically that using such similarity over the raw representations harms the performance results of the model, since seemingly, it is hard to find a shared representation that should optimize the two loss functions. Hence, we adopted linear projections, per question type, to cast the similarity learning objective into proxy linear spaces.

3.2 Incorporating Question-Type Projections

Long-context QA benchmarks often provide a question-type label per instance as an additional challenge, such as {yes, no, span} for HotpotQA. We hypothesize that maximizing question-evidence similarity under a question-type-specific sub-space can enable more flexibility and inductive bias to the model, for producing better representations and further improving the performance. Following Iter et al. (2020), we define the following similarity function:

$$sim_k(s, q) = s^\top W_k q, \quad (2)$$

where k is the expected question type and W_k is the corresponding learnable projection matrix. Such linear projections ensure that a specific subspace per question type exists. We additionally incorporate different temperature hyperparameters τ_k per question type in Eq. 1 (see the ablation in Table 4 for their effect). The dimensions of the proposed W_k tend to be large, in accordance with the dimensions of the transformer’s hidden-layers.² Hence, following Barkan et al. (2020), we apply new non-square linear projections instead of using W_k :

$$sim_k(s, q) = \frac{s_k^\top q_k}{\|s_k\| \|q_k\|}, \quad (3)$$

where we set

$$s_k := W_k^S s, \quad q_k := W_k^Q q,$$

and W_k^S (or W_k^Q) is the matrix that projects s (or q) into a lower dimension, in the k^{th} question-type sub-space. In order to improve the separation between the different sub-spaces induced by different question types, we generated additional negative instances per sentence, as follows.

We projected every question-sentence pair using all the mappings according to the available question types, and computed their cosine similarity (according to Eq. 3). Then, all the obtained scores were considered as negative, except the ones that belong to question-evidence pairs projected using the correct question-type mapping.

4 Evaluation and Analysis

In this section, we provide details about the experiments that we conducted and their outcomes.

²Empirically, we end up with $W_k \in \mathbb{R}^{768 \times 768}$ for base-sized models and $W_k \in \mathbb{R}^{1024 \times 1024}$ for large-sized models.

| Method | Yes | No | Span |
|---|------|------|------|
| Trained only with \mathcal{L}_{QA} | 8.4 | 4.9 | 46.4 |
| + \mathcal{L}_{QE} using Eq. 3 with a single projection | 72.4 | 65.2 | 77.5 |
| + \mathcal{L}_{QE} using Eq. 3 | 86.3 | 84.2 | 87.6 |

Table 1: Mean Average Precision (mAP) question-evidence cosine similarity results per question type for the HotpotQA-distractor dev set.

4.1 Experimental Setup

In order to demonstrate the contribution of our method, we evaluated it over the recent QAsper dataset (Dasigi et al., 2021) and the well-known HotpotQA dataset (Yang et al., 2018), which share the same setup (§2).

QAsper (Dasigi et al., 2021) is a long-document QA dataset which was built over academic papers, where NLP practitioners were recruited to (abstractedly) generate questions following the title and the abstract of a particular paper, as well as creating the the correct evidence and answers to those questions based on the entire paper. More than half of the examples in QAsper require collecting evidences from multiple evidences in the given paper. For this benchmark, \mathcal{L}_{QA} represents the sum of the teacher-forced text generation and evidence classification loss functions, in a multi-task training setup.

HotpotQA (Yang et al., 2018) introduced the task of multihop extractive question answering, in the reading comprehension setting, where the inputs are a question and multiple paragraphs from various related and non-related documents. A model is queried to extract answer spans and evidence sentences, where it should handle challenging questions that require finding and reasoning over multiple supporting documents. For the models we applied to this benchmark, \mathcal{L}_{QA} represents the standard cross-entropy answer extraction loss.

To test the contribution of our \mathcal{L}_{QE} loss, We replicated the experiments described in Dasigi et al. (2021) and Caciularu et al. (2021) for QAsper and HotpotQA, respectively. For QAsper, we finetuned the LED-base model,³ and evaluated it on the question answering and evidence selection tasks. For HotpotQA, we used the Longformer model and CDLM⁴ as the backbone long-sequence language models for this task. Since CDLM was provided

³According to Dasigi et al. (2021), LED-base outperforms LED-large over QAsper.

⁴CDLM was shown to be an effective long-range cross-encoder model for HotpotQA.

only as a base-sized model, we pretrained a larger version of the CDLM model, and hence used both the base and large versions of both Longformer and CDLM. For further details see Appendix B.1 and B.2.

For both benchmarks, we performed a grid search for determining the hyper-parameters of the contrastive loss (more details in Appendix C).

4.2 Results and Analysis

Question-Evidence Similarity Analysis. As a preliminary assessment, we first investigate the question and evidences representations of models trained on the HopotQA dataset. We motivate the use of our method by presenting the mean Average Precision (mAP) ranking results produced according to the question-sentence cosine similarities for the marker tokens trained representations of the CDLM-large model. From Table 1, we observe that without applying our additional loss term, the question representations are distant from the evidence representations. Using a single learned projection increases this desirable similarity, and using a learned projection per question type yields the highest mAP scores. Hence, integrating question-type aware linear projections can be a beneficial part of our contrastive loss, and overall it further improves the QA results as we show next. An additional illustration of this effect, where we visualize the marker representations, appears in Appendix A.

Main results. We adopted the F1 evaluation metrics corresponding to the original works (Dasigi et al., 2021; Beltagy et al., 2020). Tables 2 and 3 present the evaluation results over the QAsper and HotpotQA datasets, respectively. We show the performance difference when adding our additional loss term with “+ \mathcal{L}_{QE} ” (and question-type similarity function from Eq. 3).

As shown in the table, the addition of \mathcal{L}_{QE} exhibits the best performance across all examined models and benchmarks, clearly demonstrating its consistent advantage. Note that maximizing the question-evidence similarity resulted also in evidence detection improvement – see the “Evidence” and the “Sup” metrics in Tables 2 and 3, respectively. All the results are statistically significant using the bootstrap test with $p < 0.01$ (Dror et al., 2018).

Ablations. Table 4 demonstrates ablation study results for evaluating the effectiveness of our design

| Input | Extractive | | Abstractive | | Yes/No | | Unanswerable | | Evidence | | Overall | |
|-----------------------------------|------------|------|-------------|------|--------|------|--------------|------|----------|------|---------|------|
| | Dev. | Test | Dev. | Test | Dev. | Test | Dev. | Test | Dev. | Test | Dev. | Test |
| LED (2020) | 26.1 | 31.0 | 16.6 | 15.8 | 67.5 | 70.3 | 28.6 | 26.2 | 23.9 | 29.9 | 29.1 | 32.8 |
| + \mathcal{L}_{QE} (Δ) | +0.2 | +0.2 | +0.8 | +1.0 | +1.6 | +0.2 | +1.5 | +1.9 | +1.0 | +0.7 | +0.9 | +0.7 |

Table 2: Performance change when applying our additional loss \mathcal{L}_{QE} to the LED SOTA model on QAsper.

| | Model | Ans | Sup | Joint |
|------------|-----------------------------------|------|------|-------|
| base size | Longformer-base (2020) | 74.5 | 83.9 | 64.5 |
| | + \mathcal{L}_{QE} (Δ) | +0.9 | +0.2 | +1.0 |
| base size | CDLM-base (2021) | 74.7 | 86.3 | 66.3 |
| | + \mathcal{L}_{QE} (Δ) | +0.9 | +0.3 | +0.2 |
| large size | Longformer-large (2020) | 81.3 | 88.3 | 73.2 |
| | + \mathcal{L}_{QE} (Δ) | +0.3 | +0.1 | +0.9 |
| large size | CDLM-large (2021) | 81.3 | 89.1 | 73.8 |
| | + \mathcal{L}_{QE} (Δ) | +0.3 | +0.6 | +0.8 |

Table 3: HotpotQA-distractor results (F_1) for the dev set. We use the “base” and “large” model size results of CDLM and the Longformer for direct comparison. Ans: answer span, Sup: Supporting facts.

| | Joint | Δ |
|---|-------|----------|
| CDLM-large + \mathcal{L}_{QE} and $\text{sim}(\cdot)$ is Eq. 3 | 74.6 | |
| – using a single τ parameter for all the question types | 74.2 | -0.4 |
| – $\text{sim}(\cdot)$ is the dot product | 73.1 | -1.5 |
| – $\text{sim}(\cdot)$ is the cosine similarity | 74.0 | -0.6 |
| – $\text{sim}_k(\cdot)$ is the bilinear distance (Eq. 2) | 73.7 | -0.9 |
| – $\text{sim}(\cdot)$ is Eq. 3 with a single projection ($W_k = W$) | 74.0 | -0.6 |
| – w/o incorrect question type projected negatives | 74.2 | -0.4 |

Table 4: Similarity function ablation results (Joint F_1) of CDLM-large and our loss term on the HotpotQA-distractor dev set.

decisions. Using a constant temperature parameter for all question types, as well as using different degenerated similarity functions, exhibits lower performance. Further, the last row in Table 4 shows that treating correct answers that are projected to the wrong question type as negatives also improves the results. Overall, the ablation study shows the advantage of using Eq. 3 as a similarity function that provides fine-grained expressive modeling for each question type, in its own sub-space.

Discussion. An additional theoretical justification to our contrastive learning is provided in (Gao et al., 2021), where we can imply that our loss term improves the uniformity and therefore the expressiveness of the question and evidence representations. Moreover, one can attribute the success of our contrastive loss to the fact that long-range transformer models lack long-range signals during pretraining,

and hence such explicit modeling as we suggest is necessary. In fact, comparing CDLM’s results to the Longformer illustrates that our cotractive term has higher impact on models without global attention-based pretraining.

5 Conclusion

In this work, we proposed an effective sequence-level contrastive loss for improving the performance of long-range transformers in solving QA tasks that require reasoning over long contexts. We demonstrate consistent improvement when using our approach on three different models over two different benchmarks. For future work, we propose exploring variations of our proposed supervised loss on other long-context tasks, such as long-document and multi-document summarization, and integrating our method into information retrieval re-ranker models.

Acknowledgments

We thank the BIU-NLP lab and the Semantic Scholar research team at AI2 for fruitful discussions and helpful feedback. The work described herein was supported by the PBC fellowship for outstanding PhD candidates in data science, in part by the Israel Science Foundation grant 2827/21, and by a grant from the Israel Ministry of Science and Technology.

Ethical Considerations

Our work in understanding the role of maximizing the similarity between question and evidence pairs. Therefore, there is a limited risk associated with the quality of annotated evidence sentences in the dataset, as there is no guarantee that our model will always generate non-biased and factual content. Therefore, caution must be exercised when the model is deployed in practical settings, where the evidence quality is vague and cannot be verified.

References

- Oren Barkan, Avi Caciularu, and Ido Dagan. 2020. [Within-between lexical relation classification](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3521–3527, Online. Association for Computational Linguistics.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Avi Caciularu, Arman Cohan, Iz Beltagy, Matthew Peters, Arie Cattan, and Ido Dagan. 2021. [CDLM: Cross-document language modeling](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2648–2662, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. 2016. Training deep nets with sublinear memory cost. *arXiv preprint*, abs/1604.06174.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning (ICML)*.
- Xinlei Chen and Kaiming He. 2021. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758.
- Eunsol Choi, Daniel Hewlett, Jakob Uszkoreit, Illia Polosukhin, Alexandre Lacoste, and Jonathan Berant. 2017. [Coarse-to-fine question answering for long documents](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 209–220, Vancouver, Canada. Association for Computational Linguistics.
- Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith, and Matt Gardner. 2021. [A dataset of information-seeking questions and answers anchored in research papers](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4599–4610, Online. Association for Computational Linguistics.
- Rotem Dror, Gili Baumer, Segev Shlomov, and Roi Reichart. 2018. [The hitchhiker’s guide to testing statistical significance in natural language processing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1383–1392, Melbourne, Australia. Association for Computational Linguistics.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Beliz Gunel, Jingfei Du, Alexis Conneau, and Veselin Stoyanov. 2021. Supervised contrastive learning for pre-trained language model fine-tuning. In *International Conference on Learning Representations (ICLR)*.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- Dan Iter, Kelvin Guu, Larry Lansing, and Dan Jurafsky. 2020. [Pretraining with contrastive sentence objectives improves discourse performance of language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4859–4870, Online. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). In *International Conference on Learning Representations (ICLR)*.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Richard Yuanzhe Pang, Alicia Parrish, Nitish Joshi, Nikita Nangia, Jason Phang, Angelica Chen, Vishakh Padmakumar, Johnny Ma, Jana Thompson, He He, et al. 2021. Quality: Question answering with long input texts, yes! *arXiv preprint arXiv:2112.08608*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. Big Bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems (NeurIPS)*.

Appendix

A Question-Evidence Similarity Demonstration

In this section, we further interpret the outcome of our contrastive loss.

We apply PCA over the relevant normalized token representations of the validation data of HotpotQA (i.e., the question and answers representations in Fig. 1), and depicted them in 2. The projected representations of the correct and wrong answers are equally distributed at the beginning of the training (left figure). After several epochs when the model converged (right figure), the answer representations’ manifold got closer to the questions’ representations (in terms of radial distance). Each beam in the figure corresponds to a different question type (there are 3 in HotpotQA). The correct evidence representations (green dots) are the closest among the whole answer representations, confirming that our additional contrastive loss term generalizes and maximizes the question-evidence similarity.

B Datasets and Finetuning Details

In this section, we provide details, regrading finetuning and hyper-parameter configuration, over the benchmarks we used during our experiments.

B.1 QAsper

Since some of the questions included in QAsper are not answerable, we apply our contrastive loss only over examples that are answerable and contain at least one evidence sentence.

We train all models using the Adam optimizer (Kingma and Ba, 2014) and a triangular learning rate scheduler (Howard and Ruder, 2018) with 10% warmup. To determine number of epochs, peak learning rate, and batch size, we performed manual hyperparameter search on a subset of the training data. We searched over $\{1, 3, 5\}$ epochs with learning rates $\{1e^{-5}, 3e^{-5}, 5e^{-5}, 9e^{-5}\}$, and found that smaller batch sizes generally work better than larger ones. Our final configuration was 10 epochs, peak learning rate of $5e^{-5}$, and batch size of 2, which we used for all reported experimental settings. When handling full text, we use gradient checkpointing (Chen et al., 2016) to reduce memory consumption. We run our experiments on a single RTX 8000 GPU, and each experiment takes 30–60 minutes per epoch.

B.2 HotpotQA

We used the HotpotQA-distractor dataset (Yang et al., 2018). Each example in the dataset is includes a question and 10 paragraphs from different documents, extracted from Wikipedia. Two gold paragraphs include the relevant information for properly answering the question, mixed and shuffled with eight distractor paragraphs (for the full dataset statistics, see Yang et al. (2018)). There are two goals for this task: detecting the supporting facts, i.e., evidence sentences, as well as extraction of the correct answer span.

For preparing the data for training and evaluation, we follow the same finetuning scheme of the CDLM (Caciularu et al., 2021) and the Longformer (Beltagy et al., 2020); For each example, we concatenate the question and all the 10 paragraphs in one long context. We particularly use the following input format with special tokens and our document separators: “[CLS] [q] question [/q] <doc-s><t> title₁ </t> <s> sent_{1,1} </s> <s> sent_{1,2} </s> </doc-s> ... <t> <doc-s> title₂ </t> sent_{2,1} </s> <s> sent_{2,2} </s> <s> ...” where [q], [/q], <t>, </t>, <s>, </s>, [p] are special tokens representing, question start and end, paragraph title start and end, and sentence start and end, respectively. The new special tokens were added to the models vocabulary and randomly initialized before task finetuning. We use global attention to question tokens, paragraph title start tokens as well as sentence tokens. The model’s structure is taken from Beltagy et al. (2020).

As in Beltagy et al. (2020); Caciularu et al. (2021), we finetune our models for 5 epochs, using a batch size of 32, learning rate of $1e^{-4}$, 100 warmup steps. Finetuning on our models took ~ 6 hours per epoch, using four 48GB RTX8000 GPUs for finetuning our models. For generating the CDLM-large results, we pretrained our version using the code from <https://github.com/aviclucdLM/tree/main/pretraining>.

C Contrastive Loss Details

In this section, we provide the details for reproducing our contrastive term, which is relevant for both QAsper and HotpotQA.

We searched over $d \times \{d, \frac{d}{2}, \frac{d}{4}, \frac{d}{8}\}$ to determine the linear projections’ dimensions, where d is the model’s hidden layer representation dimension (it

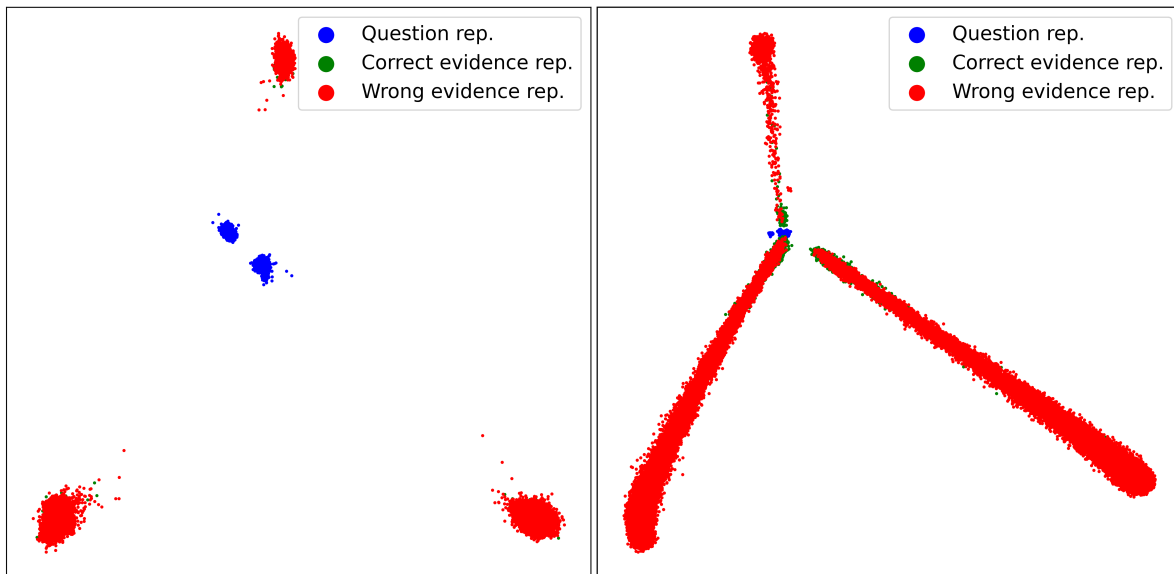


Figure 2: PCA plots of the learned question and answer token embeddings on the HotpotQA validation set, comparing early training epochs results (left) and results after model convergence (right). The wrong evidence representations correspond to both wrong evidences, or correct evidences using the wrong question type projections (our soft negatives).

depends on the size of the model). In order to determine the temperature hyperparameter τ , we searched over $\{0.2, 0.4, 0.6, 0.8, 1.0\}$ per question type (if applicable). We also applied dropout with a rate of $p = 0.1$ over the linear projections, which consistently improved the results over all the benchmarks. Finally, we searched for the best performing λ hyperparameter over the values of $\{0.2, 0.4, 0.6, 0.8, 1.0\}$.