

Generating a Common Question from Multiple Documents using Multi-source Encoder-Decoder Models

Woon Sang Cho[◇] Yizhe Zhang[†] Sudha Rao[†] Chris Brockett[†] Sungjin Lee^{§*}

[◇]Princeton University

[†]Microsoft Research AI

[§]Amazon Alexa AI

[◇]{woonsang}@princeton.edu, [§]{sungjinl}@amazon.com

[†]{yizhang, sudhra, chrisbkt}@microsoft.com

Abstract

Ambiguous user queries in search engines result in the retrieval of documents that often span multiple topics. One potential solution is for the search engine to generate multiple refined queries, each of which relates to a subset of the documents spanning the same topic. A preliminary step towards this goal is to generate a question that captures common concepts of multiple documents. We propose a new task of generating common question from multiple documents and present simple variant of an existing multi-source encoder-decoder framework, called the Multi-Source Question Generator (MSQG). We first train an RNN-based single encoder-decoder generator from (single document, question) pairs. At test time, given multiple documents, the *Distribute* step of our MSQG model predicts target word distributions for each document using the trained model. The *Aggregate* step aggregates these distributions to generate a common question. This simple yet effective strategy significantly outperforms several existing baseline models applied to the new task when evaluated using automated metrics and human judgments on the MS-MARCO-QA dataset.

1 Introduction

Search engines return a list of results in response to a user query. In the case of ambiguous queries, retrieved results often span multiple topics and might benefit from further clarification from the user. One approach to disambiguate such queries is to first partition the retrieved results by topic and then ask the user to choose from queries refined for each partition.

For example, a query *‘how good is apple?’* could retrieve documents some of which relate to apple the fruit, and some of which relate to Apple

the company. In such a scenario, if the search engine generates two refinement queries *‘how good is apple the fruit?’* and *‘how good is the company apple?’*, the user could then choose one of it as a way to clarify her initial query.

In this work, we take a step towards this aim by proposing a model that generates a common question that is relevant to a set of documents. At training time, we train a standard sequence-to-sequence model (Sutskever et al., 2014) with a large number of (single document, question) pairs to generate a relevant question given a single document. At test time, given multiple (N) input documents, we use our model, called the Multi-Source Question Generator (MSQG), to allow document-specific decoders to collaboratively generate a common question. We first encode the N input documents separately using the trained encoder. Then, we perform an iterative procedure to i) (Distribute step) compute predictive word distributions from each document-specific decoder based on previous context and generation ii) (Aggregate step) aggregate predictive word distributions by voting and generate a single shared word for all decoders. These two steps are repeated until an end-of-sentence token is generated. We train and test our model on the MS-MARCO-QA dataset and evaluate it by assessing whether the original passages can be retrieved from the generated question, as well as human judgments for fluency, relevancy, and answerability. Our model significantly outperforms multiple baselines. Our main contributions are:

- i) a new task of generating a common question from multiple documents, where a common question target is does *not* exist, unlike multi-lingual sources to common language translation tasks.
- ii) an extensive evaluation of an existing multi-source encoder-decoder models including

* Work was done when affiliated with Microsoft Research AI

our simple variant model for generating a common question.

- iii) an empirical evaluation framework based on automated metrics and human judgments on *answerability*, *relevancy*, and *fluency* to extensively evaluate our proposed MSQG model against the baselines.

2 Related Work

The use of neural networks to generate natural language questions has mostly focused on question answering (Labutov et al., 2015; Serban et al., 2016; Rothe et al., 2016; Song et al., 2017; Duan et al., 2017; Du et al., 2017; Buck et al., 2017; Song et al., 2018; Harrison and Walker, 2018; Sun et al., 2018). A number of works process multiple passages by concatenating, adding, or attention-weight-summing among passage features into a single feature, and use it for downstream tasks (Zoph and Knight, 2016; Garmash and Monz, 2016; Libovický and Helcl, 2017; Wang et al., 2018; Yan et al., 2018; Lebanoff et al., 2018; Celikyilmaz et al., 2018; Nishimura et al., 2018; Libovický et al., 2018; Li et al., 2018b; Nishida et al., 2019). Our processing mechanisms are similar to Garmash and Monz (2016), Firat et al. (2016), and Dong and Smith (2018). The information retrieval literature is primarily concerned with reformulating queries, by either selecting expansion terms from candidates as in pseudo-relevance feedback (Salton, 1971; Zhai and Lafferty, 2001; Xu and Croft, 1996; Metzler and Croft, 2007; Cao et al., 2008; Bernhard, 2010; Nogueira and Cho, 2017; Li et al., 2018a). Our task differs because there is no *supervision* unlike multi-lingual translation tasks where a single target translation is available given sources from multiple languages.

3 Method

3.1 Multi-Source Question Generator

Our Multi-Source Question Generator (MSQG) model introduces a mechanism to generate a common question given multiple documents. At *training* time, it employs a standard sequence-to-sequence (S2S) model using a large number of (single document, question) pairs. At *test* time, it generates a common question given multiple documents, similar to Garmash and Monz (2016) and Firat et al. (2016). Specifically, our MSQG model iterates over two interleaved steps, until an end-of-sentence (EOS) token is generated:

Distribute Step During the Distribute step, we take an instance of the trained S2S model, and perform inference with N different input documents. Each document is then encoded using one copy of the model to generate a unique target vocabulary distribution $\mathcal{P}_{i,t}^{\text{dec}}$ (for document i , at time t) for the next word. Note that source information comes from not only encoded latent representation from a source document, but also the cross-attention between source and generation.

Aggregate Step During the Aggregate step, we aggregate the N different target distributions into one distribution by averaging them as below:

$$\tilde{\mathcal{P}}_t^{\text{dec}} = \frac{1}{N}(\beta_1\mathcal{P}_{1,t}^{\text{dec}} + \beta_2\mathcal{P}_{2,t}^{\text{dec}} + \dots + \beta_N\mathcal{P}_{N,t}^{\text{dec}})$$

where $\tilde{\mathcal{P}}_t^{\text{dec}}$ is the final decoding distribution at time t , and $\sum_i^N \beta_i = N$. In our experiments, we weight all the decoding distributions equally ($\beta_i = 1$) to smooth out features that are distinct in each document i , where $i \in \{1, \dots, N\}$.

Note that the *average* Aggregate can be perceived as a *majority voting* scheme, in that each document-specific decoder will vote over the vocabulary and the final decision is made in a collaborative manner. We also experimented with different Aggregate functions: (i) $\text{MSQG}^{\text{mult}}$ multiplies the distributions, which is analogous to a *unanimous voting* scheme. However, it led to sub-optimal results since one unfavorable distribution can discourage decoding of certain common words. (ii) MSQG^{max} takes the maximum probability of each word across N distributions and normalizes them into a single distribution, but it could not generate sensible questions so we excluded from our pool of baselines.

3.2 Model Variants

Avoiding repetitive generation We observed that naively averaging the target distributions at every decoding time continually emphasized the common topic, thereby decoding repetitive topic words. To increase the diversity of generated tokens, we mask those tokens that have already been decoded in subsequent decoding steps. This strategy is reasonable for our task since questions generally tend to be short and rarely have repeated words. This mechanism can be viewed as a hard counterpart of the coverage models developed in Tu et al. (2016) and See et al. (2017). We denote this feature by *rmrep* in subscript.

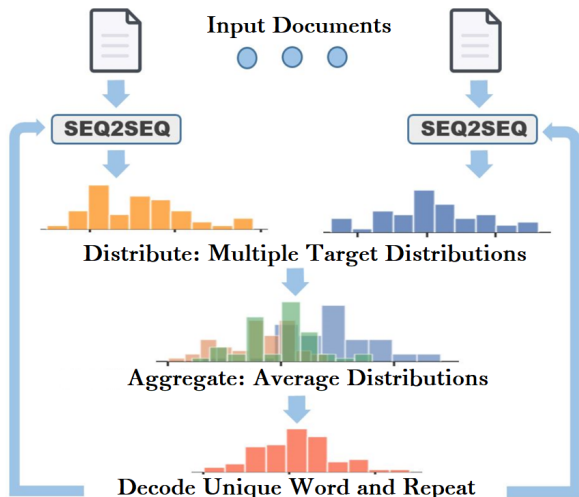


Figure 1: Multi-Source Question Generator (MSQG) model at test time. The simple architecture significantly outperforms the baselines for generating common questions, based on a number of metrics.

Shared encoder feature To initialize multiple decoders with the common meaning of the documents in a partition, we broadcast the *mean* of encoded latent representation to each decoder and denote this variant by the subscript *shared*. Note that the source document can affect the generated target vocabulary distribution $\mathcal{P}_{i,t}^{\text{dec}}$ at Distribute step through source-generation cross-attention.

4 Results

4.1 Experimental setup

Our training method uses the standard LSTM-based (Hochreiter and Schmidhuber, 1997) S2S with bi-linear attention (Luong et al., 2015). An input to our encoder is a concatenation of 100-dim GloVe (Pennington et al., 2014) vector, 100-dim predicate location vector, and 1024-dim ELMo (Peters et al., 2018) vector. Targets are embedded into 100-dim vectors. The S2S is bi-directional with a 256-dim bi-linear attention in each direction with ReLU (Nair and Hinton, 2010). Our encoder has two layers and we use an Adam (Kingma and Ba, 2014) with a learning rate of 2×10^{-5} .

4.2 Baselines

S2S We compare our model with a standard S2S baseline where we *concatenate* the N documents into a single document to generate a question. We provide detailed discussions about the effect of document order in supplementary material (SM). Two variants are considered (S2S and S2S_{rmrep}). Beam size is set to 5.

MESD We also compare our model with the multi-encoder single-decoder (MESD) baseline where documents are encoded individually into $\{v_i\}_{i=1}^N$. The single decoder’s initial hidden state is initialized by the mean of $\{v_i\}_{i=1}^N$, following (Dong and Smith, 2018).

4.3 Dataset

We use the Microsoft MACHine Reading COMprehension Question-Answering Dataset (MS-MARCO-QA) (Nguyen et al., 2016), where a single data instance consists of an anonymized Bing search query q and top-10 retrieved passages. Among the 10 passages, a passage is labelled *is-selected:True* if annotators used it, if any, to construct answers, and most instances contain one or two selected passages. For training S2S, we use a single selected passage $p^* \in \{p_1, p_2, \dots, p_{10}\}$ as input, and the query q as target output.

4.4 Constructing Evaluation Sets

For automatic evaluation, we follow the standard evaluation method from the MS-MARCO Re-Ranking task. For each generated question \tilde{q} , we construct an evaluation set that contains 100 passages in total.¹

First, using the 10-passage sets from the MS-MARCO-QA development dataset as inputs, we generate common questions with the baselines and our MSQG models, decoded for a maximum length of 25 words. A sample generation is provided in the SM. Secondly, we evaluate the generations by using the pre-trained BERT-based MS-MARCO passage re-ranker \mathcal{R} , which is publicly available and state-of-the-art as of April 1, 2019 (Nogueira and Cho, 2019). We assess whether the 10-passage set used to generate the question ranks higher than 90 other passages drawn from a pool of ~ 8.8 million MS-MARCO passages using the generated question. These 90 passages are retrieved via a different criterion: BM25 (Robertson and Zaragoza, 2009) using Lucene². Note that there are multiple 10-passage sets that generate the same question \tilde{q} . For each of these 10-passage sets, we construct a 100-passage evaluation set using the same 90 passages retrieved via the BM25 criterion.

¹A small number of evaluation sets had less than 100 passages because of duplicates between the source 10-passage set and the 90 passages retrieved via BM25.

²<https://lucene.apache.org/>

Fluency				Relevancy				Answerability			
	MSQG	S2S	Human		MSQG	S2S	Human		MSQG	S2S	Human
Completely grammatical	81.37%	71.28%	82.48%	Completely relevant	84.00%	71.00%	86.40%	Completely answered	71.89%	48.45%	72.79%
Comprehensible	16.58%	21.81%	16.35%	Somewhat relevant	7.06%	7.56%	6.69%	Somewhat answered	7.28%	5.83%	6.83%
Not comprehensible	2.05%	6.91%	1.17%	Not relevant	8.94%	21.44%	6.91%	Not answered	20.83%	45.72%	20.38%

Human judges preferred:				Human judges preferred:				Human judges preferred:						
Our Method	Neutral	Comparison		Our Method	Neutral	Comparison		Our Method	Neutral	Comparison				
MSQG	75.77%	9.74%	14.49%	S2S	MSQG	79.22%	8.06%	12.72%	S2S	MSQG	78.50%	9.09%	12.40%	S2S
MSQG	42.11%	10.66%	47.24%	Human	MSQG	40.81%	9.67%	49.52%	Human	MSQG	40.66%	10.26%	49.08%	Human

Table 1: Human evaluation of fluency, relevancy, and answerability. We used the top-ranked 30% of judges provided by a crowdsourcing service. Three judges performed each hit. Spammers were blocked at runtime. Agreement with most common was 81% overall. MSQG refers to MSQG_{sharedh,rmrep}. The upper table shows evaluations of individual models and the lower shows pairwise comparisons: (Human ↔ MSQG_{sharedh,rmrep}) and (MSQG_{sharedh,rmrep} ↔ S2S). Comparison results are significant at $p < 0.00001$.

Model	Retrieval Statistics			Unique \tilde{q} % dev.
	MRR	nDCG	MRR@10	
S2S	0.0520	0.2147	0.0266	70.6
S2S _{rmrep}	0.0540	0.2152	0.0284	80.4
MESD	0.0509	0.2141	0.0248	68.6
MEMD ^{mult}	0.0513	0.2142	0.0256	61.4
MEMD	0.0560	0.2209	0.0287	66.9
MSQG _{sharedh}	0.0569	0.2220	0.0298	67.0
MSQG _{sharedh,rmrep}	0.0704	0.2337	0.0441	70.3

Table 2: Our proposed model MSQG_{sharedh,rmrep} significantly outperforms baselines, based on the automated retrieval statistics. Discussion of the proportion of unique questions is dealt in supplementary material.

4.5 Evaluation Metrics

MRR, MRR@10, nDCG An input to the re-ranker \mathcal{R} is a concatenation of the generated question and one passage i.e. $[\tilde{q}, p]$. For each pair, it returns a score $\in (0, 1)$ where 1 denotes that the input passage is the most suitable for \tilde{q} . We score all 100 pairs in an evaluation set. For the source 10-passage set, we average the 10 scores into one score as one combined document and obtain the retrieval statistics MRR, MRR@10 (Voorhees, 2001; Radev et al., 2002), and nDCG (Järvelin and Kekäläinen, 2002) (see the SM for details).

Human Judgments We also conduct human evaluation where we compare questions generated by MSQG_{sharedh,rmrep} and the S2S baseline, and the reference question using three criteria: *fluency*, *relevancy*, and *answerability* to the original 10 passages. We randomly select 200 (10-passage, reference question) sets from which we generate questions, yielding 2,000 (passage, question) evaluation pairs for our model, baseline, and reference, respectively (see the SM for details).

4.6 Results

Table 3 shows the mean retrieval statistics and their proportion of unique generated questions

from 55,065 10-passage instances. Notice that our proposed MSQG models are more effective in terms of retrieving the source 10-passage sets. Particularly, MSQG_{sharedh,rmrep} outperforms the baselines in all metrics, indicating that broadcasting the mean of the document vectors to initialize the decoders (*sharedh*), and increasing the coverage of vocabulary (*rmrep*) are effective mechanisms for generating common questions.

Overall, the retrieval statistics are relatively low. Most 100 passages in the evaluation sets have high pair-wise cosine similarities. We computed similarities of passage pairs for a significant portion of the dataset until convergence. A random set of 10 passages has an average pair-wise similarity of 0.80, whereas the top-10 re-ranked passages have an average of 0.85 based on BERT (Devlin et al., 2018) embeddings. Given the small similarity margin, the retrieval task is challenging. Despite of low statistics, we obtained statistical significance based on MRR with $p < 0.00001$ between *all* model pairs (see the SM for details).

Human evaluation results are shown in Table 1. In the comparison tasks, our proposed model significantly outperforms the strong baseline by a large margin. Nevertheless, judges preferred the reference over our model on all three aspects. The individual tasks corroborate our observations.

5 Conclusion

We present a new task of generating common questions based on shared concepts among documents, and extensively evaluated multi-source encoder-decoder framework models, including our variant model MSQG applied to this new task. We also provide an empirical evaluation framework based on automated metrics and human judgments to evaluated multi-source generation framework for generating common questions.

References

- Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and Bruce Croft. 2019. Asking clarifying questions in open-domain information-seeking conversations. In *SIGIR '19*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. [Neural machine translation by jointly learning to align and translate](#). *CoRR*, abs/1409.0473.
- Delphine Bernhard. 2010. [Query expansion based on pseudo relevance feedback from definition clusters](#). In *Coling 2010: Posters*, pages 54–62, Beijing, China. Coling 2010 Organizing Committee.
- Christian Buck, Jannis Bulian, Massimiliano Ciaramita, Wojciech Gajewski, Andrea Gesmundo, Neil Houlsby, and Wei Wang. 2017. [Ask the right questions: Active question reformulation with reinforcement learning](#).
- Guihong Cao, Jian-Yun Nie, Jianfeng Gao, and Stephen Robertson. 2008. [Selecting good expansion terms for pseudo-relevance feedback](#). In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '08*, pages 243–250, New York, NY, USA. ACM.
- Asli Celikyilmaz, Antoine Bosselut, Xiaodong He, and Yejin Choi. 2018. [Deep communicating agents for abstractive summarization](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1662–1675, New Orleans, Louisiana. Association for Computational Linguistics.
- Eric Chu and Peter J. Liu. 2018. [Unsupervised neural multi-document abstractive summarization](#). *CoRR*, abs/1810.05739.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Rui Dong and David Smith. 2018. [Multi-input attention for unsupervised OCR correction](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2363–2372, Melbourne, Australia. Association for Computational Linguistics.
- Xinya Du, Junru Shao, and Claire Cardie. 2017. [Learning to ask: Neural question generation for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1342–1352, Vancouver, Canada. Association for Computational Linguistics.
- Nan Duan, Duyu Tang, Peng Chen, and Ming Zhou. 2017. [Question generation for question answering](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 866–874, Copenhagen, Denmark. Association for Computational Linguistics.
- Orhan Firat, Baskaran Sankaran, Yaser Al-Onaizan, Fatos T. Yarman-Vural, and Kyunghyun Cho. 2016. [Zero-resource translation with multi-lingual neural machine translation](#). *CoRR*, abs/1606.04164.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. [Neural approaches to conversational AI](#). *arXiv preprint arXiv:1809.08267*.
- Ekaterina Garmash and Christof Monz. 2016. [Ensemble learning for multi-source neural machine translation](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1409–1418, Osaka, Japan. The COLING 2016 Organizing Committee.
- Vrindavan Harrison and Marilyn Walker. 2018. [Neural generation of diverse questions using answer focus, contextual and linguistic features](#).
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Comput.*, 9(8):1735–1780.
- Kalervo Järvelin and Jaana Kekäläinen. 2002. [Cumulated gain-based evaluation of ir techniques](#). *ACM Trans. Inf. Syst.*, 20(4):422–446.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Igor Labutov, Sumit Basu, and Lucy Vanderwende. 2015. [Deep questions without deep understanding](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 889–898, Beijing, China. Association for Computational Linguistics.
- Logan Lebanoff, Kaiqiang Song, and Fei Liu. 2018. [Adapting the neural encoder-decoder framework from single to multi-document summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4131–4141, Brussels, Belgium. Association for Computational Linguistics.
- Canjia Li, Yingfei Sun, Ben He, Le Wang, Kai Hui, Andrew Yates, Le Sun, and Jungang Xu. 2018a. [NPRF: A neural pseudo relevance feedback framework for ad-hoc information retrieval](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4482–4491, Brussels, Belgium. Association for Computational Linguistics.

- Ruizhi Li, Xiaofei Wang, Sri Harish Reddy Mallidi, Takaaki Hori, Shinji Watanabe, and Hynek Herman-sky. 2018b. [Multi-encoder multi-resolution frame-work for end-to-end speech recognition](#). *CoRR*, abs/1811.04897.
- Jindřich Libovický and Jindřich Helcl. 2017. [Attention strategies for multi-source sequence-to-sequence learning](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 196–202, Vancouver, Canada. Association for Computational Lin-guistics.
- Jindřich Libovický, Jindřich Helcl, and David Mareček. 2018. [Input combination strategies for multi-source transformer decoder](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 253–260, Belgium, Brus-sels. Association for Computational Linguistics.
- Yi Luan, Yangfeng Ji, Hannaneh Hajishirzi, and Boyang Li. 2016. Multiplicative representations for unsupervised semantic role induction. In *ACL*.
- Thang Luong, Hieu Pham, and Christopher D. Man-ning. 2015. [Effective approaches to attention-based neural machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natu-ral Language Processing*, pages 1412–1421, Lis-bon, Portugal. Association for Computational Lin-guistics.
- Donald Metzler and W. Bruce Croft. 2007. [Latent con-cept expansion using markov random fields](#). In *Pro-ceedings of the 30th Annual International ACM SI-GIR Conference on Research and Development in Information Retrieval, SIGIR '07*, pages 311–318, New York, NY, USA. ACM.
- Nasrin Mostafazadeh, Chris Brockett, Bill Dolan, Michel Galley, Jianfeng Gao, Georgios Sp-ithourakis, and Lucy Vanderwende. 2017. [Image-grounded conversations: Multimodal context for natural question and response generation](#). In *Pro-ceedings of the Eighth International Joint Confer-ence on Natural Language Processing (Volume 1: Long Papers)*, pages 462–472, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Vinod Nair and Geoffrey E. Hinton. 2010. [Rectified linear units improve restricted boltzmann machines](#). In *Proceedings of the 27th International Conference on International Conference on Machine Learning, ICML'10*, pages 807–814, USA. Omnipress.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. [MS MARCO: A human generated machine reading comprehension dataset](#). In *Proceedings of the Workshop on Cognitive Computation: Inte-grating neural and symbolic approaches 2016 co-located with the 30th Annual Conference on Neu-ral Information Processing Systems (NIPS 2016), Barcelona, Spain, December 9, 2016*.
- Kyosuke Nishida, Itsumi Saito, Kosuke Nishida, Kazu-toshi Shinoda, Atsushi Otsuka, Hisako Asano, and Junji Tomita. 2019. [Multi-style generative reading comprehension](#). *CoRR*, abs/1901.02262.
- Yuta Nishimura, Katsuhito Sudoh, Graham Neubig, and Satoshi Nakamura. 2018. [Multi-source neural machine translation with missing data](#). In *Proceed-ings of the 2nd Workshop on Neural Machine Trans-lation and Generation*, pages 92–99, Melbourne, Australia. Association for Computational Lin-guistics.
- Rodrigo Nogueira and Kyunghyun Cho. 2017. [Task-oriented query reformulation with reinforcement learning](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Process-ing*, pages 574–583, Copenhagen, Denmark. Asso-ciation for Computational Linguistics.
- Rodrigo Nogueira and Kyunghyun Cho. 2019. [Passage re-ranking with BERT](#). *CoRR*, abs/1901.04085.
- Jeffrey Pennington, Richard Socher, and Christo-pher D. Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Lan-guage Processing*, pages 1532–1543.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word rep-resentations. In *Proc. of NAACL*.
- Dragomir R. Radev, Hong Qi, Harris Wu, and Weiguo Fan. 2002. [Evaluating web-based question answer-ing systems](#). In *Proceedings of the Third Interna-tional Conference on Language Resources and Eval-uation (LREC'02)*, Las Palmas, Canary Islands - Spain. European Language Resources Association (ELRA).
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: Bm25 and be-yond](#). *Found. Trends Inf. Retr.*, 3(4):333–389.
- A Rothe, Brenden Lake, and Todd Gureckis. 2016. Asking and evaluating natural language questions. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*.
- G. Salton. 1971. *The SMART Retrieval System—Experiments in Automatic Document Processing*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th An-nual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Compu-tational Linguistics.
- Iulian Vlad Serban, Alberto García-Durán, Caglar Gulcehre, Sungjin Ahn, Sarath Chandar, Aaron Courville, and Yoshua Bengio. 2016. [Generating](#)

- factoid questions with recurrent neural networks: The 30M factoid question-answer corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 588–598, Berlin, Germany. Association for Computational Linguistics.
- Linfeng Song, Zhiguo Wang, and Wael Hamza. 2017. [A unified query-based generative model for question generation and question answering](#).
- Linfeng Song, Zhiguo Wang, Wael Hamza, Yue Zhang, and Daniel Gildea. 2018. [Leveraging context information for natural question generation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 569–574, New Orleans, Louisiana. Association for Computational Linguistics.
- Xingwu Sun, Jing Liu, Yajuan Lyu, Wei He, Yanjun Ma, and Shi Wang. 2018. [Answer-focused and position-aware neural question generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3930–3939, Brussels, Belgium. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*.
- Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. [Modeling coverage for neural machine translation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 76–85, Berlin, Germany. Association for Computational Linguistics.
- Ellen M. Voorhees. 2001. [The trec question answering track](#). *Nat. Lang. Eng.*, 7(4):361–378.
- Yizhong Wang, Kai Liu, Jing Liu, Wei He, Yajuan Lyu, Hua Wu, Sujian Li, and Haifeng Wang. 2018. [Multi-passage machine reading comprehension with cross-passage answer verification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1918–1927, Melbourne, Australia. Association for Computational Linguistics.
- Jinxi Xu and W. Bruce Croft. 1996. [Query expansion using local and global document analysis](#). In *Proceedings of the 19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '96*, pages 4–11, New York, NY, USA. ACM.
- Ming Yan, Jiangnan Xia, Chen Wu, Bin Bi, Zhongzhou Zhao, Ji Zhang, Luo Si, Rui Wang, Wei Wang, and Haiqing Chen. 2018. [A deep cascade model for multi-document reading comprehension](#). *CoRR*, abs/1811.11374.
- Chengxiang Zhai and John Lafferty. 2001. [Model-based feedback in the language modeling approach to information retrieval](#). In *Proceedings of the Tenth International Conference on Information and Knowledge Management, CIKM '01*, pages 403–410, New York, NY, USA. ACM.
- Barret Zoph and Kevin Knight. 2016. [Multi-source neural translation](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 30–34, San Diego, California. Association for Computational Linguistics.

Model	Retrieval Statistics			Unique \tilde{q} % dev.
	MRR	nDCG	MRR@10	
S2S- \mathcal{M}_{256}	0.0368	0.1943	0.0147	66.3
S2S- \mathcal{M}_{512}	0.0393	0.1980	0.0165	71.1
S2S	0.0520	0.2147	0.0266	70.6
S2S _{rmrep}	0.0540	0.2152	0.0284	80.4
MESD- \mathcal{M}_{512}	0.0367	0.1941	0.0147	72.6
MESD	0.0509	0.2141	0.0248	68.6
MSQG- \mathcal{M}_{512}	0.0450	0.2056	0.0210	72.0
MSQG ^{mult}	0.0513	0.2142	0.0256	61.4
MSQG	0.0560	0.2209	0.0287	66.9
MSQG _{sharedh}	0.0569	0.2220	0.0298	67.0
MSQG _{sharedh,rmrep}	0.0704	0.2337	0.0441	70.3

Table 3: Full results, comparing models constructed with \mathcal{M}_{256} , \mathcal{M}_{512} , and \mathcal{M}_{256}^{attn} . \mathcal{M}_{512} has the most number of parameters among the three considered.

Supplementary Material

A Full Experiment Results

Table 3 shows the retrieval results on a larger set of baselines and MSQG models. \mathcal{M}_{256}^{attn} is an attention-based encoder-decoder with hidden size 256 for both encoder and decoder. \mathcal{M}_{256} and \mathcal{M}_{512} are non-attention encoder-decoders with hidden sizes 256 and 512. S2S denotes \mathcal{M}_{256}^{attn} , as in the main paper. It shows that models constructed using \mathcal{M}_{256}^{attn} are more effective as opposed to models using \mathcal{M}_{512} which has more parameters. Furthermore, we see that the averaging scheme in the *Reduction* step, broadcasting the same encoder mean, and increasing coverage of vocabulary tokens are important features to generating common questions using MSQG models.

B Effect of Document Order on S2S

To examine if the order of multiple input documents are critical for S2S, we obtain the attention weights at each decoding time, gathered across the development dataset. Next, we perform a simple ordinary least squares regression, where the predictors are indexed word positions in a concatenated input, and responses are assumed noisy attention weights over the development dataset for each word position.

The slope coefficient fell within the 95% confidence interval that includes the null: $[-2.75 \times 10^{-5}, 3.03 \times 10^{-5}]$ and a statistically significant intercept value of 0.0021. The result also validates that an average 10-passage string is approximately 476 ($\approx \frac{1}{0.0021}$) words long. Thus, we conclude that the attention weights are evenly distributed across multiple document at test time, and the document ordering is not critical to the




Figure 2: Agglomerative Clustering of 55,065 source 10-passage sets. Each set is represented by the mean of 10 BERT embeddings. Both max and average linkages yield the same inflection point at 0.0326, corresponding to 35,928 and 32,871 clusters. This method implies that the target proportion of unique generations should be at least 65% or 60%, which all models but MSQG^{mult} achieve.

performance of S2S.

C Clustering Duplicate 10-passage Sets

In the MS-MARCO-QA dataset, there are many highly similar 10-passage sets retrieved from semantically close MS-MARCO queries. Examples of semantically close MS-MARCO queries include [“*symptoms blood sugar is low*”, “*low blood sugar symptoms*”, “*symptoms of low blood sugar levels*”, “*signs and symptoms of low blood sugar*”, “*what symptoms from low blood sugar*”, ...], from which we expect duplicate generated questions, thus in sum, less than 55,065 different questions.

Therefore, to estimate the *target* proportion of unique generations, we examine the number of semantically similar 10-passage sets through agglomerative clustering. Figure 2 shows cluster results with varying degrees of affinity thresholds, and observe that the effective models should generate at least 65% unique questions from the development dataset. This, together with the low retrieval statistics of MSQG^{mult}, implies that multiplying the distributions is not an appropriate *Reduction* step.

On the other hand, generating the most number of unique questions does not imply that the model better generates common questions. In particular, S2S_{rmrep} generates the most diverse questions, however, its retrieval statistics are significantly lower than its MSQG counterparts.

D Statistical Significance Tests

Retrieval evaluation on $\sim 55K$ evaluation sets using the re-ranker \mathcal{R} is compute-intensive. Thus, for each model, we randomly sample and obtain retrieval statistics for 15K evaluation sets which are enough to mimic the true evaluation set distribution.

Then, to assess statistical significance, we use a non-parametric two-sample test, such as Mann-Whitney (MW) or Kolmogorov-Smirnov statistic, and test whether any pair of 15K retrieval sets between two models come from the same distribution. In our task, both tests reached the same conclusion. MW two-sample tests on MRR results showed statistical significance at $p < 0.00001$ for all model pairs dealt in the main paper, in spite of the relatively low retrieval statistics.

E Human Evaluation Templates

UHRS comparison and individual task instructions and shown in the next pages.

F Generated Questions Sample

Passage 1: *cucumbers and zucchini look similar but have nutritional differences . photo credit martin poole / digital vision / getty images . do n't let the similarities between cucumbers and zucchini confuse you . even though both cylindrical vegetables are dark green with white flesh , they are distinctively different species . both cucumbers and zucchini belong to the curcubit family , which also counts gourds , melons , pumpkins and squash among its members . cucumbers and zucchini differ both in how people commonly eat them and in their nutritional values . people almost always eat cukes raw , while zucchini is more often cooked .*

Passage 2: *cucumber and squash seedlings both have elongated foliage for the first set of leaves after they emerge from the soil . the second set of leaves on a seedling varies . cucumber leaves are in the shape of a triangle and are flat in the center and rough to the touch . squash plants vary in shape as to the particular variety , but have three to five lobes and are larger than cucumber leaves . zucchini squash has elongated serrated leaves .*

Passage 3: *zucchini vs cucumber . zucchini*

and cucumber are two vegetables that look mightily similar and hard to distinguish from each other . but in close inspection , they are actually very different . so read on . zucchini . zucchini is defined to be the kind of vegetable that is long , green colored and has many seeds .

Passage 4: *as a general rule , we prefer cucumbers raw and zucchini cooked . while you ca n't replace one with the other , zucchini and cucumbers do complement one another . slice two cucumbers , two zucchinis and one sweet onion , and soak them all in rice vinegar for at least an hour in the refrigerator .*

Passage 5: *cucumber and zucchini are popular vegetables that are similar in appearance and botanical classification . but they differ significantly in taste , texture and culinary application . zucchini and cucumber are both members of the botanical family cucurbitaceae , which includes melons , squashes and gourds .*

Passage 6: *melon vs. squash . the cucumber is not particularly sweet , but it shares a genus with the cantaloupe and is botanically classified as a melon . the zucchini is a variety of summer squash and is of the same species as crookneck squash .*

Passage 7: *cucumber vs. zucchini . side by side , they might fool you : cucumbers and zucchinis share the same dark green skin , pale seedy flesh , and long cylindrical shape . to the touch , however , these near - twins are not the same : cucumbers are cold and waxy , while zucchinis are rough and dry . the two vegetables also perform very differently when cooked .*

Passage 8: *the second set of squash leaves grow much quicker and larger than cucumber leaves in the same time . squash leaves may be up to four times as large as a cucumber leaf when they are the same age .*

Passage 9: *in reality , zucchini is really defined as a vegetable so when it comes to the preparation of it , it has different temperament . cucumber . cucumber is both classified as a fruit and a vegetable . it is long and is green in color , too . it is part of what they call the gourd family .*

Passage 10: *zucchini 's flowers are edible ; cucumber 's flowers are not . zucchini is generally considered as a vegetable ; cucumber is classified as both a fruit and a vegetable . yes , they can fool the eye because of their similar look but as you go deeper , they are very different in so many ways .*

Question generated by MSQG_{sharedh,rmrep}:
what are the difference between cucumber and zucchini

Question generated by S2S:
different types of zucchini

Reference question:
difference between cucumber and zucchini

Title: Compare two questions

Instructions

The goal of this task is to understand which of the two given questions is better when evaluated on the criteria of answerability, fluency and relevancy to the given passage.

You will be shown a passage and two questions (Q1 & Q2).

Your task is to answer the following:

Answerability: Which of the two questions is better answered in the passage?

(Note: Ignore minor errors and disfluencies in English here.)

Question A No preference Question B

Fluency: Which of the two questions is more fluent/grammatical English?

Question A No preference Question B

Relevance: Which of the two questions is more relevant to the passage?

(Note: a question can be relevant even if it is not answerable in the passage.)

Question A No preference Question B

Example 1

Passage:

egg prices soar across india nov 20 , 2017 , 19:29 ist . after vegetables , now it 's eggs that will burn a hole in your pocket . the prices of eggs have reached a record high across india . hit by tight supply in the many states , egg prices in the country have increased by up to 40 per cent . egg prices in retail markets in the national capital are ruling at rs 7 - 7.50 per piece , up from rs 4 - 5 last year , according to trade data .

Question A: *average price of egg in india*

Question B: *how much does it cost to raise a egg egg*

Answerability: Question A is much better answered than Question B

Reason: Question A is answered in the passage whereas Question B is only relevant to passage (but not answered in the passage).

Fluency: Question A is much more fluent than Question B

Reason: Because Question A is grammatical whereas Question B contains grammatical errors.

Relevancy: Question A is much more relevant than Question B

Reason: Because Question A is answered in the passage whereas the question Question B does not make sense.

Title: Rate how good is the question on basis of answerability, fluency and relevancy.

Instructions

The goal of this task is to how good is the question on basis of answerability, fluency and relevancy to a given passage.

You will be shown a passage and a question asked to answer the following:

- a. How well is the question answered in the given passage?
Choose from the following:
 - Question is completely answered in the passage
 - Question is somewhat answered in the passage
 - Question is not answered in the passage

- b. How fluent is the question?
Choose from the following:
 - Question is completely grammatical
 - Question is comprehensible but includes a few grammatical errors
 - Question is not comprehensible

- c. How relevant is the question to the given passage?
Choose from the following:
 - Question is completely relevant to the passage
 - Question is somewhat relevant to the passage
 - Question is not relevant to the passage

Example 1

Passage:

wells fargo center philadelphia parking lots location , transportation options , parking info , and seat maps available here . where is wells fargo center philadelphia parking lots in philadelphia ? find out at stubhub .

Question: *how much is parking by wells fargo Philadelphia*

Answerability: Question is somewhat answered in the passage

Reason: Although the passage does not directly include the parking cost, it points to a website that could contain this information. Hence the question is somewhat answered in the passage.

Fluency: Question is comprehensible but includes a few grammatical errors

Reason: A fluent version of the question would look something like "how much *would* parking by wells fargo philadelphia *cost*"

Relevancy: Question is completely relevant to the passage

Reason: Since the question is asking about parking by Wells Fargo and the passage is talking about the same.