

# Neural document expansion for ad-hoc information retrieval

Cheng Tang <sup>\*\*\*</sup>, Andrew Arnold <sup>†</sup>

Amazon

**Abstract.** Recently, Nogueira et al. [2019] proposed a new approach to document expansion based on a neural Seq2Seq model, showing significant improvement on short text retrieval task. However, this approach needs a large amount of in-domain training data. In this paper, we show that this neural document expansion approach can be effectively adapted to standard IR tasks, where labels are scarce and many long documents are present.

## 1 Introduction

In traditional ad-hoc retrieval, queries and documents are represented by variants of bag-of-words representations. This leads to the so called vocabulary mismatch problem: when a query contains words that do not exactly match words in a relevant document, the search engine may fail to retrieve this document. Query expansion and document expansion, the methods of adding additional terms to the original query or document, are two popular solution to alleviate the vocabulary mismatch problem.

Document expansion has been shown to be particularly effective for short text retrieval and language-model based retrieval [Tao et al., 2006, Efron et al., 2012]. Most of the existing works in document expansion are unsupervised: using information from the corpus to augment document representation, e.g., retrieval based [Efron et al., 2012] and clustering based [Liu and Croft, 2004, Tao et al., 2006], or using external information to augment document representation [Agirre et al., 2010, Sherman and Efron, 2017].

Recently, Nogueira et al. [2019] proposed a new approach to document expansion, which is based on a popular generative sequence-to-sequence model (Seq2Seq) in NLP, transformers [Wolf et al., 2020]. It leverages supervision to train the model to predict expansion terms conditional on each document. The paper has shown significant improvement on passage (short document) datasets, when trained in-domain. In this paper, we follow this line of supervised neural document expansion approach and explore its performance on standard IR benchmarking dataset. Our main contributions are: 1. Adapting the method to unlabeled datasets by exploring transfer learning and weak-supervision approaches. 2. Adapting the method to traditional IR datasets, where a large number of long documents are present.

---

<sup>\*\*\*</sup> tcheng@amazon.com

<sup>†</sup> anarnd@amazon.com




Fig. 1: Left: document lengths distribution in Robust04; Right: distribution of relevant passage (answer-containing passage) indices in MSMarco-documents

## 2 Document retrieval with Seq2Seq models

In this section, we discuss our approach to document expansion.

**Train a passage level document expansion model** We follow the approach of Nogueira et al. [2019]: we use data of format  $(p, q)$ , where  $q$  is a question and  $p$  is a passage *relevant* to it. A Seq2Seq model is trained so that conditional on input  $p$ , it learns to generate the ground-truth question  $q$ . The dataset we found most effective for training the document expansion model is the MSMarco passage-ranking dataset [Bajaj et al., 2018].

**Domain adaptation** Since most IR datasets are short of annotated queries, it is important to explore ways to adapt the document expansion model to out-of-domain unlabeled datasets. We explored two simple approaches: zero-shot transfer and weak supervision (retrieval-based pseudo annotation). In zero-shot transfer, we directly apply a model trained from one domain to another; in retrieval-based pseudo annotation, we issue out-of-domain queries to a target corpus and treat the top- $k$  returned documents as “relevant”.

**Adaptation to long documents** Due to the memory overhead of transformer-based models, the maximal input length of this family of models are constrained to be short. We trained our expansion models on passage-level datasets. However, during inference time, a standard IR dataset is usually a mixture of short and long documents, and expanding on an arbitrarily capped portion of a document may yield sub-optimal performance. Figure 1-Left shows the document lengths distribution of Robust04 dataset. We see that most of the documents are around 1000-tokens long, which has exceeded the typical input length transformer based models can take [Hofstätter et al., 2020]. To deal with this problem, we explored three strategies (see experiment results at Table 3: “Long document generation methods” section):

1. **concatenation of paragraph generation (CONCAT in Table 3)**: Split each long document into overlapping passages, respecting natural sentence

boundaries (with sliding window size to be around 60 tokens, since this is roughly the median size of passages in MSMarco dataset, which we used for training) and run expansion inference on each passage. All generated expansions are directly appended to the original documents.

2. **first k sentences (FIRST-K in Table 3)**: Run expansion model on the first  $k$  whole sentences of each document. This strategy is based on our analysis of MSMarco datasets. From MSMarco-QA dataset, we can obtain  $(q, p, a)$  triplets so that we know passage  $p$  contains an answer for question  $q$ . We treat this as a signal of question-passage relevance. Then we can trace these passages back to the document they belong to using the MSMarco-documents dataset. Comparing “Capped passage generation methods” to “Long document generation methods” in Table 3, we see that the “Long document generation methods”, and in particular, the CONCAT and FIRST-K generation methods are more effective.
3. **passage-importance indicator (PI in Table 3)**: Before seeing any questions, we predict which passage will be queried using a unsupervised predictor [Bendersky and Kurland, 2008], and generate expansion exclusively based on this passage. In our experiments, this method do not yield good performance, likely due to the low quality of the passage-importance predictor.

### 3 Experiments

To train the document expansion models, we selected datasets that are typically used for passage-ranking and question-answering tasks, i.e., MSMarco passage-ranking [Bajaj et al., 2018], Natural Questions [Kwiatkowski et al., 2019], and SQuAD [Rajpurkar et al., 2018]. In addition, we also used a TREC collection short-document dataset, microblogs [Lin et al., 2014]. To evaluate the performance of document expansion method on information retrieval, we additionally add the standard IR benchmarking dataset Robust04 [Voorhees, 2004]. Our baseline retrieval methods, BM25 and Query Likelihood (QL), use implementation from the open-source Anserini toolkit [Yang et al., 2017]; for Seq2Seq models, we experimented with transformers (OpenNMT implementation [Klein et al., 2017]) and pre-trained BART (Huggingface Pytorch implementation [Wolf et al., 2020]).

#### 3.1 Retrieval performance on passage datasets

From Table 1, we see that retrieval-based weak supervision, while showing good performance as training signal for neural retrieval models [Dehghani et al., 2017], does not yield good document expansion models. Instead, the BART-based zero-shot transfer model has competitive performance to in-domain trained-from-scratch models. Once decided on the zero-shot transfer learning approach, we tried several fine-tuning strategies with the BART model (Table 2), drawing inspiration from [Yilmaz et al., 2019]. We found that fine-tuning pre-trained

BART with MSMarco-passages dataset and with a mixed MSMarco-passages and microblogs dataset produces the best document expansion models. Although less effective, our experiment suggests that other passage-level datasets such as Natural Questions and SQuAD, can also be used as sources to train document expansion models.

DE models	MSMarco-passages		Natural Questions		SQuAD-v2	
	Trec-DL (DEV)	Trec-DL (DEV)	R@100	MAP	R@100	MAP
	R@100	MAP	R@100	MAP	R@100	MAP
BM25	0.4531 (0.7619)	0.3773 <b>(0.3616)</b>	0.7619	0.3616	0.9737	0.7997
<b>In-domain trained transformer</b>	0.4545 (0.7127)	0.3872 (0.2203)	<b>0.8671</b>	0.4450	0.9754	0.7915
<b>Weakly supervised transformer</b>	NA	NA	0.7649	0.3608	0.9717	0.7913
<b>Zero-shot transfer</b> (transformer) (fine-tuning BART)	NA	NA	0.7773	0.3879	0.9764	0.8056
	<b>(0.5297)</b>	<b>(0.465)</b>	0.8302	<b>0.4501</b>	<b>0.9787</b>	<b>0.8121</b>
			<b>(0.7949)</b>	<b>(0.2674)</b>		

Table 1: In-domain trained and weakly-supervised document expansion model; for MSMarco-passages, we have two test sets: DEV and Trec-DL [Craswell et al., 2020]

DE models	Natural Questions		SQuAD-v2		Robust04	
	R@100	MAP	R@100	MAP	R@100	MAP
BM25	0.7619	0.3616	0.9737	0.7997	0.4152	0.2553
MSMarco-passages	<b>0.8302</b>	<b>0.4501</b>	<b>0.9787</b>	0.8121	<b>0.4229</b>	0.2620
MSMarco-passage → microblogs	0.7931	0.4	0.9757	0.7962	0.4206	0.2533
MSMarco-passages + microblogs	0.8239	0.4437	<b>0.9787</b>	<b>0.8133</b>	0.4212	<b>0.2630</b>
Natural Questions	<b>0.9031</b>	<b>0.5271</b>	0.9782	0.8099	0.4190	0.2626
SQuAD-v2	0.8173	0.4228	<b>0.9798</b>	<b>0.8156</b>	0.4218	0.2616

Table 2: Zero-shot transfer of passage-level DE model

### 3.2 Retrieval performance on Robust04

To test the performance of passage-level trained document expansion model on standard IR corpus, we fixed the passage-level model to be “MSMaroc-passage + microblog” trained. Then we explored three expansion generation methods as

described in Section 2. The result of applying passage-level trained document expansion model to Robust04 can be found in Table 3.

**Traditional IR baselines** In addition to tf-idf based BM25, we applied document expansion to two popular language-model based retrieval methods: query-likelihood with Dirichlet Smooth (QLD) and query-likelihood with Jelinek-Mercer smoothing (QLJM). We found that document expansion has good performance with QLD. We speculate that this is because the way our current document expansion model works is similar to the Dirichlet smoothing model, which models the process of adding unseen words to the original documents by drawing from a background distribution based on the entire corpus. Here, the document expansion model in addition samples words from the query distribution (conditional on each document). Since document expansion does not significantly improve QLJM, we did not include it in the rest of our experiments.

**Capped passage generation vs Long document generation** Comparing “Capped passage generation methods” to “Long document generation methods”, we see that for both BM25 and QLD, the two long document expansion methods, CONCAT and FIRST-K, have better performance. The fact the CONCAT has the best performance suggests that even if most queries are targeting the head portion of a document (Figure 1-Right), using information from the entire document may still be beneficial to expansion generation.

**Performance with pseudo-relevance feedback and BERT re-ranker** Since document expansion is one of several techniques that can improve the document retrieval performance, we also want to understand how it works when combined with other techniques. In our experiments, we first explored combining the best performing document expansion model with RM3 [Abdul-jaleel et al., 2004], a popular query expansion method based on pseudo-relevance feedback. While RM3 alone significantly improves the baseline retrieval performance, DE\*\*<sup>1</sup> can still add an additional boost on top. We want to point out that comparing to RM3, which requires two rounds of retrieval at query time, document expansion model is applied offline and does not add additional computational overhead at run time. BM25 with document and query expansion makes a first-stage ranker in the ranking pipeline. In our last experiment, we test its end-to-end performance when combined with a second-stage neural ranker, BERT [Nogueira and Cho, 2019]. To evaluate the end-to-end result, we used metrics  $R@k$  and  $P@k$  for small  $k$ , mimicking what the ranking system presents to a user (the top- $k$  ranked documents). Our experiment results indicate that our document expansion models are complementary to query expansion as a first-stage ranker and can improve the end-to-end ranking performance when combined with a second-stage ranker.

---

<sup>1</sup> DE\*\* indicates concatenation of paragraph generation with MSMarco+microblog trained passage model.

Experiment category	Methods	Robust04	
		R@100	MAP
Traditional IR baselines	BM25	0.4152	0.2553
	QLD	0.4157	0.2502
	QLJM	0.3995	0.2287
Capped passage generation methods	BM25+passage-DE*	0.4212	<b>0.2630</b>
	QLD+passage-DE*	<b>0.4270</b>	0.2620
	QLJM+passage-DE*	0.4058	0.2350
Long document generation methods	BM25+DE* (CONCAT)	<b>0.4283</b>	<b>0.2631</b>
	BM25+DE* (FIRST-K)	0.4226	0.2625
	BM25+DE* (PI)	0.4212	0.2588
	QLD+DE* (CONCAT)	<b>0.4290</b>	0.2615
	QLD+DE* (FIRST-K)	0.4272	<b>0.2625</b>
	QLD+DE* (PI)	0.4259	0.2577
DE+pseudo-relevance feedback	BM25+RM3	0.4517	0.2941
	BM25+RM3+DE**	<b>0.4641</b>	<b>0.3035</b>
DE + BERT reranker	End-to-end metrics		
	R@10	P@10	P@5
BM25 + BERT	0.2771	0.3731	0.4137
BM25 + RM3 + BERT	0.2608	0.3803	0.4048
BM25 + DE** + BERT	<b>0.2824</b>	0.3767	0.4161
BM25 + RM3 + DE** + BERT	0.2726	<b>0.3944</b>	<b>0.4241</b>

Table 3: Robust04 experiments (DE\*: MSMarco-passage+microblog trained passage model; DE\*\*: concatenation of paragraph generation, CONCAT, with DE\*)

## 4 Conclusion

We showed that a document expansion model trained on passage-level datasets of (question, relevant passage) pairs can be directly applied to out-of-domain IR datasets to improve retrieval performance. We explored simple approaches to adapt the model to standard IR dataset (Robust04) where a large number of long documents are present, and we showed that adapting the passage-level trained model to long documents further improves retrieval performance. However, our current simple adaptations to long documents do not significantly improve the model performance (see Table 3, Long document generation methods). We cannot conclude whether this is due to the nature of the relevant passage distribution over long documents (i.e., they tend to be the first few passages of any document, according to Figure 1-Right). Hence, it may be worth exploring model architectures that allow longer input sequences, for example, by switching to sparse attention layers [Beltagy et al., 2020].

## Bibliography

- Rodrigo Nogueira, Wei Yang, Jimmy Lin, and Kyunghyun Cho. Document expansion by query prediction, 2019.
- Tao Tao, Xuanhui Wang, Qiaozhu Mei, and ChengXiang Zhai. Language model information retrieval with document expansion. In *Proceedings of the Main Conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics, HLT-NAACL '06*, page 407–414, USA, 2006. Association for Computational Linguistics. <https://doi.org/10.3115/1220835.1220887>. URL <https://doi.org/10.3115/1220835.1220887>.
- Miles Efron, Peter Organisciak, and Katrina Fenlon. Improving retrieval of short texts through document expansion. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '12*, page 911–920, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450314725. <https://doi.org/10.1145/2348283.2348405>. URL <https://doi.org/10.1145/2348283.2348405>.
- Xiaoyong Liu and W. Bruce Croft. Cluster-based retrieval using language models. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '04*, page 186–193, New York, NY, USA, 2004. Association for Computing Machinery. ISBN 1581138814. <https://doi.org/10.1145/1008992.1009026>. URL <https://doi.org/10.1145/1008992.1009026>.
- Eneko Agirre, Xabier Arregi, and Arantxa Otegi. Document expansion based on wordnet for robust ir. In *COLING*, 2010.
- Garrick Sherman and Miles Efron. Document expansion using external collections. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '17*, page 1045–1048, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450350228. <https://doi.org/10.1145/3077136.3080716>. URL <https://doi.org/10.1145/3077136.3080716>.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Huggingface’s transformers: State-of-the-art natural language processing, 2020.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. Ms marco: A human generated machine reading comprehension dataset, 2018.
- Sebastian Hofstätter, Hamed Zamani, Bhaskar Mitra, Nick Craswell, and Allan Hanbury. Local self-attention over long text for efficient document re-

- trieval. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '20, page 2021–2024, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450380164. <https://doi.org/10.1145/3397271.3401224>. URL <https://doi.org/10.1145/3397271.3401224>.
- Michael Bendersky and Oren Kurland. Utilizing passage-based language models for document retrieval. In *Proceedings of the IR Research, 30th European Conference on Advances in Information Retrieval*, ECIR'08, page 162–174, Berlin, Heidelberg, 2008. Springer-Verlag. ISBN 3540786457.
- 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. Natural questions: a benchmark for question answering research. *Transactions of the Association of Computational Linguistics*, 2019.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. Know what you don't know: Unanswerable questions for SQuAD. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia, July 2018. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-2124>. URL <https://www.aclweb.org/anthology/P18-2124>.
- J. Lin, Y. Wang, M. Efron, and Garrick Sherman. Overview of the trec-2014 microblog track. In *TREC*, 2014.
- E. Voorhees. Overview of the trec 2004 robust track. 2004.
- Peilin Yang, Hui Fang, and Jimmy Lin. Anserini: Enabling the use of lucene for information retrieval research. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '17, page 1253–1256, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450350228. <https://doi.org/10.1145/3077136.3080721>. URL <https://doi.org/10.1145/3077136.3080721>.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. OpenNMT: Open-source toolkit for neural machine translation. In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada, July 2017. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/P17-4012>.
- Mostafa Dehghani, Hamed Zamani, Aliaksei Severyn, Jaap Kamps, and W. Bruce Croft. Neural ranking models with weak supervision. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '17, page 65–74, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450350228. <https://doi.org/10.1145/3077136.3080832>. URL <https://doi.org/10.1145/3077136.3080832>.
- Zeynep Akkalyoncu Yilmaz, W. Yang, Haotian Zhang, and Jimmy Lin. Cross-domain modeling of sentence-level evidence for document retrieval. In *EMNLP/IJCNLP*, 2019.

- Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M. Voorhees. Overview of the trec 2019 deep learning track, 2020.
- Nasreen Abdul-jaleel, James Allan, W. Bruce Croft, O Diaz, Leah Larkey, Xiaoyan Li, Mark D. Smucker, and Courtney Wade. Umass at trec 2004: Novelty and hard. In *In Proceedings of TREC-13*, 2004.
- Rodrigo Nogueira and Kyunghyun Cho. Passage re-ranking with bert. *arXiv preprint arXiv:1901.04085*, 2019.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. Longformer: The long-document transformer, 2020.