

Simple Entity-Centric Questions Challenge Dense Retrievers

Christopher Sciavolino* Zexuan Zhong* Jinhyuk Lee Danqi Chen

Department of Computer Science, Princeton University

{cds3, zzhong, jinhyuklee, danqic}@cs.princeton.edu

Abstract

Open-domain question answering has exploded in popularity recently due to the success of dense retrieval models, which have surpassed sparse models using only a few supervised training examples. However, in this paper, we demonstrate current dense models are not yet the holy grail of retrieval. We first construct EntityQuestions, a set of simple, entity-rich questions based on facts from Wikidata (e.g., “Where was Arve Furset born?”), and observe that dense retrievers drastically underperform sparse methods. We investigate this issue and uncover that dense retrievers can only generalize to common entities unless the question pattern is explicitly observed during training. We discuss two simple solutions towards addressing this critical problem. First, we demonstrate that data augmentation is unable to fix the generalization problem. Second, we argue a more robust passage encoder helps facilitate better question adaptation using specialized question encoders. We hope our work can shed light on the challenges in creating a robust, universal dense retriever that works well across different input distributions.¹

1 Introduction

Recent dense passage retrievers outperform traditional sparse retrieval methods like TF-IDF and BM25 (Robertson and Zaragoza, 2009) by a large margin on popular question answering datasets (Lee et al. 2019, Guu et al. 2020, Karpukhin et al. 2020, Xiong et al. 2021). These dense models are trained using supervised datasets and the dense passage retriever (DPR) model (Karpukhin et al., 2020) demonstrates that only training 1,000 supervised examples on top of BERT (Devlin et al., 2019) already outperforms BM25, making it very appealing in practical use.

*The first two authors contributed equally.

¹Our dataset and code are publicly available at <https://github.com/princeton-nlp/EntityQuestions>.

	DPR (NQ)	DPR (multi)	BM25 -
Natural Questions	80.1	79.4	64.5
EntityQuestions (this work)	49.7	56.7	71.2
What is the capital of [E]?	77.3	78.9	90.0
Who is [E] married to?	35.6	48.1	85.9
Where is the headquarter of [E]?	70.0	72.0	85.0
Where was [E] born?	25.4	41.8	75.2
Where was [E] educated?	26.4	41.8	73.0
Who was [E] created by?	54.1	57.7	71.7
Who is [E]’s child?	19.2	33.8	82.9
(17 more types of questions)

Table 1: Top-20 retrieval accuracy for dense and sparse retrieval models on Natural Questions (Kwiatkowski et al., 2019) and our EntityQuestions along with a set of sampled questions (full results in Appendix A). We test two DPR models: (1) trained on NQ only; (2) trained on 4 datasets (NQ, TQA, WebQ, TREC) combined.²

In this work, we argue that dense retrieval models are not yet robust enough to replace sparse methods, and investigate some of the key shortcomings dense retrievers still face.

We first construct EntityQuestions, an evaluation benchmark of simple, entity-centric questions like “Where was Arve Furset born?”, and show dense retrieval methods generalize very poorly. As shown in Table 1, a DPR model trained on either a single dataset Natural Questions (NQ) (Kwiatkowski et al., 2019) or a combination of common QA datasets drastically underperforms the sparse BM25 baseline (49.7% vs 71.2% on average), with the gap on some question patterns reaching 60% absolute!

Based on these results, we perform a deep dive into why a single dense model performs so poorly on these simple questions. We decouple the two distinct aspects of these questions: the entities and the question pattern, and identify what about these

²NQ: Natural Questions (Kwiatkowski et al., 2019), TQA: TriviaQA (Joshi et al., 2017), WebQ: WebQuestions (Berant et al., 2013), CuratedTREC (Baudiš and Šedivý, 2015). We use the pre-trained models from <https://github.com/facebookresearch/DPR>.

questions gives dense models such a hard time. We discover the dense model is only able to successfully answer questions based on common entities, quickly degrading on rarer entities. We also observe that dense models can generalize to unseen entities only when the question pattern is explicitly observed during training.

We end with two investigations of practical solutions towards addressing this crucial problem. First, we consider data augmentation and analyze the trade-off between single- and multi-task fine-tuning. Second, we consider a single fixed passage index and fine-tune specialized question encoders, leading to memory-efficient transfer to new questions. We find that data augmentation, while able to close gaps on a single domain, is unable to consistently improve performance on unseen domains. We also find that building a robust passage encoder is crucial in order to successfully adapt to new domains. We view this study as one important step towards building universal dense retrieval models.

2 Background and Related Work

Sparse retrieval Before the emergence of dense retrievers, traditional sparse retrievers such as TF-IDF or BM25 were the de facto method in open-domain question-answering systems (Chen et al., 2017; Yang et al., 2019). These sparse models measure similarity using weighted term-matching between questions and passages and do not train on a particular data distribution. It is well-known that sparse models are great at lexical matching, but fail to capture synonyms and paraphrases.

Dense retrieval On the contrary, dense models (Lee et al., 2019; Karpukhin et al., 2020; Guu et al., 2020) measure similarity using learned representations from supervised QA datasets, leveraging pre-trained language models like BERT. In this paper, we use the popular dense passage retriever (DPR) model (Karpukhin et al., 2020) as our main evaluation,³ and we also report the evaluation of REALM (Guu et al., 2020) in Appendix A. DPR models the retrieval problem using two encoders, namely the question and the passage encoders, initialized using BERT. DPR uses a contrastive objective during training, with in-batch negatives and hard negatives mined from BM25. During inference, a pre-defined large set of passages (e.g., 21-million passages in English Wikipedia) are encoded

³The detailed experimental settings are in Appendix B.

and pre-indexed—for any test question, the top passages with the highest similarity scores are returned. Recently, other advances have been made in improving dense retrieval, including incorporating better hard negatives (Xiong et al., 2021; Qu et al., 2021), or fine-grained phrase retrieval (Lee et al., 2021). We leave them for future investigation.

Generalization problem Despite the impressive in-domain performance of dense retrievers, their capability of generalizing to unseen questions still remains relatively under-explored. Recently, Lewis et al. (2021a) discover that there is a large overlap between training and testing sets on popular QA benchmarks, concluding that current models tend to memorize training questions and perform significantly worse on non-overlapping questions. AmbER (Chen et al., 2021) test sets are designed to study the entity disambiguation capacities of passage retrievers and entity linkers. They find models perform much worse on rare entities compared to common entities. Similar to this work, our results show dense retrieval models generalize poorly, especially on rare entities. We further conduct a series of analyses to dissect the problem and investigate potential approaches for learning robust dense retrieval models. Finally, another concurrent work (Thakur et al., 2021) introduces the BEIR benchmark for zero-shot evaluation of retrieval models and shows that dense retrieval models underperform BM25 on most of their datasets.

3 EntityQuestions

In this section, we build a new benchmark EntityQuestions, a set of simple, entity-centric questions and compare dense and sparse retrievers.

Dataset collection We select 24 common relations from Wikidata (Vrandečić and Krötzsch, 2014) and convert fact (*subject, relation, object*) triples into natural language questions using manually defined templates (Appendix A). To ensure the converted natural language questions are answerable from Wikipedia, we sample triples from the T-REx dataset (Elsahar et al., 2018), where triples are aligned with a sentence as evidence in Wikipedia. We select relations following the criteria: (1) there are enough triples (>2k) in the T-REx; (2) it is easy enough to formulate clear questions for the relation; (3) we do not select relations with only a few answer candidates (e.g., *gender*), which may cause too many false negatives when we evaluate

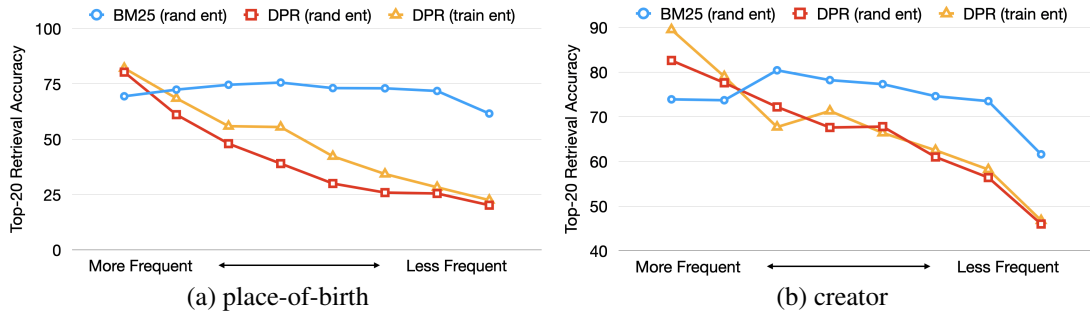


Figure 1: Top-20 retrieval accuracy on (a) *place-of-birth* questions, (b) *creator* questions, grouped by the entity’s frequency in Wikipedia. We group entities into 8 buckets according to their frequency in Wikipedia. *rand ent*: randomly selected entities from Wikidata; *train ent*: entities in the NQ training set.

the retriever; (4) we include both person-related relations (e.g., *place-of-birth*) and non-person relations (e.g., *headquarter*). For each relation, we randomly sample up to 1,000 facts to form the evaluation set. We report the averaged accuracy over all relations of EntityQuestions.

Results We evaluate DPR and BM25 on the EntityQuestions dataset and report results in Table 1 (see full results and examples in Appendix A). DPR trained on NQ significantly underperforms BM25 on almost all sets of questions. For example, on the question “Where was [E] born?”, BM25 outperforms DPR by 49.8% absolute using top-20 retrieval accuracy.⁴ Although training DPR on multiple datasets can improve the performance (i.e., from 49.7% to 56.7% on average), it still clearly pales in comparison to BM25. We note the gaps are especially large on questions about person entities.

In order to test the generality of our findings, we also evaluate the retrieval performance of REALM (Guu et al., 2020) on EntityQuestions. Compared to DPR, REALM adopts a pre-training task called salient span masking (SSM), along with an inverse cloze task from Lee et al. (2019). We include the evaluation results in Appendix A.⁵ We find that REALM still scores much lower than BM25 over all relations (19.6% on average). This suggests that incorporating pre-training tasks such as SSM still does not solve the generalization problem on these simple entity-centric questions.

⁴For our entire analysis, we consider top-20 retrieval accuracy for brevity. However, trends still hold for top-1, top-5, and top-100 retrieval accuracy.

⁵We cannot directly compare the retrieval accuracy of REALM to DPR, as the REALM index uses 288 BPE token blocks while DPR uses 100 word passages.

4 Dissecting the Problem: Entities vs. Question Patterns

In this section, we investigate why dense retrievers do not perform well on these questions. Specifically, we want to understand whether the poor generalization should be attributed to (a) novel entities, or (b) unseen question patterns. To do this, we study DPR trained on the NQ dataset and evaluate on three representative question templates: *place-of-birth*, *headquarter*, and *creator*.⁶

4.1 Dense retrievers exhibit popularity bias

We first determine how the entity [E] in the question affects DPR’s ability to retrieve relevant passages. To do this, we consider all triples in Wikidata that are associated with a particular relation, and order them based on frequency of the subject entity in Wikipedia. In our analysis, we use the Wikipedia hyperlink count as a proxy for an entity’s frequency. Next, we group the triples into 8 buckets such that each bucket has approximately the same cumulative frequency.

Using these buckets, we consider two new evaluation sets for each relation. The first (denoted “rand ent”) randomly samples at most 1,000 triples from each bucket. The second (denoted “train ent”) selects all triples within each bucket that have subject entities observed in questions within the NQ training set, as identified by ELQ (Li et al., 2020).

We evaluate DPR and BM25 on these evaluation sets and plot the top-20 accuracy in Figure 1. DPR performs well on the most common entities but quickly degrades on rarer entities, while BM25 is less sensitive to entity frequency. It is also notable that DPR performs generally better on entities

⁶The question templates for these relations are: *place-of-birth*: “Where was [E] born?”; *headquarter*: “Where is the headquarters of [E]?”; *creator*: “Who was [E] created by?”.

	p-of-birth	headquarter	creator
DPR-NQ	25.4	70.0	54.1
FT	73.9	84.0	80.0
FT w/ similar	74.7	79.9	76.2
FT OnlyP	72.8	84.2	78.0
FT OnlyQ	45.4	72.8	73.4
BM25	75.2	85.0	71.7

Table 2: Top-20 retrieval accuracy on EntityQuestions test sets after fine-tuning. FT: fine-tuning on each individual question pattern. w/ similar: fine-tuning on a similar, semantically equivalent question pattern. OnlyP and OnlyQ: fixing the weights of the question encoder and only updating the passage encoder, or vice versa.

seen during NQ training than on randomly selected entities. This suggests that DPR representations are much better at representing the most common entities as well as entities observed during training.

4.2 Observing questions helps generalization

We next investigate whether DPR generalizes to unseen entities when trained on the question pattern. For each relation considered, we build a training set with at most 8,000 triples. We ensure no *tokens* from training triples overlap with tokens from triples in the corresponding test set. In addition to using the question template used during evaluation to generate training questions, we also build a training set based on a syntactically different but semantically equal question template.⁷ We fine-tune DPR models on the training set for each relation and test on the evaluation set of EntityQuestions for the particular relation and report results in Table 2.

Clearly, observing the question pattern during training allows DPR to generalize well on unseen entities. On all three relations, DPR can match or even outperform BM25 in terms of retrieval accuracy. Training on the equivalent question pattern achieves comparable performance to the exact pattern, showing dense models do not rely on specific phrasing of the question. We also attempt fine-tuning the question encoder and passage encoder separately. As shown in Table 2, surprisingly, there is a significant discrepancy between only training the passage encoder (OnlyP) and only training the question encoder (OnlyQ): for example, on place-of-birth, DPR achieves 72.8% accuracy with the fine-tuned passage encoder, while it

⁷place-of-birth: “What is the birthplace of [E]?”; headquarter: “Where is [E] headquartered?”; creator: “Who is the creator of [E]?”.

	NQ	Rel	EntityQ.
DPR-NQ	80.1	25.4	49.7
+ FT <i>p-of-birth</i>	62.8	74.3	56.2
+ FT $NQ \cup p\text{-of-birth}$	70.8	52.0	47.4
DPR-NQ	80.1	70.0	49.7
+ FT <i>headquarter</i>	71.6	80.3	53.3
+ FT $NQ \cup headquarter$	75.1	81.3	49.5
DPR-NQ	80.1	54.1	49.7
+ FT <i>creator</i>	70.8	80.8	52.3
+ FT $NQ \cup creator$	72.6	72.3	44.1
BM25	64.5	-	71.2

Table 3: Top-20 retrieval accuracy on NQ and EntityQuestions. FT: fine-tuning. Rel: the performance on the relation that is used during fine-tuning.

achieves 45.4% if only the question encoder is fine-tuned. This suggests that passage representations might be the culprit for model generalization.

To understand what passage representations have learned from fine-tuning, we visualize the DPR passage space before and after fine-tuning using t-SNE (Van der Maaten and Hinton, 2008). We plot the representations of positive passages sampled from NQ and *place-of-birth* in Figure 2. Before fine-tuning, positive passages for *place-of-birth* questions are clustered together. Discriminating passages in this clustered space is more difficult using an inner product, which explains why only fine-tuning the question encoder yields minimal gains. After fine-tuning, the passages are distributed more sparsely, making differentiation much easier.

5 Towards Robust Dense Retrieval

Equipped with a clear understanding of the issues, we explore some simple techniques aimed at fixing the generalization problem.

Data augmentation We first explore whether fine-tuning on questions from a *single* EntityQuestions relation can help generalize on the full set of EntityQuestions as well as other QA datasets such as NQ. We construct a training set of questions for a single relation and consider two training regimes: one where we fine-tune on relation questions alone; and a second where we fine-tune on both relation questions and NQ in a multi-task fashion. We perform this analysis for three relations and report top-20 retrieval accuracy in Table 3.

We find that fine-tuning only on a single relation improves EntityQuestions meaningfully, but degrades performance on NQ and still largely falls

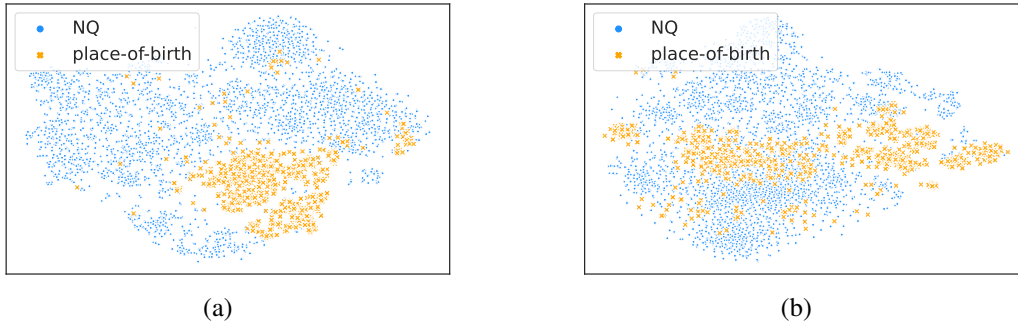


Figure 2: Visualization of positive passage embeddings returned by DPR before and after fine-tuning on the *place-of-birth* questions. (a): Positive passage embeddings returned by DPR trained on NQ; (b) Positive passage embeddings returned by DPR after fine-tuning.

behind BM25 on average. When fine-tuning on both relation questions and NQ together, most of the performance on NQ is retained, but the gains on EntityQuestions are much more muted. Clearly, fine-tuning on one type of entity-centric question does not necessarily fix the generalization problem for other relations. This trade-off between accuracy on the original distribution and improvement on the new questions presents an interesting tension for universal dense encoders to grapple with.

Specialized question encoders While it is challenging to have one retrieval model for all unseen question distributions, we consider an alternative approach of having a single passage index and adapting specialized question encoders. Since the passage index is fixed across different question patterns and cannot be adapted using fine-tuning, having a robust passage encoder is crucial.

We compare two DPR passage encoders: one based on NQ and the other on the PAQ dataset (Lewis et al., 2021b).⁸ We expect a question encoder trained on PAQ is more robust because (a) 10M passages are sampled in PAQ, which is arguably more varied than NQ, and (b) all the plausible answer spans are identified using automatic tools. We fine-tune a question encoder for each relation in EntityQuestions, keeping the passage encoder fixed. As shown in Table 4,⁹ fine-tuning the encoder trained on PAQ improves performance over fine-tuning the encoder trained on NQ. This suggests the DPR-PAQ encoder is more robust and adaptable, nearly closing the gap with BM25 using a single passage index. We believe constructing a robust passage index is an encouraging avenue for

⁸PAQ dataset sampling scheme is described in Appendix B.

⁹Per-relation accuracy can be found in Appendix C.

	EntityQ.
DPR-NQ [†]	45.1
+ Per-relation FT (OnlyQ)	61.6
+ EntityQuestions FT (OnlyQ)	53.0
DPR-PAQ [†]	59.3
+ Per-relation FT (OnlyQ)	68.4
+ EntityQuestions FT (OnlyQ)	65.4
BM25	71.2

Table 4: Averaged top-20 retrieval accuracy on EntityQuestions. We fix the passage encoder and fine-tune the question encoder. *Per-relation FT*: fine-tuning an individual question encoder for each relation, *EntityQuestions FT*: fine-tuning on all questions in EntityQuestions. [†]: we re-train the models from scratch.

future work towards a more general retriever.

6 Conclusion

In this study, we show that DPR significantly underperforms BM25 on EntityQuestions, a dataset of simple questions based on facts mined from Wikidata. We derive key insights about why DPR performs so poorly on this dataset. We learn that DPR remembers robust representations for common entities, but has trouble differentiating rarer entities without explicitly observing the question pattern during training.

We suggest future work in incorporating explicit entity memory into dense retrievers to help differentiate rare entities. Numerous recent works (Wu et al. 2020; Li et al. 2020; Cao et al. 2021) demonstrate retrievers can easily learn dense representations for a large number of Wikipedia entities. DPR could also leverage entity-aware embedding models like EaE (Férvy et al., 2020) or LUKE (Yamada et al., 2020) to better recall long-tail entities.

Acknowledgements

We thank the members of the Princeton NLP group for helpful discussion and valuable feedback. This research is supported by gift awards from Apple and Amazon.

Ethical Considerations

Our proposed dataset, EntityQuestions, is constructed by sampling (subject, relation, object) triples from Wikidata, which is dedicated to the public domain under the Creative Commons CC0 License. In general, machine learning has the ability to amplify biases presented implicitly and explicitly in the training data. Models that we reference in our study are based on BERT, which has been shown to learn and exacerbate stereotypes during training (e.g., Kurita et al. 2019, Tan and Celis 2019, Nadeem et al. 2021). We further train these models on Wikidata triples, which again has the potential to amplify harmful and toxic biases.

In the space of open-domain question answering, deployed systems leveraging biased pre-trained models like BERT will likely be less accurate or biased when asked questions related to stereotyped and marginalized groups. We acknowledge this fact and caution those who build on our work to consider and study this implication before deploying systems in the real world.

References

- Petr Baudiš and Jan Šedivý. 2015. [Modeling of the question answering task in the yodaqa system](#). In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 222–228. Springer.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on Freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, Seattle, Washington, USA. Association for Computational Linguistics.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021. [Autoregressive entity retrieval](#). In *International Conference on Learning Representations (ICLR)*.
- Anthony Chen, Pallavi Gudipati, Shayne Longpre, Xiao Ling, and Sameer Singh. 2021. [Evaluating entity disambiguation and the role of popularity in retrieval-based NLP](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4472–4485, Online. Association for Computational Linguistics.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading Wikipedia to answer open-domain questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hady Elsahar, Pavlos Vougiouklis, Arslan Remaci, Christophe Gravier, Jonathon Hare, Frederique Laforest, and Elena Simperl. 2018. [T-REx: A large scale alignment of natural language with knowledge base triples](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Thibault Févry, Livio Baldini Soares, Nicholas FitzGerald, Eunsol Choi, and Tom Kwiatkowski. 2020. [Entities as experts: Sparse memory access with entity supervision](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4937–4951, Online. Association for Computational Linguistics.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Papatat, and Ming-Wei Chang. 2020. [REALM: Retrieval-augmented language model pre-training](#). In *International Conference on Machine Learning (ICML)*.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.

- Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. 2019. [Measuring bias in contextualized word representations](#). In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 166–172, Florence, Italy. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Jinhyuk Lee, Mujeen Sung, Jaewoo Kang, and Danqi Chen. 2021. [Learning dense representations of phrases at scale](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6634–6647, Online. Association for Computational Linguistics.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy. Association for Computational Linguistics.
- Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. 2021a. [Question and answer test-train overlap in open-domain question answering datasets](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1000–1008, Online. Association for Computational Linguistics.
- Patrick Lewis, Yuxiang Wu, Linqing Liu, Pasquale Minervini, Heinrich Küttler, Aleksandra Piktus, Pontus Stenetorp, and Sebastian Riedel. 2021b. [PAQ: 65 million probably-asked questions and what you can do with them](#). *ArXiv preprint*, abs/2102.07033.
- Belinda Z. Li, Sewon Min, Srinivasan Iyer, Yashar Mehdad, and Wen-tau Yih. 2020. [Efficient one-pass end-to-end entity linking for questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6433–6441, Online. Association for Computational Linguistics.
- Jimmy Lin, Xueguang Ma, Sheng-Chieh Lin, Jheng-Hong Yang, Ronak Pradeep, and Rodrigo Nogueira. 2021. [Pyserini: An easy-to-use python toolkit to support replicable ir research with sparse and dense representations](#). *ArXiv preprint*, abs/2102.10073.
- Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. [StereoSet: Measuring stereotypical bias in pre-trained language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5356–5371, Online. Association for Computational Linguistics.
- Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2021. [RocketQA: An optimized training approach to dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5835–5847, Online. Association for Computational Linguistics.
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: BM25 and beyond](#). *Foundations and Trends in Information Retrieval*, 3(4):333–389.
- Yi Chern Tan and L. Elisa Celis. 2019. [Assessing social and intersectional biases in contextualized word representations](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13209–13220.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. [BEIR: A heterogenous benchmark for zero-shot evaluation of information retrieval models](#). *ArXiv preprint*, abs/2104.08663.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. [Visualizing data using t-sne](#). *Journal of machine learning research*, 9(11).
- Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: a free collaborative knowledgebase](#). *Communications of the ACM*, 57(10):78–85.
- Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2020. [Zero-shot entity linking with dense entity retrieval](#). In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul Bennett, Junaid Ahmed, and Arnold Overwijk. 2021. [Approximate nearest neighbor negative contrastive learning for dense text retrieval](#). In *International Conference on Learning Representations (ICLR)*.
- Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. [LUKE: Deep contextualized entity representations with entity-aware self-attention](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6442–6454, Online. Association for Computational Linguistics.
- Wei Yang, Yuqing Xie, Aileen Lin, Xingyu Li, Luchen Tan, Kun Xiong, Ming Li, and Jimmy Lin. 2019. [End-to-end open-domain question answering with](#)

BERTserini. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 72–77, Minneapolis, Minnesota. Association for Computational Linguistics.

Relation	DPR	DPR	BM25
	(NQ)	(multi)	
P36 What is the capital of [E]?	77.3	78.9	90.0
P407 Which language was [E] written in?	77.1	82.5	86.1
P26 Who is [E] married to?	35.6	48.1	85.9
P159 Where is the headquarter of [E]?	70.0	72.0	85.0
P276 Where is [E] located?	74.9	77.3	84.3
P40 Who is [E]’s child?	19.2	33.8	82.9
P176 Which company is [E] produced by?	61.7	73.7	80.6
P20 Where did [E] die?	34.4	45.1	80.3
P112 Who founded [E]?	77.1	75.7	78.2
P127 Who owns [E]?	60.7	63.8	77.5
P19 Where was [E] born?	25.4	41.8	75.2
P740 Where was [E] founded?	59.9	61.6	74.3
P413 What is [E] famous for?	19.0	33.9	74.2
P800 What position does [E] play?	75.7	71.5	74.2
P69 Where was [E] educated?	26.4	41.8	73.0
P50 Who is the author of [E]?	75.7	77.8	72.2
P170 Who was [E] created by?	54.1	57.7	71.7
P106 What kind of work does [E] do?	25.9	52.9	70.7
P131 Where is [E] located?	45.7	44.2	62.6
P17 Which country is [E] located in?	64.2	67.7	61.0
P175 Who performed [E]?	47.6	51.5	54.1
P136 What type of music does [E] play?	37.4	36.8	48.5
P264 What music label is [E] represented by?	25.3	43.2	45.0
P495 Which country was [E] created in?	21.6	28.0	21.8
<i>Average</i>	49.7	56.7	71.2

Table 5: Top-20 retrieval accuracy (percentage of retrieved passages that contain the answer) for dense and sparse retrieval models on different sets of questions of EntityQuestions. We test two DPR models: (1) trained on NQ only; (2) trained on 4 datasets (NQ, TQA, WebQ, TREC) combined.

A Full Results on EntityQuestions

DPR vs. BM25 The evaluation results are shown in Table 5. BM25 significantly outperforms DPR models trained on either a single dataset NQ or a combination of common QA datasets.

REALM vs. BM25 We also evaluate the retrieval performance of REALM (Guu et al., 2020) on EntityQuestions. Specifically, we use REALM to retrieve 20 passages and check if the gold answer is a sub-string of the retrieved passages. We also evaluate BM25 on the same 288-token blocks that are used in REALM model. As shown in Table 6, the results show that REALM still significantly underperforms BM25 on EntityQuestions, even with the extra pre-training tasks.

Examples of DPR retrieved passages Table 7 shows examples of DPR retrieved results on three representative questions. DPR makes clear mistakes like confusing entities with similar names or missing the presence of an entity, causing it to retrieve irrelevant passages on these simple, entity-centric questions.

Relation	REALM		BM25
	REALM	BM25	
P36 What is the capital of [E]?	91.7	91.9	
P407 Which language was [E] written in?	81.9	92.0	
P26 Who is [E] married to?	47.1	90.0	
P159 Where is the headquarter of [E]?	70.4	90.7	
P276 Where is [E] located?	77.1	89.5	
P40 Who is [E]’s child?	39.7	87.1	
P176 Which company is [E] produced by?	69.2	83.2	
P20 Where did [E] die?	61.9	89.2	
P112 Who founded [E]?	77.3	85.5	
P127 Who owns [E]?	73.6	84.4	
P19 Where was [E] born?	52.9	90.3	
P740 Where was [E] founded?	50.9	77.5	
P413 What position does [E] play?	53.8	90.4	
P800 What is [E] famous for?	45.3	81.9	
P69 Where was [E] educated?	38.6	84.1	
P50 Who is the author of [E]?	77.2	76.2	
P170 Who was [E] created by?	56.8	78.5	
P106 What kind of work does [E] do?	53.6	83.4	
P131 Where is [E] located?	63.9	86.8	
P17 Which country is [E] located in?	70.6	76.0	
P175 Who performed [E]?	53.1	65.8	
P136 What type of music does [E] play?	42.6	53.4	
P264 What music label is [E] represented by?	53.2	55.3	
P495 Which country was [E] created in?	34.8	24.8	
<i>Average</i>	59.9	79.5	

Table 6: Top-20 retrieval accuracy of REALM and BM25 on different sets of questions of EntityQuestions.

B Experimental Details

Experimental settings of DPR In our experiments, we use either pre-trained DPR models released by the authors, or the DPR models re-trained by ourself (Table 4). All our experiments are carried out on $4 \times 11\text{Gb}$ Nvidia RTX 2080Ti GPUs. For all our fine-tuning experiments, we fine-tune for 10 epochs, with a learning rate 2×10^{-5} and a batch size of 24. When we retrain DPR from scratch, we train for 20 epochs with a batch size of 24 (the original DPR models were trained on $8 \times 32\text{Gb}$ GPUs with a batch size of 128 and we have to reduce the batch size due to the limited computational resources) and a learning rate of 2×10^{-5} .

Experimental settings of BM25 In our experiments, we use the Pyserini (Lin et al., 2021) implementation of unigram BM25 with default parameters. We build an index using the same Wikipedia passage splits provided in the official DPR release.

PAQ dataset sampling Lewis et al. (2021b) introduce Probably Asked Questions (PAQ), a large question repository constructed using a question generation model on Wikipedia passages. We group all of the questions asked about a particular passage and filter out any passages that have

Question	Retrieved results from DPR-NQ
Q: Where was Arve Furset born?	<p>Gold: [Title: Arve Furset] Arve Furset Arve Eilif Furset (born 5 December 1964 in Askvoll, Western Norway) is a Norwegian composer, jazz musician (piano, keyboards) and music producer</p> <p>Top-1: [Title: Gard Agdi] Gard Agdi ("Old Norse" Garðr Agði) appears in the legendary genealogies of "Hversu Noregr byggdist" as one of the three sons of . . .</p> <p>Top-2: [Title: Yrsa] kidnapped the queen for a while during which time he made her pregnant. Having returned to her kingdom, the queen bore a child, a girl which she named Yrsa after her dog. Yrsa was sent. . .</p> <p>Top-3: [Title: Arvid Noe] Sailor" and the anagram "Arvid Noe" to conceal his identity; his true name, Arne Vidar Røed, became known after his death. Røed began his career as a sailor in 1961, when he. . .</p>
Q: Where was Mason Proper founded?	<p>Gold: [Title: Mason Proper] Mason Proper is an American rock band formed in Alpena, Michigan, in 2004. The band consists of Jonathan Visger (singer), Matt. . .</p> <p>Top-1: [Title: George Mason] with the world. Thus, colonial Virginia initially developed few towns, since estates were largely self-sufficient, and could get what they needed without the need to purchase locally. Even the. . .</p> <p>Top-2: [Title: History of Freemasonry] so on through an elaborate path to Athelstan. This myth formed the basis for subsequent manuscript constitutions, all tracing masonry back to biblical times, and fixing its institutional. . .</p> <p>Top-3: [Title: Mason (surname)] Masullo, Masutti, Masutto, and several others. Mason (surname) Mason is an Italian, French Irish or English surname that refers to someone who did stonemasonry work, or it derives from. . .</p>
Q: Who was Veep created by?	<p>Gold: [Title: Veep] Veep is an American political satire comedy television series... created by Scottish writer Armando Iannucci as an...</p> <p>Top-1: [Title: Beep the Meep] (DWM #250). The Meeps (although not Beep himself) appeared in the spin-off short story "Wish Upon A Star Beast" by Steve Lyons, published in the charity anthology "Perfect Timing". Beep. . .</p> <p>Top-2: [Title: Beep the Meep] Beep the Meep is a fictional alien who appeared in the "Doctor Who Weekly" comic strip based on the long-running British science fiction television series "Doctor Who". The cute and cuddly. . .</p> <p>Top-3: [Title: Mister Fantastic] Mister Fantastic (Reed Richards) is a fictional superhero appearing in American comic books published by Marvel Comics. The character is a founding member of the Fantastic Four. Richards. . .</p>

Table 7: Examples of DPR retrieved results on simple entity-centric questions.

less than 3 generated questions. We then sample 100K such passages and sample one question asked about each. We split this dataset into 70K/15K/15K for train/dev/test splits, although we do not evaluate on this dataset. Following [Karpukhin et al. \(2020\)](#), we use BM25 to mine hard negative examples.

C Per-relation Accuracy with Different Passage Encoders

We fine-tune DPR with the passage encoder fixed on either NQ or PAQ. Table 8 compares the per-relation accuracy of DPR with fixed passage encoder fine-tuned on NQ and PAQ. As is shown, the passage encoder trained on PAQ is much more robust than the passage encoder trained on NQ. For many non-person relations, using a PAQ-based passage encoder can outperform BM25.

Relation	DPR (NQ)	Per-rel FT (OnlyQ)	EQ FT (OnlyQ)	DPR (PAQ)	Per-rel FT (OnlyQ)	EQ FT (OnlyQ)	BM25
P106 What kind of work does [E] do?	19.9	59.6	19.9	47.7	71.6	47.7	70.7
P112 Who founded [E]?	74.7	72.2	73.3	75.1	74.9	76.3	78.2
P127 Who owns [E]?	46.5	70.3	46.5	63.4	73.6	63.4	77.5
P131 Where is [E] located?	44.1	50.6	49.9	42.1	49.5	50.8	62.6
P136 What type of music does [E] play?	34.7	57.3	54.8	44.7	57.6	56.0	48.5
P159 Where is the headquarter of [E]?	69.0	77.7	78.3	72.2	75.5	78.4	85.0
P17 Which country is [E] located in?	56.6	63.9	64.2	58	65.2	64.3	61.0
P170 Who was [E] created by?	33.4	64.8	33.4	66.1	75.6	66.1	71.7
P175 Who performed [E]?	41.6	56.2	41.6	51.4	57.8	51.4	54.1
P176 Which company is [E] produced by?	43.0	81.0	43.0	73.9	82.2	73.9	80.6
P19 Where was [E] born?	26.0	48.1	53.8	54.6	63.9	64.4	75.2
P20 Where did [E] die?	32.8	61.1	65.4	63.1	71.8	70.7	80.3
P26 Who is [E] married to?	25.1	32.7	38.5	60.8	69.4	69.4	85.9
P264 What music label is [E] represented by?	27.6	47.9	27.6	47.1	53.8	52.2	45.0
P276 Where is [E] located?	71.4	80.8	76.8	73.8	80.6	80.2	84.3
P36 What is the capital of [E]?	74.9	82.2	74.9	76.6	82.4	85.7	90.0
P40 Who is [E]’s child?	16.5	46.0	16.5	49.7	63.0	63.9	82.9
P407 Which language was [E] written in?	72.9	81.6	85.1	73.7	84.2	86.5	86.1
P413 What position does [E] play?	75.7	85.3	75.7	69.3	85.7	69.3	74.2
P495 Which country was [E] created in?	19.4	24.2	35.1	20.2	26.0	30.0	21.8
P50 Who is the author of [E]?	75.7	79.4	79.6	74.8	79.8	79.2	72.2
P69 Where was [E] educated?	19.9	55.9	55.8	48.9	68.1	69.5	73.0
P740 Where was [E] founded?	57.0	77.8	57.0	67.6	79.9	67.6	74.3
P800 What is [E] famous for?	24.4	22.6	25.8	49.3	49.8	52.0	74.2
<i>Average</i>	45.1	61.6	53.0	59.3	68.4	65.4	71.2

Table 8: Top-20 retrieval accuracy on NQ and EntityQuestions (EQ). *Per-rel FT*: we fine-tune an individual question encoder for each relation. *EQ FT*: we fine-tune a single question encoder on all relations in EntityQuestions.