
Autoregressive Search Engines: Generating Substrings as Document Identifiers

Michele Bevilacqua^{1,2} Giuseppe Ottaviano² Patrick Lewis²
Wen-tau Yih² Sebastian Riedel^{2,3} Fabio Petroni²
¹Sapienza University of Rome ²Meta AI ³University College London

Abstract

Knowledge-intensive language tasks require NLP systems to both provide the correct answer and retrieve supporting evidence for it in a given corpus. Autoregressive language models are emerging as the de-facto standard for generating answers, with newer and more powerful systems emerging at an astonishing pace. In this paper we argue that all this (and future) progress can be directly applied to the retrieval problem with minimal intervention to the models' architecture. Previous work has explored ways to partition the search space into hierarchical structures and retrieve documents by autoregressively generating their unique identifier. In this work we propose an alternative that doesn't force any structure in the search space: using all ngrams in a passage as its possible identifiers. This setup allows us to use an autoregressive model to generate and score distinctive ngrams, that are then mapped to full passages through an efficient data structure. Empirically, we show this not only outperforms prior autoregressive approaches but also leads to an average improvement of at least 10 points over more established retrieval solutions for passage-level retrieval on the KILT benchmark, establishing new state-of-the-art downstream performance on some datasets, while using a considerably lighter memory footprint than competing systems. Code and pre-trained models are available at <https://github.com/facebookresearch/SEAL>.

1 Introduction

Surfacing knowledge from large corpora is a crucial step when dealing with knowledge intensive language tasks [Levy et al., 2017, Dinan et al., 2019, Elsayar et al., 2018, Petroni et al., 2021], such as open-domain question answering [Voorhees et al., 1999, Joshi et al., 2017, Yang et al., 2018, Kwiatkowski et al., 2019] and fact checking [Thorne et al., 2018]. A popular paradigm to approach such tasks is to combine a search engine with a machine reader component. The former retrieves relevant context, usually in the form of short passages, which the latter then examines to produce answers [Chen et al., 2017, Lewis et al., 2020, Izacard and Grave, 2021].

In recent years we have witnessed a surge of research and development in autoregressive language models [Radford et al., 2019, Lewis et al., 2019, Raffel et al., 2019, Brown et al., 2020, Rae et al., 2021, Artetxe et al., 2021, Smith et al., 2022], with ever increasing size and natural language understanding (NLU) capabilities. Such models are currently the de-facto implementation of the machine reader component in retrieval-reader architectures, and have contributed to rapid progress on a wide range of benchmarks [Joshi et al., 2017, Kwiatkowski et al., 2019, Petroni et al., 2021]. However, these tremendous advances in aggressive modelling has yet to bring similar transformational changes in how retrieval is approached.

Transferring the NLU capabilities of modern autoregressive models to retrieval is non-trivial. Some works have demonstrated that knowledge stored in the parameters of these models can be retrieved to some extent by directly generating evidence given a query [Petroni et al., 2019, 2020, Roberts




Figure 1: High-level SEAL architecture, composed of an autoregressive LM paired with an FM-Index, for which we show the first (F) and last (L) columns of the underlying matrix (more details in Sec 3.1). The FM-index constrains the autoregressive generation (*e.g.*, after *carbon* the model is constrained to generate either *tax*, *dioxide* or *atom* in the example) and provides the documents matching (*i.e.*, containing) the generated ngram (at each decoding step).

et al., 2020]. However, such approaches have been shown to be unreliable because of their tendency to hallucinate non-factual content [Massarelli et al., 2019, Metzler et al., 2021, Ji et al., 2022]. To alleviate this issue, previous work proposes to only use generation for query expansion in traditional search engines [Mao et al., 2021], but these solutions don’t exploit the full potential of autoregressive architecture, such as word order sensitivity and conditional probability modeling, and still lag behind vector-based approaches [Karpukhin et al., 2020].

Recently, another line of work has investigated using autoregressive language models to generate identifier strings for documents as an intermediate target for retrieval, such as Wikipedia page titles [De Cao et al., 2021b], or root-to-leaf paths in a hierarchical cluster tree [Tay et al., 2022]. Employing identifiers, rather than generating evidence directly, induces structure in the search space, (*i.e.*, index documents by their title or their cluster tree) which can be easier to memorize, learn, and retrieve from, than full unstructured passages. Moreover, it is relatively easy to constrain decoding with a prefix tree “index” so that only valid identifiers are generated. As a downside, if appropriate metadata (*e.g.*, titles) are not available, one needs to create the identifiers, hence the structure (*e.g.*, with hierarchical clustering), which has not been thoroughly evaluated on a large-scale benchmark.

In this work, we propose a solution that does not force any structure in the search space, but rather uses all the ngrams occurring in a document as its identifiers. Concretely, we introduce Search Engines with Autoregressive LMs (SEAL), a retrieval solution that combines an autoregressive model, *i.e.*, BART [Lewis et al., 2019], with a compressed full-text substring index, *i.e.*, the FM-Index [Ferragina and Manzini, 2000] — see Figure 1 for an high-level overview. This configuration comes with a twofold benefit: i) we can constrain BART’s generations with the FM-Index, hence preventing the generation of invalid identifiers (*i.e.*, ngrams not occurring in any document); ii) the FM-Index provides information on all documents in the corpus containing a specific ngram (for every decoding step), thus allowing to retrieve them. This setup allows SEAL to generate *any span* from *any position* in the corpus, without needing to explicitly encode all substrings in a document. Moreover, we design a novel scoring function to intersect the results of multiple ngrams combining LM probabilities with FM-index frequencies (*i.e.*, number of occurrences of the ngram in the whole corpus).

Our experimental evaluation shows that SEAL matches or outperforms recent retrieval solutions (including autoregressive ones) on Natural Questions [Kwiatkowski et al., 2019], while requiring substantially less memory (~2 to 7 times smaller in footprint). Moreover, SEAL’s intersection formulation improves the state-of-the-art on passage-level retrieval by more than 10 points on the KILT benchmark [Petroni et al., 2021], contributing in establishing new state-of-the-art downstream results on multiple datasets when paired with existing reader technologies.

2 Related Work

One way to approach retrieval with autoregressive models makes use of unique identifiers, *i.e.*, string pointers to documents that are in some way easier to generate than the full document itself [De Cao

et al., 2021b, Tay et al., 2022]. In our work the identifiers are generated ngrams, which do not necessarily occur in just one document. Note that the idea of using phrases to retrieve passages has shown promise already in the context of dense retrieval [Lee et al., 2021a,b].

Our approach is quite conceptually distinct from document [Nogueira et al., 2019, Nogueira and Lin, 2021] and query expansion approaches [Mao et al., 2021], that use autoregressive models to boost search, but still rely on black-box systems like BM25 [Robertson and Zaragoza, 2009] for the retrieval: in our work the boundary between generation and retrieval is blurred, since we generate *grounded* passage spans. The same reasoning applies to models that generate queries for web search engines [Komeili et al., 2021, Shuster et al., 2022, Nakano et al., 2021, Lazaridou et al., 2022].

Virtually all modern approaches to string-matching-based retrieval use bag-of-words representations. Many recent works propose learned contextualized weighting for both queries and documents terms [Dai and Callan, 2019, Gao et al., 2021, Lin and Ma, 2021, Mallia et al., 2021, Dai and Callan, 2020, Bai et al., 2020, Zhao et al., 2021, Formal et al., 2021b,a]. Many of these methods can also weigh terms that are not present in the query, addressing so-called vocabulary mismatch. In contrast, SEAL generates (and scores) ngrams of arbitrary size. These approaches are partly orthogonal to SEAL, as many of the proposed techniques could be used to rescore higher-order ngrams.

Finally, a connected strand of research is that of query likelihood models, which use autoregressive models to (re)rank passages according to the probability $P(q|p)$ of a query q given the passage p [Nogueira dos Santos et al., 2020, Zhuang and Zuccon, 2021, Lesota et al., 2021, Sachan et al., 2022]. In our case, the autoregressive architecture models the likelihood of an ngram given the query, *i.e.*, $P(n|q)$.

3 Background

In retrieval, the automatic system is required to return an ordered list of documents d_1, d_2, \dots, d_n from a retrieval corpus \mathcal{R} , given a query q . Both queries and documents are texts, *i.e.*, lists of tokens $\langle t_1, t_2, \dots, t_N \rangle$, where each token t is drawn from a vocabulary V . A span of tokens in a text is called an ngram; ngrams of size 1 are known as unigrams. We denote with $F(n, \mathcal{R})$ the frequency of an ngram n in \mathcal{R} , *i.e.*, the total number of times it appears in the whole retrieval corpus.

3.1 The FM-Index

Our method requires a data structure that can support the efficient identification of occurring substrings to guarantee that all decoded sequences are located somewhere in the retrieval corpus. Moreover, to perform retrieval, we require the ability to identify which documents the generated ngrams appear in. Neither inverted indices (which have no efficient way to search for phrases of arbitrary length), nor prefix trees (which would force us to explicitly encode all k suffixes in a document), are viable options. The core data structure that satisfies our requirements is the FM-index [Ferragina and Manzini, 2000], *i.e.*, a compressed suffix array that, as a self-index, requires no additional storage for the original text. FM-index space requirements are linear in the size of the corpus, and, with small vocabularies such as those used by modern subword-based language models, is thus usually *significantly smaller* than the uncompressed corpus. The FM-index can be used to count the frequency of any sequence of tokens n in $O(|n|\log|V|)$, *i.e.*, independently from the size of the corpus itself. For constrained decoding, the list of possible token successors can be obtained in $O(|V|\log|V|)$. Internally, the FM-index relies on the Burrows-Wheeler Transform [Burrows and Wheeler, 1994], or *BWT*, an invertible transformation that permutes a string to make it easier to compress, defined as follows: all the rotations of the string are sorted lexicographically and laid out in a matrix; the last column of the matrix is the strings's *BWT*.¹ For example, given the string *CABAC*, the corresponding matrix would be:

F						L
$\6	<i>C</i>	<i>A</i>	<i>B</i>	<i>A</i>	<i>C</i>	C^5
A^2	<i>B</i>	<i>A</i>	<i>C</i>	$\$$	<i>C</i>	C^1
A^4	<i>C</i>	$\$$	<i>C</i>	<i>A</i>	<i>B</i>	B^3
B^3	<i>A</i>	<i>C</i>	$\$$	<i>C</i>	<i>A</i>	A^2
C^5	$\$$	<i>C</i>	<i>A</i>	<i>B</i>	<i>A</i>	A^4
C^1	<i>A</i>	<i>B</i>	<i>A</i>	<i>C</i>	$\$$	$\6

¹Since our corpus contains multiple documents, we concatenate them with a separator token.

where $\$$ is a special end-of-string token. The first (**F**) and last (**L**) columns are the only ones that will be explicitly stored in the FM-index; **F** is just an array of runs (*i.e.*, sequences of repeated tokens), due to the rotations being sorted, so it can be represented with one count for each alphabet symbol; **L**, the string's BWT, will be stored in a data structure known as the Wavelet Tree [Grossi et al., 2003], which allows efficient rank-select-access queries, while exploiting the compressibility induced by the transformation. FM-indices have the useful property that for each symbol, the relative rank stays the same: that is, the i th occurrence of a symbol σ in **F** points to the same location in the corpus of the i th occurrence of σ in **L**. Thanks to this property, we can locate any string $\langle \sigma_1, \sigma_2, \dots, \sigma_n \rangle$ in the index by starting from σ_n and going backwards. First, we select the contiguous range of rows corresponding to the symbol σ_n in **F**, then we check the ranks of the first and last occurrences of σ_{n-1} in the same range of rows in **L**. We use the ranks to select a new, smaller or equal range of rows looking up the symbol σ_{n-1} in **F**. The procedure can be applied iteratively to find ngrams of any size.

4 Method

In our retrieval methodology, SEAL, we generate multiple ngrams, conditioning on a query. The ngrams are then used to find the documents they appear in within the corpus, which are then returned to the user. In Figure 1 we show this process at a high-level. We use our indexing structure, *i.e.*, the FM-index to do constrained decoding so that each ngram occurs at least once in the retrieval corpus. Jointly, we use the FM-index to efficiently find matching documents. Documents are ranked using the scores of the generated ngrams.

Autoregressive Retrieval We generate ngrams identifiers with constrained beam search, using the FM-index to identify the set of possible next tokens in at most $O(|V|\log|V|)$: tokens corresponding to unattested continuations are blocked by masking the logit to $-\infty$. As a result, after a single decoding pass, we get a set of ngrams (K), along with their autoregressively-computed probabilities according to the model. It is also trivial to find the positions in the corpus where the decoded ngrams appear, as constrained decoding already requires selecting the relevant range of rows in the FM-index. Note that autoregressive scoring entails monotonically decreasing scores—any string will be assigned a lower probability than any of its prefixes. To address this issue, we use fixed-length ngrams. Each document is assigned the score ($P(n|q)$) of its most probable decoded occurring ngram, *i.e.*, the probability of the ngram as autoregressively computed by the decoder, conditioned on the encoder input. We refer to this as the **LM** scoring.

Factoring in FM-index frequencies To counterbalance the monotonic probability decrease, we integrate in scoring unconditional ngram probabilities, computed as normalized index frequencies:

$$P(n) = \frac{F(n, \mathcal{R})}{\sum_{d \in \mathcal{R}} |d|} \quad (1)$$

This also enables us to promote *distinctive* ngrams, *i.e.*, those that have high probability according to the model and low probability according to the FM-index. We take inspiration from the theory behind TF-IDF and BM25 [Robertson and Zaragoza, 2009] and use the following scoring function:

$$w(n, q) = \max \left(0, \log \frac{P(n|q)(1 - P(n))}{P(n)(1 - P(n|q))} \right) \quad (2)$$

This formulation addresses the problem of length, as the unconditional probability of an ngram will also be equal or lower than that of any of its prefixes. To make better use of the computational resources, we slightly modify the beam search implementation to keep track of all the partially decoded sequences that have been considered. Thanks to this, we score a larger number of ngrams than the size of the beam. We refer to this formulation as the **LM+FM** scoring.

An Intersective Scoring for Multiple Ngrams One problem with the previous scoring formulations is that it is impossible to break ties among documents whose highest scoring ngram is the same, as they receive exactly the same score. Moreover, it might be difficult to capture all relevant information within a document by considering only a single ngram, for instance when salient ngrams are non-contiguous (*e.g.*, separated by unrelated text). To address these issues we propose a novel scoring

formulation that aggregates the contribution of multiple ngrams contained in the same document. To avoid repeated scoring of overlapping ngrams, for each document $d \in \mathcal{R}$ we only consider a subset of the generated ngrams $K^{(d)} \subset K$. An ngram n belongs to $K^{(d)}$ if there is at least one occurrence of n in d that does *not* overlap in the corpus with an occurrence of another ngram n' such that a) $n' \in K^{(d)}$ b) $w(n', q) > w(n, q)$. The document-level score, then, is the weighted sum of all ngrams in $K^{(d)}$:

$$W(d, q) = \sum_{n \in K^{(d)}} w(n, q)^\alpha \cdot \text{cover}(n, K^{(d)}) \quad (3)$$

where α is a hyperparameter and the coverage weight $\text{cover}(n, K)$ (controlled by the second hyperparameter β) is a function of how many ngram tokens are not included in the coverage set $C(n, K) \subset V$, *i.e.*, the union of all tokens in ngrams with a higher score. We define this coverage weight as follows:

$$\text{cover}(n, K) = 1 - \beta + \beta \cdot \frac{|\text{set}(n) \setminus C(n, K)|}{|\text{set}(n)|} \quad (4)$$

where $\text{set}(n)$ is the set of all tokens in n . The purpose of the coverage weight is to avoid the overscoring of very repetitive documents, where many similar ngrams are matched. Note that by saving the probability distribution at the first decoding step we can compute scores for all unigrams with no additional forward pass. We refer to this last approach, which can be thought of as a higher-order generalization of the bag-of-words assumption, as the **LM+FM intersective** scoring.

5 Experimental Setting

Our experimental setting evaluates SEAL on English knowledge-intensive NLP tasks. Each considered dataset is a collection of queries, each of which can be answered by looking for piece(s) of evidence in the corpus. We consider both an *in vivo* evaluation, in which we assess the model by looking at how well the document ranking matches with the ground truth, and, in addition, we perform a downstream evaluation, in which we feed the retrieved documents to a trained reader, that uses the documents to generate the answer.

5.1 Data

Natural Questions Natural Questions (NQ) is dataset containing query-document pairs, where the query is a question (*e.g.*, “who wrote photograph by ringo starr”), and the document is a Wikipedia page, in which a span is marked as an answer [Kwiatkowski et al., 2019]. We experiment on both the customary retrieval setup used by, among others, Karpukhin et al. [2020] and Mao et al. [2021], and the substantially different setup used by Tay et al. [2022]. We refer to these two settings as, respectively, **NQ** and **NQ320k**. In NQ, retrieval is performed on an entire Wikipedia dump, chunked in around 21M passages of 100 tokens. Performance is measured as $\text{accuracy}@k$, *i.e.*, the fraction of instances for which at least one of the top- k retrieved passages contains the answer. **NQ320k** is a much more restricted setting, in which the retrieval set is limited to the union of all ground truth document in the training, dev or test set. Different revisions of the same Wikipedia page count as different documents. Note that the exact splits used by Tay et al. [2022], the retrieval corpus and the preprocessing code have not been yet released at the time of writing. Therefore, we have tried to replicate the setting as closely as possible, but the exact numbers are not precisely comparable with those reported in the original paper. In **NQ320k**, performance is measured as $\text{hits}@k$, *i.e.*, the fraction of instances for which at least one of the top- k retrieved passages is in the ground truth.

KILT is a comprehensive benchmark collecting different datasets including question answering, fact checking, dialogue, slot filling, and entity linking [Petroni et al., 2021]. All these tasks are solvable by retrieving information from a unified corpus — a Wikipedia dump. In KILT, the evidence is usually the paragraph that contains the answer. Following Maillard et al. [2021], we have rechunked KILT’s retrieval corpus, which is originally paragraph-based, in around 36M passages of 100 tokens. We do not use the entity linking and ELI5 KILT tasks, where a ground truth passage is not provided in the training set. KILT’s retrieval performance is measured with R-precision, a precision-oriented measure that considers only gold documents as correct answers, not just any document containing the answer. R-precision can be computed at either passage level or at page level.

Table 1: Language model and index size on Natural Questions (around 21M passages). SEAL’s index is ~1.5 times smaller than uncompressed plain text.

System	Model Params	Size	Index Params	GPU?
<i>plain text</i>	-	13.4GB	-	-
DPR	220M	64.6 GB	16.1B	✓
BM25	-	18.8 GB	-	✗
GAR	406M	18.8 GB	-	✗
DSI-BART	406M	-	-	-
SEAL	406M	8.8GB	-	✗

5.2 SEAL configuration

Training We finetune BART large [Lewis et al., 2019] to generate ngrams of length $k = 10$ from the ground truth document. Since there are $|d| - k$ ngrams in a document d , we sample (with replacement) 10 ngrams from it, biasing the distribution in favor of ngrams with a high character overlap with the query. We also add the title of the document to the set of training ngrams. To expose the model to more possible pieces of evidence, we also add different “unsupervised” examples for each document in the retrieval corpus to the training set. In each of these examples the model takes as input a uniformly sampled span from the document, and predicts either another sampled span, or the title of the page. We append special tokens to the input to signal to the model a) whether the pair comes from the supervised or unsupervised training pairs (in the same spirit as the co-training task prompts used by Tay et al. [2022]) b) whether a title or span is expected as output. On KILT we train SEAL on all datasets at once. We report training and inference hyperparameters in Appendix (§A).

How to choose ngrams In preliminary experiments, we have found that training SEAL with ngrams of fixed length sampled uniformly from the training chunk did not lead to good performances. Instead, we have biased the ngram distribution towards ngrams that were relevant to the query. As a simple proxy for relevance, we have used the Levenshtein (character-based) distance between the query and the ngram overlap. The final distribution to sample from is defined by the following formula:

$$\frac{e^{L(q, d_{i:i+k})/\tau}}{\sum_{j=1}^{|d|-k+1} e^{L(q, d_{j:j+k})/\tau}} \quad (5)$$

where L is the Levenshtein distance and τ ($= 1.5$ in our experiments) is a temperature parameter, controlling the peakiness of the distribution.

Index We use the C++ FM-index implementation in `sds1-lite`. While the FM-index construction (which requires a sort of all rotations) takes around 6 hours in our single-threaded implementation, parallel algorithms are available [Labeit et al., 2017]. Each document is encoded as the subword tokenization of the concatenation of the title and the passage, separated by a special token. We report in Table 1 the index statistics for Natural Questions. As can be seen, SEAL’s FM-index is more than 7 times lighter compared to DPR’s full document embeddings for exact inner product search, and needs neither a GPU for search on top of that, nor separate storage for the text itself. While vector compression methods can reduce dense retrievers’ index size, this still comes at the expense of performance [Yamada et al., 2021, Lewis et al., 2021a]. In addition, our the size of our index is less than 50% of that of the well-optimized Lucene BM25 index used by `pyserini`, but also roughly 65% of the uncompressed plain text itself.

5.3 Retriever Baselines

We compare SEAL against well-established systems in the literature on each benchmark. On NQ and NQ320k we also compare against our BART-based replication of DSI [Tay et al., 2022, **DSI-BART**].² On NQ320k, a page-level benchmark, we include our own replication of GENRE [De Cao et al., 2021b]. Unless otherwise specified, we use `pyserini` to compute the BM25 baseline. For other systems, we either take figures from the literature, or use publicly released model predictions.

²We report reproduction details in Appendix (§B).

Table 2: Results on NQ320k. Reporting hits@1 and hits@10. Best in bold.

System	hits@k	
	1	10
BM25 (gensim)	15.3	44.5
BM25	22.7	59.0
DSI-BART	25.0	63.6
GENRE	26.3	71.2
SEAL (LM, n = 3)	21.3	66.5
SEAL (LM, n = 4)	22.2	68.2
SEAL (LM, n = 5)	22.6	68.7
SEAL (LM+FM)	25.3	72.0
SEAL (LM+FM, intersect.)	26.3	74.5

Table 3: Retrieval results on the NQ test set. Column blocks (left to right): retrieval results (accuracy@5/20/100); retrieval results on the test splits of Lewis et al. [2021b], partitioned according to whether the query/answer is a paraphrase of one in the training set; downstream performances (exact match). Except for Izacard and Grave [2021], all downstream results are computed with the same FiD reader trained on DPR predictions. Best in bold.

System	accuracy@k			Overlap? (A@100)				EM
	5	20	100	ans. ✓	✗	ques. ✓	✗	
BM25	43.6	62.9	78.1	82.9	70.1	80.9	76.6	40.4
DPR [Karpukhin et al., 2020]	68.3	80.1	86.1	91.4	76.8	93.2	83.2	47.2
GAR [Mao et al., 2021]	59.3	73.9	85.0	91.6	74.4	94.1	80.4	46.2
DSI-BART	28.3	47.3	65.5	77.8	44.2	84.9	57.7	31.4
Izacard and Grave [2021]	-	-	-	-	-	-	-	48.2
SEAL (LM, n = 5)	40.5	60.2	73.1	82.2	57.1	85.2	64.9	36.0
SEAL (LM+FM)	43.9	65.8	81.1	86.9	70.9	89.5	78.1	42.9
SEAL (LM+FM, intersepective)	61.3	76.2	86.3	91.2	77.7	93.2	84.1	48.0

5.4 Reader

For downstream results, we use the Fusion-in-Decoder abstractive reader [Izacard and Grave, 2021], which takes in the query along with 100 contexts and produces a task-specific answer. We train FiD on training set predictions.

6 Results

NQ320k We report results on NQ320k in Table 2. SEAL outperforms BM25 and DSI-BART in hits@10 in all its formulations. When taking into account ngram frequencies (*i.e.*, LM+FM), SEAL achieves even higher results than GENRE, despite the fact that this benchmark only requires page-level retrieval capabilities (that is the focus of GENRE). Finally, our intersepective formulation achieves the highest results, both in hits@1 and @10, indicating that multiple ngrams identifiers might capture complementary information, which can be aggregated for stronger performances.

Natural Questions We report in Table 3 the results of our evaluation on Natural Questions, a passage-level retrieval benchmark with a larger collection of documents (*i.e.*, ~21M w.r.t. 200k in NQ320k). In this setting, the gap in performance between DSI-BART and SEAL is larger, possibly because memorizing documents identifiers in the parameters of the model becomes more challenging with larger corpora. Remarkably, the intersepective formulation of SEAL achieves results comparable or superior to more established retrieval paradigms (*e.g.*, BM25, DPR and GAR), at high-recall (accuracy@100). To better understand the generalization capabilities of our retrieval solution we use the question/answer overlap split of Lewis et al. [2021b]. This study reveals that SEAL achieves the highest performance for question/answer pairs never seen during training (*i.e.*, no overlap), suggesting a better ability to generalize to completely novel questions with novel answers (*e.g.*, 3.5 points better than GAR on average).

Table 4: Retrieval results on individual KILT dev set(s), with the average in the rightmost column. Reporting passage-level R-precision (higher is better). We mark model that are also trained on additional synthetic data [Lewis et al., 2021c] with †. All SEAL models are multitask. Best among models trained only on KILT queries in bold.

Model	FEV	T-REx	zsRE	NQ	HoPo	TQA	WoW	AVG
BM25	40.1	51.6	53.0	14.2	38.4	16.2	18.4	33.1
DPR Maillard et al. [2021]	43.9	58.5	78.8	28.1	43.5	23.8	20.7	42.5
MT-DPR [Maillard et al., 2021]	52.1	53.5	41.7	28.8	38.4	34.2	24.1	39.0
MT-DPR [Oğuz et al., 2021]	52.1	61.4	54.1	40.1	41.0	34.2	24.6	43.9
MT-DPR† [Oğuz et al., 2021]	61.4	68.4	73.3	44.1	44.6	38.9	26.5	51.0
MT-DPR† (large) [Oğuz et al., 2021]	62.8	66.6	66.9	42.6	42.1	37.9	23.4	48.9
SEAL (LM+FM)	31.5	42.0	34.0	21.7	24.7	21.4	17.6	27.6
SEAL (LM+FM, intersective)	67.8	58.9	78.8	43.6	54.3	41.8	36.0	54.5

Table 5: Downstream results on the KILT test set(s). Downstream metrics are accuracy (FEVER, T-REx, zero-shot RE), exact match (Natural Questions, HotpotQA, TriviaQA), or F1 (Wizard of Wikipedia). Best in bold. †: result taken from the eval.ai KILT leaderboard.

System	FEV	T-REx	zsRE	NQ	HoPo	TQA	WoW
	ACC	ACC	ACC	EM	EM	EM	F1
KGI [Glass et al., 2021]†	85.6	84.4	72.6	45.2	-	61.0	18.6
Hindsight [Paranjape et al., 2021]	-	-	-	-	-	-	19.2
DPR+BART [Petroni et al., 2021]	86.7	59.2	30.4	41.3	25.2	58.6	15.2
RAG [Petroni et al., 2021]	86.3	59.2	44.7	44.4	27.0	71.3	13.1
MT-DPR+BART [Maillard et al., 2021]	86.3	-	58.0	39.8	31.8	59.6	15.3
MT-DPR+FiD [Piktus et al., 2021]	89.0	82.5	71.7	49.9	36.9	71.0	15.7
MT-DPR-WEB+FiD [Piktus et al., 2021]	89.0	81.7	74.2	51.6	38.3	72.7	15.5
SEAL+FiD (LM+FM)	87.9	83.7	74.2	47.3	37.6	65.8	17.5
SEAL+FiD (LM+FM, intersective)	89.5	83.6	74.7	53.7	40.5	70.9	18.3

Table 6: Retrieval results on the NQ test set with different model sizes. DPR (large) performance from [Oğuz et al., 2021].

System	# Par.	A@20	A@100
DPR (base)	~220M	80.1	86.1
DPR (large)	~350M	80.2	86.7
SEAL (large)	~400M	76.2	86.3

KILT We report retrieval results at passage level on the KILT benchmark in Table 4.³ SEAL outperforms DPR by more than 10 points on average in passage-level R-precision, indicating that our method is more precise in surfacing ground truth evidence as the first result. Moreover, SEAL also performs better than MT-DPR (multi-task DPR) even when the latter is pretrained on tens of millions of questions from PAQ [Lewis et al., 2021c], a technique that can drastically improve results and that could potentially bring benefits to our method as well (a task we leave for future work). When it comes to downstream performances (Table 5), FiD with passages retrieved by intersective SEAL establishes a new state-of-the-art on 4 datasets out of 7 (FEVER, zsRE, NQ, HoPo), and achieves very competitive results on the remaining 3.

Speed and constrained decoding The inference speed of SEAL is directly proportional to the beam size, with a limited overhead added by constrained decoding. On the Natural Questions test set, for instance, retrieval with the intersective scoring requires on our 1 GPU evaluation setup ~16 minutes and ~35 minutes with, respectively, a beam size of 5 or 15. Mao et al. [2021] report a lower runtime for GAR (~5 minutes), and a comparable one for DPR (~30 minutes). Note that more efficient approaches to constrained decoding have been proposed (*e.g.*, De Cao et al. [2021a])

³We report page-level and KILT-score results in the Appendix (§C).

Table 7: Ablation on Natural Questions. SEAL when using (✓) or not using (✗) FM-index constrained decoding, for beam size values in {3, 5, 10, 15}. Reporting accuracy@k.

System	Constr.	Beam	A@20	A@100
SEAL	✓	15	65.8	81.1
(LM+FM)	✗	15	65.3	80.1
	✓	3	63.3	78.0
	✓	5	64.7	79.9
	✓	10	65.4	80.8
SEAL	✓	15	76.2	86.3
(LM+FM, intersective)	✗	15	76.2	86.2
	✓	3	75.2	84.9
	✓	5	75.9	85.8
	✓	10	76.4	86.4

Table 8: Performance on Natural Questions with different max ngrams sizes, using SEAL (LM+FM, intersective).

Length	A@20	A@100
3	64.7	74.8
5	73.6	83.7
10	76.2	86.3

and we leave their application to SEAL as future work. Moreover, generation is becoming the de facto standard approach to NLP, not just as the method for materializing final outputs, but also for modeling the computational process needed before computing the answer [Wei et al., 2022]. As such, we expect generation latency will improve significantly and increasingly over time as a result of this growing interest. Any improvement will be directly applicable to SEAL.

Model size In Table 4 we show that parameter size is not a crucial factor behind the good performance of the method that we propose: SEAL (~400M) outperforms all model from Oğuz et al. [2021], including the MT-DPR large model (~350M) trained on NQ and PAQ [Lewis et al., 2021c]. In fact, the results of MT-DPR are *lower* with a bigger backbone—which could point to an increased difficulty in training larger models for dense retrieval. In the same direction, on vanilla NQ (Table 6), where DPR large only slightly outperforms DPR base, the performance of SEAL is in the same ballpark.

Ablation studies In Table 7 we report results on NQ for various configurations of SEAL. While, in general, performances increase with a larger beam, diminishing returns (or even a performance decrease) are found between a value of 10 and 15. Disabling constrained decoding and discarding a posteriori all generated ngrams that don’t appear in the corpus results in slightly lower performances. We have found the decoding maximum length to have a crucial impact on the retrieval performance of SEAL, since shorter ngrams tend to be less informative than longer ones. We report in Table 8 performances on the NQ test set when using, respectively, 3, 5, and 10 as ngram maximum length. Every other SEAL figure reported in this paper uses 10.

Qualitative Analysis In Table 9, we show examples of ngrams predicted by SEAL (trained on KILT) given the query “can you predict earthquakes”. SEAL is able to rephrase the query in ways that preserve its lexical material producing ngrams such as *earthquakes can be predicted, used to predict earthquakes* etc. Moreover, the model is also able to explore more diverse regions of the output space, overcoming the vocabulary mismatch problem: ngrams contain related tokens like the subword *seism-* and the word *forecast*. SEAL’s LM+FM scoring is also able to assign a score below 0 (and, thus, exclude from the search), unrelated ngrams that are considered by the beam because of their promising start, such as “Seismic risk in Malta @@”.

Table 9: Best (top) and worst (bottom) generated keys for the query “can you predict earthquakes” (left), and retrieved documents (right). Matched ngrams in bold. “@@” separates title and body.

score	#	identifier	doc #1	doc #2
273.2	1	earthquakes can be predicted	Seismology @@ for precise earthquake predictions , including the VAN method. Most seismologists	Earthquake prediction @@ reliably identified across significant spatial and temporal scales. While part of the
272.7	75	Earthquake prediction @@	do not believe that a system to provide timely warnings for individual	scientific community hold that, taking into account non- seismic precursors
269.9	3	predicted earthquakes	earthquakes has yet been developed, and many believe that such a system would be unlikely to give useful	and given enough resources to study them extensively, prediction might
229.7	11	Earthquake forecasting @@	warning of impending seismic events. However, more general forecasts routinely predict seismic hazard . Such	be possible , most scientists are pessimistic and some maintain that earthquake prediction is inherently impossible. Predictions are deemed significant if they can be shown to be successful beyond random chance.[...]
217.2	2	prediction Earthquake	forecasts estimate the probability of an earthquake of a particular [...]	
211.5	1	used to predict earthquakes		
205.3	7	earthquakes. Earthquake		
-	-	-		
-77.0	9	Seismic metamaterial @@		
-97.4	14	Seismic risk in Malta @@		
-113.4	3	Quaternary (EP) @@		
-150.3	1	used to predict the locatio[...]		
-301.5	17	Precipice (Battlestar Gala[...])		

7 Discussion

With SEAL we present solution that could potentially find applications outside information retrieval (*e.g.*, enforce generated substrings come from a white list of trusted sources). While we conduct our experiments with a model of ~400M parameters (*i.e.*, BART) for fast iterations, we believe the use of larger models could considerably improve performance. Changing the model would not affect the size of the index nor the cost of using it — $O(|n|\log|V|)$ for finding an ngram n . Moreover, we believe that indexing very large corpora (*e.g.*, the web) could be done more efficiently than existing attempts (*e.g.*, Piktus et al. [2021]) given the light memory footprint. Finally, dynamic variants [Gerlach, 2007, Salson et al., 2009] could allow the update of the FM-index on the fly without the need of re-indexing. While out of the scope of the current paper, we plan to tackle some of these scaling challenges in future work.

8 Conclusion

In this paper we present SEAL, a novel retrieval system that combines an autoregressive language model with a compressed full-text substring index. Such combination allows to constrain the generation to existing ngrams in a corpus and to jointly retrieve documents containing them. Empirically, we show an improvement by more than 10 points in average passage-level R-precision on KILT, and establish new state-of-the-art downstream performance on 4 out 7 datasets when paired with a reader model. While our results show that SEAL could already compete with more established retrieval systems, we believe there is potential in exploring existing (or yet to come) larger autoregressive models.

9 Broader Impact

While our decoding methodology provides a way to enforce corpus-based generation constraints, this does not fully prevent the generation and retrieval of non-factual or abusive text, since the corpus itself may well contain non-factual or abusive text. Moreover, even if a given corpus is misinformative or abusive, the language model could still produce undesirable text by “misquoting” material from it, *e.g.*, selectively copying an utterance that is being criticized in the original context of appearance. Applications should allow user to check the original supporting text. Furthermore, while we have only used a relatively small language model, scaling up the autoregressive model of SEAL would come with the same environmental risks that any other large system would pose, due to the more demanding energy requirements.

Acknowledgments and Disclosure of Funding

We thank Aleksandra Piktus, Edoardo Barba, Niccolò Campolungo, and Pere-Lluís Huguet Cabot for their helpful comments and suggestions.

References

- Mikel Artetxe, Shruti Bhosale, Naman Goyal, Todor Mihaylov, Myle Ott, Sam Shleifer, Xi Victoria Lin, Jingfei Du, Srinivasan Iyer, Ramakanth Pasunuru, et al. Efficient large scale language modeling with mixtures of experts. *CoRR*, abs/2112.10684, 2021.
- Yang Bai, Xiaoguang Li, Gang Wang, Chaoliang Zhang, Lifeng Shang, Jun Xu, Zhaowei Wang, Fangshan Wang, and Qun Liu. Sparterm: Learning term-based sparse representation for fast text retrieval. *CoRR*, abs/2010.00768, 2020. URL <https://arxiv.org/abs/2010.00768>.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- M. Burrows and D. J. Wheeler. A block-sorting lossless data compression algorithm. Technical report, 1994.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. Reading wikipedia to answer open-domain questions. *CoRR*, abs/1704.00051, 2017. URL <http://arxiv.org/abs/1704.00051>.
- Zhuyun Dai and Jamie Callan. Context-aware sentence/passage term importance estimation for first stage retrieval. *CoRR*, abs/1910.10687, 2019. URL <http://arxiv.org/abs/1910.10687>.
- Zhuyun Dai and Jamie Callan. Context-aware document term weighting for ad-hoc search. In Yennun Huang, Irwin King, Tie-Yan Liu, and Maarten van Steen, editors, *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, pages 1897–1907. ACM / IW3C2, 2020. doi: 10.1145/3366423.3380258. URL <https://doi.org/10.1145/3366423.3380258>.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. Highly parallel autoregressive entity linking with discriminative correction. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7662–7669, Online and Punta Cana, Dominican Republic, November 2021a. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.604. URL <https://aclanthology.org/2021.emnlp-main.604>.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. Autoregressive entity retrieval. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021b. URL <https://openreview.net/forum?id=5k8F6UU39V>.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. Wizard of wikipedia: Knowledge-powered conversational agents. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- Hady Elsahar, Pavlos Vougiouklis, Arslan Remaci, Christophe Gravier, Jonathon Hare, Elena Simperl, and Frederique Laforest. T-rex: A large scale alignment of natural language with knowledge base triples. *LREC*, 2018.
- P. Ferragina and G. Manzini. Opportunistic data structures with applications. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 390–398, 2000. doi: 10.1109/SFCS.2000.892127.
- Thibault Formal, Carlos Lassance, Benjamin Piwowarski, and Stéphane Clinchant. SPLADE v2: Sparse lexical and expansion model for information retrieval. *CoRR*, abs/2109.10086, 2021a. URL <https://arxiv.org/abs/2109.10086>.
- Thibault Formal, Benjamin Piwowarski, and Stéphane Clinchant. SPLADE: sparse lexical and expansion model for first stage ranking. In Fernando Diaz, Chirag Shah, Torsten Suel, Pablo Castells, Rosie Jones, and Tetsuya Sakai, editors, *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 2288–2292. ACM, 2021b. doi: 10.1145/3404835.3463098. URL <https://doi.org/10.1145/3404835.3463098>.

- Luyu Gao, Zhuyun Dai, and Jamie Callan. COIL: Revisit exact lexical match in information retrieval with contextualized inverted list. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3030–3042, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.241. URL <https://aclanthology.org/2021.naacl-main.241>.
- Wolfgang Gerlach. Dynamic fm-index for a collection of texts with application to space-efficient construction of the compressed suffix array diplomarbeit im fach. Master’s thesis, University of Bielefeld, 2007.
- Michael Glass, Gaetano Rossiello, Md Faisal Mahbub Chowdhury, and Alfio Gliozzo. Robust retrieval augmented generation for zero-shot slot filling. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1939–1949, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.148. URL <https://aclanthology.org/2021.emnlp-main.148>.
- Roberto Grossi, Ankur Gupta, and Jeffrey Scott Vitter. High-order entropy-compressed text indexes. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA ’03*, page 841–850, USA, 2003. Society for Industrial and Applied Mathematics. ISBN 0898715385.
- Gautier Izacard and Edouard Grave. Leveraging passage retrieval with generative models for open domain question answering. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 874–880, Online, April 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.eacl-main.74. URL <https://aclanthology.org/2021.eacl-main.74>.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *CoRR*, abs/2202.03629, 2022.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 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, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1147. URL <https://www.aclweb.org/anthology/P17-1147>.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 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, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.550. URL <https://aclanthology.org/2020.emnlp-main.550>.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Mojtaba Komeili, Kurt Shuster, and Jason Weston. Internet-augmented dialogue generation. *CoRR*, abs/2107.07566, 2021.
- 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. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466, 2019. doi: 10.1162/tacl_a_00276. URL <https://aclanthology.org/Q19-1026>.
- Julian Labeit, Julian Shun, and Guy E. Blelloch. Parallel lightweight wavelet tree, suffix array and fm-index construction. *Journal of Discrete Algorithms*, 43:2–17, 2017. ISSN 1570-8667. doi: <https://doi.org/10.1016/j.jda.2017.04.001>. URL <https://www.sciencedirect.com/science/article/pii/S1570866717300205>.

- Angeliki Lazaridou, Elena Gribovskaya, Wojciech Stokowiec, and Nikolai Grigorev. Internet-augmented language models through few-shot prompting for open-domain question answering. *CoRR*, abs/2203.05115, 2022.
- Jinhyuk Lee, Mujeen Sung, Jaewoo Kang, and Danqi Chen. 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, August 2021a. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.518. URL <https://aclanthology.org/2021.acl-long.518>.
- Jinhyuk Lee, Alexander Wettig, and Danqi Chen. Phrase retrieval learns passage retrieval, too. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3661–3672, Online and Punta Cana, Dominican Republic, November 2021b. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.297. URL <https://aclanthology.org/2021.emnlp-main.297>.
- Oleg Lesota, Navid Rekabsaz, Daniel Cohen, Klaus Antonius Grasserbauer, Carsten Eickhoff, and Markus Schedl. A modern perspective on query likelihood with deep generative retrieval models. In Faegheh Hasibi, Yi Fang, and Akiko Aizawa, editors, *ICTIR '21: The 2021 ACM SIGIR International Conference on the Theory of Information Retrieval, Virtual Event, Canada, July 11, 2021*, pages 185–195. ACM, 2021. doi: 10.1145/3471158.3472229. URL <https://doi.org/10.1145/3471158.3472229>.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. Zero-shot relation extraction via reading comprehension. *CoNLL*, 2017.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *CoRR*, abs/1910.13461, 2019.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive nlp tasks, 2020.
- Patrick Lewis, Barlas Oğuz, Wenhan Xiong, Fabio Petroni, Wen tau Yih, and Sebastian Riedel. Boosted dense retriever. *CoRR*, abs/2112.07771, 2021a.
- Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. 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, April 2021b. Association for Computational Linguistics. doi: 10.18653/v1/2021.eacl-main.86. URL <https://aclanthology.org/2021.eacl-main.86>.
- Patrick Lewis, Yuxiang Wu, Linqing Liu, Pasquale Minervini, Heinrich Küttler, Aleksandra Piktus, Pontus Stenetorp, and Sebastian Riedel. PAQ: 65 million probably-asked questions and what you can do with them. *Transactions of the Association for Computational Linguistics*, 9:1098–1115, 2021c. doi: 10.1162/tacl_a_00415. URL <https://aclanthology.org/2021.tacl-1.65>.
- Jimmy Lin and Xueguang Ma. A few brief notes on deepimpact, coil, and a conceptual framework for information retrieval techniques. *CoRR*, abs/2106.14807, 2021. URL <https://arxiv.org/abs/2106.14807>.
- Jean Maillard, Vladimir Karpukhin, Fabio Petroni, Wen-tau Yih, Barlas Oguz, Veselin Stoyanov, and Gargi Ghosh. Multi-task retrieval for knowledge-intensive tasks. 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 1098–1111, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.89. URL <https://aclanthology.org/2021.acl-long.89>.
- Antonio Mallia, Omar Khattab, Torsten Suel, and Nicola Tonellotto. Learning passage impacts for inverted indexes. In Fernando Diaz, Chirag Shah, Torsten Suel, Pablo Castells, Rosie Jones, and Tetsuya Sakai, editors, *SIGIR '21: The 44th International ACM SIGIR Conference on Research and*

- Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 1723–1727. ACM, 2021. doi: 10.1145/3404835.3463030. URL <https://doi.org/10.1145/3404835.3463030>.
- Yuning Mao, Pengcheng He, Xiaodong Liu, Yelong Shen, Jianfeng Gao, Jiawei Han, and Weizhu Chen. Generation-augmented retrieval for open-domain question answering. 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 4089–4100, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.316. URL <https://aclanthology.org/2021.acl-long.316>.
- Luca Massarelli, Fabio Petroni, Aleksandra Piktus, Myle Ott, Tim Rocktäschel, Vassilis Plachouras, Fabrizio Silvestri, and Sebastian Riedel. How decoding strategies affect the verifiability of generated text. *CoRR*, abs/1911.03587, 2019.
- Donald Metzler, Yi Tay, Dara Bahri, and Marc Najork. Rethinking search: Making domain experts out of dilettantes. *SIGIR Forum*, 55(1), jul 2021. ISSN 0163-5840. doi: 10.1145/3476415.3476428. URL <https://doi.org/10.1145/3476415.3476428>.
- Reiichiro Nakano, Jacob Hilton, S. Arun Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. Webgpt: Browser-assisted question-answering with human feedback. *CoRR*, abs/2112.09332, 2021.
- Rodrigo Nogueira and Jimmy Lin. From doc2query to docttttquery. 2021. URL https://cs.uwaterloo.ca/~jimmylin/publications/Nogueira_Lin_2019_docTTTTquery.pdf.
- Rodrigo Nogueira, Wei Yang, Jimmy Lin, and Kyunghyun Cho. Document expansion by query prediction. *CoRR*, abs/1904.08375, 2019. URL <http://arxiv.org/abs/1904.08375>.
- Cicero Nogueira dos Santos, Xiaofei Ma, Ramesh Nallapati, Zhiheng Huang, and Bing Xiang. Beyond [CLS] through ranking by generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1722–1727, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.134. URL <https://aclanthology.org/2020.emnlp-main.134>.
- Barlas Oğuz, Kushal Lakhota, Anchit Gupta, Patrick Lewis, Vladimir Karpukhin, Aleksandra Piktus, Xilun Chen, Sebastian Riedel, Wen tau Yih, Sonal Gupta, and Yashar Mehdad. Domain-matched pre-training tasks for dense retrieval. *CoRR*, abs/2107.13602, 2021.
- Ashwin Paranjape, Omar Khattab, Christopher Potts, Matei Zaharia, and Christopher D. Manning. Hindsight: Posterior-guided training of retrievers for improved open-ended generation. *CoRR*, abs/2110.07752, 2021. URL <https://arxiv.org/abs/2110.07752>.
- Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. Language models as knowledge bases? *EMNLP*, 2019.
- Fabio Petroni, Patrick Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. How context affects language models’ factual predictions. *AKBC*, 2020.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim Rocktäschel, and Sebastian Riedel. KILT: a benchmark for knowledge intensive language tasks. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2523–2544, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.200. URL <https://aclanthology.org/2021.naacl-main.200>.
- Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Dmytro Okhonko, Samuel Broscheit, Gautier Izacard, Patrick Lewis, Barlas Oğuz, Edouard Grave, Wen-tau Yih, and Sebastian Riedel. The web is your oyster - knowledge-intensive NLP against a very large web corpus. *CoRR*, abs/2112.09924, 2021. URL <https://arxiv.org/abs/2112.09924>.

- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. Scaling language models: Methods, analysis & insights from training gopher. *CoRR*, abs/2112.11446, 2021.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *CoRR*, abs/1910.10683, 2019.
- Adam Roberts, Colin Raffel, and Noam Shazeer. How much knowledge can you pack into the parameters of a language model? *CoRR*, abs/2002.08910, 2020.
- Stephen E. Robertson and Hugo Zaragoza. The probabilistic relevance framework: BM25 and beyond. *Found. Trends Inf. Retr.*, 3(4):333–389, 2009. doi: 10.1561/15000000019. URL <https://doi.org/10.1561/15000000019>.
- Devendra Singh Sachan, Mike Lewis, Mandar Joshi, Armen Aghajanyan, Wen tau Yih, Joelle Pineau, and Luke Zettlemoyer. Improving passage retrieval with zero-shot question generation. *CoRR*, abs/2204.07496, 2022. URL <https://arxiv.org/abs/2204.07496>.
- M. Salson, T. Lécroq, M. Léonard, and L. Mouchard. A four-stage algorithm for updating a Burrows-Wheeler Transform. *Theoretical Computer Science*, 410(43):4350–4359, 2009. doi: 10.1016/j.tcs.2009.07.016. URL [articles/Final_Version_LaurentMouchard_TCS_SpecialIssue_MaximeCrochemore.pdf](https://arxiv.org/abs/0907.016).
- Kurt Shuster, Mojtaba Komeili, Leonard Adolphs, Stephen Roller, Arthur D. Szlam, and Jason Weston. Language models that seek for knowledge: Modular search & generation for dialogue and prompt completion. *CoRR*, abs/2203.13224, 2022.
- Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, et al. Using deep-speed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. *abs/2201.11990*, CoRR, 2022.
- Yi Tay, Vinh Quang Tran, Mostafa Dehghani, Jianmo Ni, Dara Bahri, Harsh Mehta, Zhen Qin, Kai Hui, Zhe Zhao, Jai Gupta, Tal Schuster, William W. Cohen, and Donald Metzler. Transformer memory as a differentiable search index. *CoRR*, abs/2202.06991, 2022.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. FEVER: a large-scale dataset for fact extraction and verification. In *NAACL-HLT*, 2018.
- Ellen M Voorhees et al. The trec-8 question answering track report. In *Trec*, volume 99, pages 77–82, 1999.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed H. Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *CoRR*, abs/2201.11903, 2022. URL <https://arxiv.org/abs/2201.11903>.
- Ikuya Yamada, Akari Asai, and Hannaneh Hajishirzi. Efficient passage retrieval with hashing for open-domain question answering. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 979–986, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-short.123. URL <https://aclanthology.org/2021.acl-short.123>.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answering. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, October–November 2018. doi: 10.18653/v1/D18-1259. URL <https://www.aclweb.org/anthology/D18-1259>.

Tiancheng Zhao, Xiaopeng Lu, and Kyusong Lee. SPARTA: Efficient open-domain question answering via sparse transformer matching retrieval. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 565–575, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.47. URL <https://aclanthology.org/2021.naacl-main.47>.

Shengyao Zhuang and Guido Zuccon. TILDE: term independent likelihood model for passage re-ranking. In Fernando Diaz, Chirag Shah, Torsten Suel, Pablo Castells, Rosie Jones, and Tetsuya Sakai, editors, *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 1483–1492. ACM, 2021. doi: 10.1145/3404835.3462922. URL <https://doi.org/10.1145/3404835.3462922>.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] Discussion in the Appendix (§9)
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [N/A]
 - (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The prediction code is available in the supplementary materials. For training, we used the popular `fairseq` library. We have reported hyperparameters in §5.2 and in the Appendix (§A).
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (b) Did you mention the license of the assets? [No]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]