

DyVo: Dynamic Vocabularies for Learned Sparse Retrieval with Entities

Thong Nguyen¹, Shubham Chatterjee², Sean MacAvaney³

Ian Mackie³, Jeff Dalton², Andrew Yates¹

¹University of Amsterdam, ²University of Edinburgh, ³University of Glasgow

Correspondence: t.nguyen2@uva.nl

Abstract

Learned Sparse Retrieval (LSR) models use vocabularies from pre-trained transformers, which often split entities into nonsensical fragments. Splitting entities can reduce retrieval accuracy and limits the model’s ability to incorporate up-to-date world knowledge not included in the training data. In this work, we enhance the LSR vocabulary with Wikipedia concepts and entities, enabling the model to resolve ambiguities more effectively and stay current with evolving knowledge. Central to our approach is a Dynamic Vocabulary (DyVo) head, which leverages existing entity embeddings and an entity retrieval component that identifies entities relevant to a query or document. We use the DyVo head to generate entity weights, which are then merged with word piece weights to create joint representations for efficient indexing and retrieval using an inverted index. In experiments across three entity-rich document ranking datasets, the resulting DyVo model substantially outperforms state-of-the-art baselines.¹

1 Introduction

Neural IR methods typically operate in two stages. Initially, a set of candidate documents is retrieved using a fast, computationally-efficient first-stage retrieval method that considers sparse or dense vector representations. These candidates are then re-ranked using more computationally-intensive scoring functions, such as those involving cross-encoders (Nogueira and Cho, 2019; MacAvaney et al., 2019; Nogueira et al., 2020; Sun et al., 2023).

Learned Sparse Retrieval (LSR) (Nguyen et al., 2023b; Formal et al., 2021, 2022) is a prominent neural method for first-stage retrieval. LSR encodes queries and documents into sparse, lexically-aligned representations that can be stored in an inverted index for fast retrieval. LSR offers several advantages over Dense Retrieval (DR), another

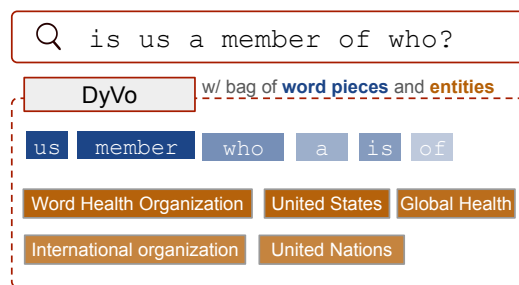


Figure 1: DyVo augments BERT’s word piece vocabulary with an entity vocabulary to help disambiguate a query (or document). Word pieces are in blue and entities are in orange. Darker terms have a higher weight in the sparse representation.

common approach for first-stage retrieval (Lin et al., 2020). LSR’s lexically grounded representations are more transparent, making it easier for users to understand the model and inspect representations for biases (Abolghasemi et al., 2024). Furthermore, LSR’s compatibility with an inverted index enables efficient and exact retrieval (Ding and Suel, 2011), while also simplifying the transition from existing lexical search infrastructure supporting methods like BM25. LSR not only performs competitively with DR in terms of performance within the same domain, but it also tends to generalize better across different domains and tasks (Formal et al., 2021).

However, LSR models lack explicit representations for entities and concepts in their vocabulary. This can pose challenges due to the tokenization process, where words are segmented into subwords or wordpieces. For instance, a word like “BioN-Tech” might be tokenized into [bio, ##nte, ##ch]. Such fragmentation can lead to ambiguity, complicating the retrieval process by obscuring the full meaning and context of the original word, which in turn may affect the accuracy and relevance of search results. Additionally, the bag of word pieces representation employed by LSR methods strug-

¹Code: <https://github.com/thongnt99/DyVo>

gles with homonyms, where different meanings or entities, such as “WHO” (World Health Organization) and “US” (United States), could be conflated when represented merely as word pieces in a query like “Is the US a member of WHO?”

Hence, while LSR provides a framework for efficient first-stage document retrieval, its current design – particularly in handling entities and complex vocabulary – poses significant challenges. We hypothesize that integrating explicit entities into the LSR vocabulary could significantly enhance its performance. This integration is especially pertinent as a large proportion of queries pertain to specific entities or are closely related to them (Kumar and Tomkins, 2010; Guo et al., 2009). Previous work indicates that hybrid models combining word and entity representations have improved both sparse retrieval (Dalton et al., 2014; Shehata et al., 2022; Mackie et al., 2024) and dense retrieval (Xiong et al., 2017a; Tran and Yates, 2022; Chatterjee et al., 2024).

To address the above limitations, we incorporate entities from Wikipedia into the vocabulary of an LSR model. The English Wikipedia contains entities spanning a diverse range of categories and disciplines, including named entities like people, organizations, and locations, as well as general concepts such as *eudaimonia*, *hot dog*, and *net income*. Integrating these Wikipedia entities into a LSR model significantly enhances its ability to handle complex semantic phrases and entities that are currently fragmented into nonsensical word pieces. By enriching query and document representations with relevant entities, we reduce ambiguity and improve the representational power of LSR. This approach is illustrated in Figure 1. Moreover, leveraging Wikipedia – a rich and continually updated knowledge base – allows the LSR model to refresh its internal knowledge, aligning it with evolving global information.

As of April 2024, the English Wikipedia hosts nearly 7 million entities and concepts, which is more than 200 times larger than the word piece vocabulary used in current state-of-the-art LSR methods. To identify relevant entities from among millions of them, we propose adding a Dynamic Vocabulary (DyVo) head with an entity candidate retrieval component. Specifically, we leverage entity retrieval techniques and Large Language Models (LLMs) to dynamically generate relevant entities. These methods aim to refine the set of highly rel-

evant entities, which are then passed to the LSR encoder for scoring. The encoder outputs a small bag of weighted entities, ignoring those that were not retrieved. The entity representation is then concatenated with the word-piece representation, forming a joint representation used for indexing and retrieval processes. Our contributions are:

- We propose the DyVo model to address the limitations of the word piece vocabulary commonly employed in LSR, which uses a Dynamic Vocabulary (DyVo) head to extend LSR to a large vocabulary (e.g., millions of Wikipedia entities and concepts) by leveraging existing entity embeddings and a candidate retrieval component that identifies a small set of entities to score.
- We introduce a few-shot generative entity retrieval approach capable of generating highly relevant entity candidates, which leads to superior performance when integrated into our DyVo framework. Furthermore, we find that document retrieval effectiveness using candidates generated by Mixtral or GPT4 is competitive with using entities identified by human annotators.
- We demonstrate that incorporating entities into LSR through a dynamic vocabulary consistently enhances the effectiveness of LSR across three entity-rich benchmark datasets (i.e., TREC Robust04, TREC Core 2018, and CODEC). Despite its simplicity, Wikipedia2Vec is a surprisingly effective source of entity embeddings. We achieve further performance gains by utilizing transformer-based dense entity encoders to encode entity descriptions into embeddings.

2 Related Work

Learned sparse retrieval. LSR encodes queries and documents into sparse lexical vectors, which are bag of words representations that are indexed and retrieved using an inverted index, akin to traditional lexical retrieval methods like BM25. One of the early works in this area proposed using neural networks to learn sparse representations that are compatible with an inverted index and demonstrated promising performance (Zamani et al., 2018). With the advent of the transformer architecture (Vaswani et al., 2017), subsequent work has successfully utilized pretrained transformers to enhance the effectiveness and efficiency of LSR models (Formal et al., 2021; Lassance and Clinchant, 2022; Formal et al., 2022; MacAvaney et al.,

2020; Zhao et al., 2020; Zhuang and Zuccon, 2021). Among these, SPLADE (Formal et al., 2021, 2022) stands out as a state-of-the-art LSR method. While SPLADE uses a word piece vocabulary, prior work has demonstrated that its vocabulary can be replaced by performing additional masked language modeling (MLM) pretraining and then exhaustively scoring all terms in the new vocabulary (Dudek et al., 2023). In this work, we dynamically augment a word piece vocabulary using pre-existing embeddings rather than performing additional pretraining. SPLADE typically employs a shared MLM encoder for both queries and documents, enabling term expansion and weighting on both sides. However, previous work (Nguyen et al., 2023b; MacAvaney et al., 2020) has shown that removing query expansion by replacing the MLM query encoder with an MLP encoder can simplify training and improve efficiency by reducing the number of query terms involved. While most LSR research has focused on ad-hoc paragraph retrieval tasks, recent efforts have explored extending LSR to other settings, such as conversational search (Hai Le et al., 2023), long documents (Nguyen et al., 2023a), and text-image search (Zhao et al., 2023; Chen et al., 2023; Nguyen et al., 2024).

Entity-oriented search. Early work in entity-oriented search primarily utilized entities for query expansion. A significant advancement in this domain was made by Meij et al. (2010), who introduced a double translation process where a query was first translated into relevant entities, and then the terms associated with these entities were used to expand the query. Dalton et al. (2014) further developed this concept with Entity Query Feature Expansion, which enhanced document retrieval by enriching the query context with entity features.

The field then recognized the more integral role of entities in search applications, transitioning from merely using entities for query expansion to treating them as a latent layer while maintaining the original document and query representations. Among these methods, Explicit Semantic Analysis (Gabrilovich and Markovitch, 2009) used “concept vectors” from knowledge repositories like Wikipedia to generate vector-based semantic representations. The Latent Entity Space model (Liu and Fang, 2015) utilized entities to assess relevance between documents and queries based on their alignments in the entity-informed dimensions. EsdRank (Xiong and Callan, 2015) leveraged semi-

structured data such as controlled vocabularies and knowledge bases to connect queries and documents, pioneering a novel approach to document representation and ranking based on interrelated entities.

This progression in research inspired a shift towards methodologies that treated entities not just as a latent layer but as explicit, integral elements of the retrieval model. For example, the creation of entity-based language models marked a significant development. Raviv et al. (2016) explored the impact of explicit entity markup within queries and documents, balancing term-based and entity-based information for document ranking. Ensan and Bagheri (2017) developed the Semantic Enabled Language Model, which ranks documents based on semantic relatedness to the query.

Xiong et al.’s line of work (Xiong et al., 2017b,a, 2018, 2017c) exemplifies a dual-layered approach that pairs a traditional bag of terms representation with a separate bag of entities representation, enhancing the document retrieval process by incorporating both term and entity-based semantics. For example, Explicit Semantic Ranking used a knowledge graph to create “soft matches” in the entity space, and the Word-Entity Duet Model captured multiple interactions between queries and documents using a mixture of term and entity vectors.

The Entity-Duet Ranking Model (EDRM) (Liu et al., 2018) represents a pioneering effort in neural entity-based search, merging the word-entity duet framework with the capabilities of neural networks and knowledge graphs (KGs). Tran and Yates (2022) advanced this area by introducing a method that clusters entities within documents to produce multiple entity “views” or perspectives, enhancing the understanding and interpretation of various facets of a document. Recently, Chatterjee et al. (2024) proposed to learn query-specific weights for entities within candidate documents to re-rank them.

Entity ranking. The task of entity ranking involves retrieving and ordering entities from a knowledge graph based on their relevance to a given query. Traditionally, this process has utilized term-based representations or descriptions derived from unstructured sources or structured knowledge bases like DBpedia (Lehmann et al., 2015). Ranking was commonly performed using models such as BM25 (Robertson and Zaragoza, 2009). Additionally, Markov Random Fields-based models like the Sequential Dependence Model (Metzler

and Croft, 2005) and its variants (Zhiltsov et al., 2015; Nikolaev et al., 2016; Hasibi et al., 2016; Raviv et al., 2012) addressed the joint distribution of entity terms from semi-structured data.

As the availability of large-scale knowledge graphs increased, semantically enriched models were developed. These models leverage aspects such as entity types (Kaptein et al., 2010; Balog et al., 2011; Garigliotti and Balog, 2017) and the relationships between entities (Tonon et al., 2012; Ciglan et al., 2012) to enhance ranking accuracy.

More recently, the focus has shifted towards Learning-To-Rank (LTR) methods (Schuhmacher et al., 2015; Graus et al., 2016; Dietz, 2019; Chatterjee and Dietz, 2021), which utilize a variety of features, particularly textual information and neighboring relationships, to re-rank entities. The introduction of graph embedding-based models like GEEER (Gerritse et al., 2020) and KEWER (Nikolaev and Kotov, 2020) has further enriched the field by incorporating Wikipedia2Vec (Yamada et al., 2020) embeddings, allowing entities and words to be jointly embedded in the same vector space.

The latest advancements in this domain have been driven by transformer-based neural models such as GENRE (Cao et al., 2021), BERT-ER++ (Chatterjee and Dietz, 2022), and EM-BERT (Gerritse et al., 2022). These models introduce sophisticated techniques including autoregressive entity ranking, blending BERT-based entity rankings with additional features, and augmenting BERT (Devlin et al., 2019) with Wikipedia2Vec embeddings.

3 Methodology

In this section, we describe our approach to producing sparse representations of queries and documents that contain both entities and terms from a word piece vocabulary. To do so, we incorporate entities into the model’s vocabulary through the use of a Dynamic Vocabulary head.

3.1 Sparse Encoders

Given a query q and a document d as input, an LSR system uses a query encoder f_q and a document encoder f_d to convert the inputs into respective sparse representations s_q and s_d . The dimensions are aligned with a vocabulary \mathcal{V} and only a small number of dimensions have non-zero values. Each dimension s_q^i or s_d^i encodes the weight of the i^{th} vocabulary item (v_i) in the input query or document, respectively. The similarity between a query and

a document is computed as the the dot product between the two corresponding sparse vectors:

$$\mathcal{S}(q, d) = f_q(q) \cdot f_d(d) = s_q \cdot s_d = \sum_{i=0}^{|\mathcal{V}|-1} s_q^i s_d^i \quad (1)$$

Various types of sparse encoders have been previously defined in the literature and summarized by Nguyen et al. (2023b). SPLADE (Formal et al., 2021, 2022; Lassance and Clinchant, 2022) is a state-of-the-art LSR method that employs the MLM architecture for both the query and document encoders. The strength of the MLM architecture is its ability to do term weighting and expansion in an end-to-end fashion, meaning that the model can itself learn from data to expand the input to semantically relevant terms and to weight the importance of individual terms. With an MLM encoder, the sparse representation of a query or document are generated as follows:

$$s_{(\cdot)}^i = \max_{0 \leq j < \mathcal{L}} \log(1 + \text{ReLU}(e_i \cdot h_j)) \quad (2)$$

where $s_{(\cdot)}^i$ and e_i are the output weight and the embedding (from the embedding layer) of the i^{th} vocabulary item respectively, \mathcal{L} is the length of the input query or document, and h_j is the the last hidden state of the j^{th} query or document input token produced by a transformer backbone, such as BERT (Devlin et al., 2019). A recent study (Nguyen et al., 2023b) found that it is not necessary to have both query and document expansion. Disabling query expansion, by replacing the MLM query encoder by an MLP encoder can improve model efficiency while keeping the model’s effectiveness. The MLP encoder weights each query input token as follows:

$$s_q^i = \sum_{0 \leq j < \mathcal{L}} \mathbb{1}_{v_i=q_j} (\mathcal{W} \cdot h_j^T + b) \quad (3)$$

where \mathcal{W} and b are parameters of a linear layer projecting a hidden state h_j to a scalar weight.

In this work, we employ this model variant with a MLP query encoder and a MLM document encoder as the baseline, and try to improve the model’s expressiveness by expanding the output vocabulary to Wikipedia entities. This model variant is similar to EPIC (MacAvaney et al., 2020) and SPLADE-v3-Lexical (Lassance et al., 2024), though it does not exactly correspond to either model. We call this model LSR-w to emphasize its use of the word piece vocabulary.

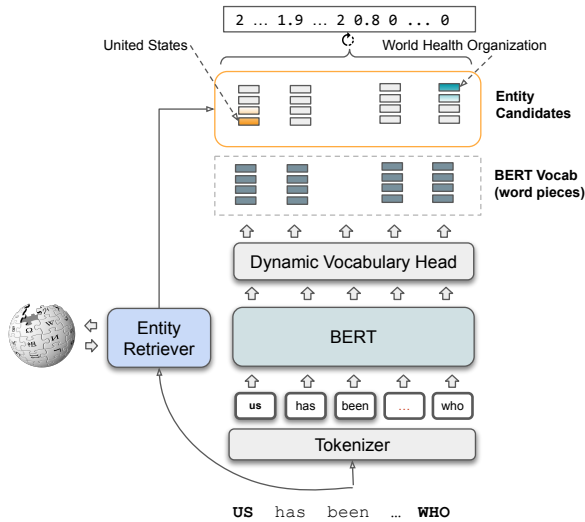


Figure 2: DyVo model with large entity vocabulary. The DyVo head scores entity candidates from an Entity Retriever component.

3.2 Entity Vocabulary

In this section, we describe our methodology to enrich the LSR vocabulary with Wikipedia entities. We build upon the MLM architecture for entity scoring in order to expand the input to any relevant items in the vocabulary, including entities which are not part of the encoder input. In the MLM head, the weight of the i -th entity with regard to an input query or document is calculated as follows:

$$s_{ent}^i = \lambda_{ent} \max_{0 \leq j < \mathcal{L}} \log(1 + \text{ReLU}(e_i^{entity} \cdot h_j)) \quad (4)$$

We calculate the dot product between the entity embedding e_i^{entity} and every hidden state h_j , and then select the maximum score. Via a ReLU gate, only positive weights are retained and then log scaled. For each query or document, only a small number of relevant entities have non-zero weights, forming a small bag of weighted entities (i.e., a sparse entity representation). This resulting entity representation is merged with the bag of words representation in the previous section to form a joint word-entity sparse representation. We add λ_{ent} (initialized as 0.05) as a trainable scaling factor to adjust the entity weights. This scaling factor is important to prevent training collapse as discussed in Appendix A.2

The final relevance score, which integrates both word and entity vocabularies, is computed as follows:

$$S(q, d) = \sum_{i=0}^{|V|-1} s_w^i(q) s_w^i(d) + \sum_{j=0}^{|E|-1} s_{ent}^j(q) s_{ent}^j(d) \quad (5)$$

where $s_w^i(\cdot)$ represents the weight of word v^i

and $s_{ent}^j(\cdot)$ represents the weight of entity e^j with regard to the input query or document.

3.3 Dynamic Vocabulary Head

It is not practical to add every entity to the existing MLM head, because the MLM head exhaustively scores every term in its vocabulary for each input vector. We propose a Dynamic Vocabulary (DyVo) head that augments an existing vocabulary using two ingredients: (1) embeddings of the new vocabulary terms (e.g., entity embeddings obtained from an external source) and (2) a candidate retrieval method that takes a query or document as input and identifies a small subset of the new vocabulary that may be present in the input (e.g., entities identified by an entity linker). We use a DyVo head to expand the sparse encoder’s vocabulary to include millions of Wikipedia entities, without the need to exhaustively score them as in Equation 4.

Entity embeddings. To produce a score for an entity in the vocabulary, the DyVo head needs to compute the dot product between the entity embedding and the hidden state of each input token. This operation requires both the entity embedding and the hidden states in the transformer backbone to have the same size and live in the same latent space. In this work, we chose DistilBERT (Sanh et al., 2019), which has proven its effectiveness in previous research, as the transformer backbone with an embedding size of 768. For our default entity embeddings, we utilize the **LaQue** pretrained dense entity encoder (Arabzadeh et al., 2024) to encode entity descriptions from KILT (Petroni et al., 2021) into entity embeddings. We choose LaQue for its consistent performance in yielding good entity weights and retrieval effectiveness in pilot experiments. We later provide detailed results comparing different types of entity embeddings.

Entity candidate retrieval. Instead of computing the weights for millions of entities in the vocabulary, we propose to add an entity candidate retrieval component (Figure 2) that aims to narrow down the search space to a small set of relevant entities, which are then scored by the LSR encoder using Equation 4. Offloading the entity retrieval task to a separate specialized component would allow the LSR model to focus entirely on the scoring task to maximize the document retrieval objective. While using linked entities is a popular option in prior research, this approach may overlook important entities that are not directly mentioned in the text. In-

stead, we introduce a few-shot generative approach that leverages the power of LLMs to generate high quality candidates, including both linked entities and relevant entities. For each query, we show two examples and prompt LLMs (Mixtral, GPT4) to generate a list of Wikipedia entities that are helpful to retrieve relevant documents. The prompt template is shown in Prompt A.2. We later compare our generative approach to various baselines.

Practical considerations. The DyVo head is memory-efficient when handling a large vocabulary, such as Wikipedia entities. At both training time and inference time, DyVo avoids instantiating sparse vectors with millions of dimensions, which would require a substantial amount of memory compared to the raw text (e.g., 10MB to store a single float16 vector with 5 million dimensions). During training, DyVo leverages the fact that the vast majority of entities do not appear in any given query (or document) to create a compact subset of the vocabulary for each batch. To do so, DyVo maintains a per-batch tensor of entity candidate IDs along with the corresponding entity weights, which are used to match entities between the query and the document. The weights of the matching entities are multiplied together and summed to produce the final relevance score. This allows DyVo to instantiate relatively small sparse vectors that contain enough dimensions to hold the entity candidates, rather than instantiating vectors that correspond to the entire vocabulary. Sparse representations are stored in an inverted index that is queried at inference time, so vocabulary-size vectors do not need to be instantiated at retrieval time.

4 Experimental setup

Datasets. Given our need for entity-rich queries and documents, we evaluate our approach using datasets containing a mix of news documents and complex information needs (i.e., TREC Robust04, TREC Core 2018, CODEC), which have also been commonly used in prior work, e.g. (Dalton et al., 2014; Chatterjee et al., 2024; Tan et al., 2023; MacAvaney et al., 2019; Nogueira et al., 2020; Li et al., 2023). **Robust04** (Voorhees et al., 2003) has 528k documents and 250 query topics where documents are news articles. All topics are deeply annotated with 1246 judged documents per topic on average. **Core 2018** contains 595k news articles or blog posts from The Washington Post with about 50 topics and 524 relevant judgements per topic.

CODEC (Mackie et al., 2022) provides 729k web documents crawled from various sources and 42 complex query topics, covering recent themes (e.g., bitcoins, NFT) from diverse domains, such as history, economics, politics. Furthermore, each topic comes with approximately 147 document judgements and 269 entity annotations.

We use all provided topics (*description* field on TREC datasets and *query* field on CODEC) for evaluation. To train models, we used synthesized dataset provided by InParsV2 (Jeronymo et al., 2023). Because CODEC is not available on InParsV2, we generate 10k queries ourselves using *Mixtral-8x7B-Instruct-v0.1* (Jiang et al., 2024).

Knowledge base and entity candidates. We use the KILT (Petroni et al., 2021) knowledge repository with 5.9 millions entities and only keep entities appearing in Wikipedia2Vec (Yamada et al., 2020), resulting in ~5.3 millions entities. To obtain linked entity candidates for queries, we use the REL (van Hulst et al., 2020) entity linker with n-gram NER tagger. For the entity retrieval approach, we experimented with different approaches, including traditional sparse retrieval (BM25), dense retrieval (LaQue), and generative retrievers (Mixtral and GPT4). For BM25, we index the entity’s description and retrieve the top 20 entities per query using Pyserini (Lin et al., 2021) with the default parameters. With LaQue, we encode both queries and entity descriptions using the LaQue (DistilBERT) dense encoder, and select the top 20 entities that have the highest dot product with the query’s dense vector. With generative approaches, we prompt Mixtral and GPT4 to generate relevant entities and remove out-of-vocabulary entities. For simplicity, we re-use the linked entities from Chatterjee et al. (2024) on the document side for all experiments.

Training configuration. Starting from a LSR checkpoint without entities trained on MSMARCO, we further fine-tune them on the three datasets using the synthesized queries, MonoT5-3b scores for distillation, KL loss (Formal et al., 2022) and BM25 negatives. To regularize vector sparsity, we apply a L1 penalty on the output sparse representations, which has previously been shown to be effective (Nguyen et al., 2023b). We experiment with different L1 weights, including [1e-3, 1e-4, 1e-5]. For each setting, we train two LSR versions: LSR-w that produces word piece representations only, and DyVo that produces joint word-entity rep-

representations. On each dataset, we train the models for 100k steps with a batch size of 16, learning rate of $5e-7$, and 16-bit precision on a single A100 GPU. Entity embeddings are pre-computed and frozen during training; only a projection layer is trained where the word and entity embedding sizes differ.

Evaluation metrics We report commonly used IR metrics, including $nDCG@10$, $nDCG@20$ and $R@1000$ on all three datasets using the *ir_measures* toolkit (MacAvaney et al., 2022).

5 Experimental results

We first consider whether incorporating linked entities in sparse representations increases effectiveness over representations containing only word pieces, finding that doing so yields consistent improvements on our entity-rich benchmarks. We then consider the impact of the entity selection component and the entity embeddings used, finding that performing entity retrieval rather than entity linking can further improve performance and that DyVo performs well with a range of entity embedding techniques.

RQ1: Does incorporating linked entities improve the effectiveness of LSR?

In this RQ, we seek to evaluate the effectiveness of LSR with linked entities. We train three different LSR versions with different sparse regularization weights ($1e-3$, $1e-4$, $1e-5$). For each LSR version, we trained two models (LSR-w and DyVo) with and without entities, respectively, using exactly the same training configuration. Although we are mainly interested in the comparison between DyVo and LSR-w, other baselines (e.g., BM25, BM25+RM3, and zero-shot single-vector dense retrieval methods) are provided in Table 1 to help readers position LSR with regard to other first-stage retrieval families.

Our first observation is that our model with linked entities (DyVo) outperforms the model without entities (LSR-w) consistently on all metrics ($nDCG@10$, $nDCG@20$, $R@1000$) across three different datasets and three different sparsity constraints. The difference between the two models is more pronounced when the document representations become more sparse. With the largest regularization weight ($reg=1e-3$), the documents are the most sparse and have the fewest terms. In this scenario, enriching the word representation with linked entities typically results in a significant gain,

notably with an increase ranging from 1.15 to 3.57 points in $nDCG@10$ across all datasets. When we relax the sparsity regularization to $1e-4$ and $1e-5$, we observe an improvement in the performance of LSR-w baseline models. However, we still consistently observe the usefulness of linked entities, albeit to a lesser degree. In the most relaxed setup ($reg=1e-5$), we often gain from 1 to 2 $nDCG$ points on all three datasets. The $R@1000$ improvement is similar, except we only observe a minimal increase on Core 2018.

Compared to other families, both LSR and DyVo demonstrate better performance than unsupervised lexical retrieval methods (BM25, BM25+RM3) and Dense Retrieval (DR) models, including DistilBERT-dot-v5 (Reimers and Gurevych, 2019a), GTR-T5-base (Ni et al., 2022b), and Sentence-T5-base (Ni et al., 2022a). Despite using models three times larger (T5-base vs. DistilBERT), both GTR-T5-base and Sentence-T5-base still show lower effectiveness than LSR models. This is due to the generalization difficulties of dense retrieval methods.

DyVo also outperforms BM25 + RM3, a traditional query expansion method using pseudo-relevance feedback. Compared to GRF, a LLM-based query expansion approach by Mackie et al. (2023), DyVo achieves a significantly higher $nDCG@10$ score (e.g., 53.40 with DyVo using the REL linker versus 40.50 with GRF on CODEC). It is important to note that the LLMs used in GRF were not fine-tuned, and doing so would present substantial computational challenges.

RQ2: Can LSR be more effective with retrieval-oriented entity candidates?

In the previous RQ, we explored how incorporating linked entities enhances LSR’s representations. However, relying solely on linked entities overlooks other relevant entities crucial for document retrieval. For instance, with the CODEC query “Why are many commentators arguing NFTs are the next big investment category?”, entities like “Cryptocurrency”, “Bitcoin”, and “Digital asset” can be valuable despite not being explicitly mentioned.

In this RQ, we aim to evaluate our few-shot generative entity retrieval approach based on Mixtral or GPT4 and compare it with other entity retrieval approaches, including entity linking (as explored in the previous RQ), sparse methods (BM25), dense entity retrieval methods (LaQue), and human annotations. The results are shown in Table 2.

Method	Reg	TREC Robust04			TREC Core 2018			CODEC		
		nDCG@10	nDCG@20	R@1k	nDCG@10	nDCG@20	R@1k	nDCG@10	nDCG@20	R@1k
<i>Unsupervised sparse retrieval</i>										
BM25		39.71	36.25	57.18	30.94	29.19	52.19	37.70	35.28	61.25
BM25 + RM3		43.77	40.64	64.21	35.82	34.79	60.09	39.93	39.96	65.70
<i>Zero-shot Dense Retrieval</i>										
DistilBERT-dot-v5		37.95	34.97	52.41	37.02	34.60	54.07	42.76	46.67	60.33
GTR-T5-base		43.79	39.33	54.35	38.81	36.51	57.62	48.42	54.01	66.96
Sentence-T5-base		44.06	39.60	57.64	43.18	39.54	60.88	44.22	32.10	65.48
<i>Learned Sparse Retrieval</i>										
LSR-w	1e-3	40.37	37.23	55.66	34.50	31.45	52.66	39.10	35.32	57.58
DyVo (REL)		41.52	38.62	56.78	37.50	34.61	54.14	42.67	38.32	59.81
LSR-w	1e-4	47.69	44.48	64.47	38.94	37.37	60.44	50.54	46.71	66.39
DyVo (REL)		48.15	44.85	64.72	43.10	39.46	60.43	51.66	47.95	68.49
LSR-w	1e-5	49.13	46.34	66.86	40.99	38.73	63.22	52.61	49.22	69.07
DyVo (REL)		51.19	47.65	68.56	43.72	40.56	63.56	53.40	51.15	70.60

Table 1: Results with linked entities. All LSR models use a DistilBERT backbone. DyVo uses entities found by the REL entity linker and LaQue entity embeddings. All documents are truncated to the first 512 tokens.

Observing the table, we note that DyVo (BM25) and DyVo (LaQue) show modest performance gains compared to the DyVo (REL) model, which incorporates linked entities, and the LSR model without entities. Employing LaQue-retrieved candidates to DyVo increases LSR-w’s R@1000 by +1.39 points (66.86 → 68.25), +1.61 points (63.22 → 64.83), and +1.8 points (from 69.07 → 70.87) on the Robust04, Core18, and CODEC datasets, respectively. This recall improvement is comparable to the gain achieved by using the REL entity linker. However, we generally observe no benefits in terms of nDCG when using the BM25 or LaQue retriever. This could be because BM25 and LaQue tend to prioritize recall, resulting in the retrieval of not only relevant entities but also noisy entities.

Our generative approach utilizing Mixtral and GPT4 represents a significant step forward in entity retrieval for document ranking. Compared to linked entities provided by REL, our approach showcases notable improvements, enhancing nDCG@10 and nDCG@20 scores by approximately +1.3 to +1.78 points across all datasets, with the exception of nDCG@10 on Core 2018. Mixtral’s effectiveness is further highlighted by its impact on R@1000 scores, with increases observed across the Robust04, Core 2018, and CODEC datasets.

Additionally, when we replace Mixtral with GPT4, we see further improvements that result in DyVo achieving the highest performance on nearly every metric and dataset. Notably, retrieval using GPT-4 generated entities is competitive with re-

trieval using human-annotated entities on CODEC, underlining the significance of enriching query representations with relevant entities beyond linked ones. We attach examples in Table 4 in the Appendix to illustrate the candidates retrieved by different systems.

RQ3: How does changing entity embeddings affect the model’s ranking performance?

Previously, we utilized the same entity encoder, LaQue (Arabzadeh et al., 2024), to generate entity embeddings. Here, our objective is to evaluate various approaches to obtain entity embeddings including Token Aggregation (i.e., splitting an entity’s surface form into word pieces and averaging their static embeddings), Wikipedia2Vec (Yamada et al., 2020), general dense passage encoders like JDS and DPR (Puran Ben Veyseh et al., 2021) and specialized dense entity encoders like LaQue and BLINK (Arabzadeh et al., 2024; Laskar et al., 2022). JDS is a joint dense ([CLS] vector) and sparse model with a shared DistilBERT backbone. We train our JDS model on MSMARCO dataset with a dual dense-sparse loss, using it to encode entity descriptions into dense embeddings. The results is shown in Table 3.

First, we observe that simply tokenizing the entity name into word pieces and averaging the transformer’s static token embeddings proves to be a viable method for creating entity embeddings. This approach typically yields a +1 point improvement over LSR-w across various metrics

Method	TREC Robust04			TREC Core 2018			CODEC		
	nDCG@10	nDCG@20	R@1k	nDCG@10	nDCG@20	R@1k	nDCG@10	nDCG@20	R@1k
LSR-w	49.13	46.34	66.86	40.99	38.73	63.22	52.61	49.22	69.07
DyVo (REL)	51.19	47.65	68.56	43.72	40.56	63.56	53.40	51.15	70.60
DyVo (BM25)	51.38	47.72	67.74	42.48	38.89	64.58	53.25	49.80	69.83
DyVo (LaQue)	49.42	46.31	68.25	40.24	38.39	64.83	53.73	50.34	70.87
DyVo (Mixtral)	52.97	49.21	69.28	43.80	41.86	68.27	54.90	52.82	73.20
DyVo (GPT4)	54.39	50.89	70.86	43.06	42.25	68.57	56.46	53.72	74.47
DyVo (Human)	-	-	-	-	-	-	56.42	52.96	75.13

Table 2: Results with entities retrieved by different retrievers. All models are trained with a DistilBERT backbone, LaQue entity embeddings, and L1 regularization (weight=1e-5).

Method	Entity Rep.	TREC Robust04			TREC Core 2018			CODEC		
		nDCG@10	nDCG@20	R@1k	nDCG@10	nDCG@20	R@1k	nDCG@10	nDCG@20	R@1k
LSR-w	-	49.13	46.34	66.86	40.99	38.73	63.22	52.61	49.22	69.07
DyVo (GPT4)	Token Aggr.	51.35	48.01	67.46	41.63	39.37	64.01	53.44	50.39	69.94
DyVo (GPT4)	DPR	48.68	45.77	75.21	40.26	37.52	70.81	53.04	49.18	75.19
DyVo (GPT4)	JDS	51.21	48.38	73.79	44.29	41.86	70.16	55.08	50.93	73.97
DyVo (GPT4)	Wiki2Vec	54.04	50.21	69.85	44.15	43.13	67.77	56.30	53.25	73.03
DyVo (GPT4)	LaQue	54.39	50.89	70.86	43.06	42.25	68.57	56.46	53.72	74.47
DyVo (GPT4)	BLINK	55.56	51.71	71.81	44.63	42.94	69.11	58.15	54.83	74.72

Table 3: Results with different entity embeddings. All models are trained with a DistilBERT backbone and L1 regularization (weight=1e-5). Entity candidates generated by GPT4 are used on queries for inference.

and datasets. We hypothesize that this improvement mainly stems from phrase matching through entity name matching, as we believe the token static embeddings do not encode much entity knowledge.

Interestingly, in terms of nDCG scores, this simple method outperforms the DPR and JDS methods, which rely on generic dense passage encoders trained for ad-hoc passage retrieval tasks to encode entity descriptions. DPR and JDS, however, demonstrate strong recall, suggesting that these encoders may prioritize encoding abstract entity information, which enables them to pull relevant documents within the top 1000 results. However, they may lack fine-grained entity knowledge necessary for more nuanced weighting.

Wikipedia2Vec (Wiki2Vec, dim=300), LaQue, and BLINK are specialized for entity representation learning or entity ranking tasks. As indicated in the last three rows of Table 3, using them to generate entity embeddings enhances document retrieval performance across all metrics and datasets. Despite being trained using a simple skip-gram model, Wikipedia2Vec effectively supports LSR in document retrieval, outperforming models utilizing aggregated token embeddings and dense passage encoders. The robustness of Wikipedia2Vec has been documented in prior research (Oza and Dietz, 2023). Substituting Wikipedia2Vec with more advanced transformer-based entity encoders such

as LaQue and BLINK results in the strongest overall performance. LaQue, based on the lightweight DistilBERT backbone, shows a slight improvement over Wikipedia2Vec. Using a larger transformer model (BERT-large), BLINK usually achieves a +1 nDCG point increase compared to LaQue, topping all datasets in terms of nDCG@10 and nDCG@20.

6 Conclusion

LSR has emerged as a competitive method for first-stage retrieval. In this work, we observed that relying on only word pieces for lexical grounding can create ambiguity in sparse representations—especially when entities are split into subwords. We explored whether learned sparse representations can include entity dimensions in addition to word piece dimensions. In order to facilitate modeling millions of potential entities, we proposed a Dynamic Vocabulary (DyVo) head that leverages entity retrieval to identify potential entity candidates and entity embeddings to represent them. We find that while both linked entities and LLM-generated entities are effective, LLM-generated entities ultimately yield higher LSR effectiveness. The approach is largely robust to the choice of entity embedding. Our work sets the stage for other LSR models that go beyond word piece vocabularies.

Limitations

While our approach is highly effective on the document retrieval benchmarks considered, it is important to note that its reliance on large language models (LLMs) like Mixtral and GPT4 can pose computational and cost inefficiencies. This challenge is not unique to our methodology; rather, it is a common concern across various research pursuits employing LLMs for retrieval purposes. One potential avenue for mitigating these costs involves leveraging LLMs to generate synthetic datasets and distill their internal knowledge into a more streamlined entity ranker or re-ranker. Addressing this issue extends beyond the scope of our current work.

Ethics Statement

We constructed our LSR encoder using a pretrained DistilBERT and employed Large Language Models such as Mixtral and GPT4 to generate entity candidates. Consequently, our models may inherit biases (e.g., preferences towards certain entities) encoded within these language models. Our evaluation encompasses both open-source models (Mixtral, DistilBERT, LaQue, BLINK, REL, Wikipedia2Vec) and proprietary ones (GPT4), which do not always disclose their training data.

Acknowledgements

This research was supported by the Hybrid Intelligence Center, a 10-year program funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research, <https://hybrid-intelligence-centre.nl>, and project VI.Vidi.223.166 of the NWO Talent Programme which is (partly) financed by the Dutch Research Council (NWO).

References

- Amin Abolghasemi, Leif Azzopardi, Arian Askari, Maarten de Rijke, and Suzan Verberne. 2024. Measuring bias in a ranked list using term-based representations. In *European Conference on Information Retrieval*, pages 3–19. Springer.
- Negar Arabzadeh, Amin Bigdeli, and Ebrahim Bagheri. 2024. Laque: Enabling entity search at scale. In *European Conference on Information Retrieval*, pages 270–285. Springer.
- Krisztian Balog, Marc Bron, and Maarten De Rijke. 2011. [Query modeling for entity search based on terms, categories, and examples](#). *ACM Transactions on Information Systems*, 29(4).
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021. [Autoregressive entity retrieval](#). In *International Conference on Learning Representations*.
- Shubham Chatterjee and Laura Dietz. 2021. [Entity Retrieval Using Fine-Grained Entity Aspects](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, page 1662–1666, New York, NY, USA. Association for Computing Machinery.
- Shubham Chatterjee and Laura Dietz. 2022. [Bert-er: Query-specific bert entity representations for entity ranking](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*, page 1466–1477, New York, NY, USA. Association for Computing Machinery.
- Shubham Chatterjee, Iain Mackie, and Jeff Dalton. 2024. Dreg: Document re-ranking using entity-based query understanding. *arXiv preprint arXiv:2401.05939*.
- Chen Chen, Bowen Zhang, Liangliang Cao, Jiguang Shen, Tom Gunter, Albin Madappally Jose, Alexander Toshev, Jonathon Shlens, Ruoming Pang, and Yinfei Yang. 2023. Stair: Learning sparse text and image representation in grounded tokens. *arXiv preprint arXiv:2301.13081*.
- Marek Ciglan, Kjetil Nørvåg, and Ladislav Hluchý. 2012. [The semsets model for ad-hoc semantic list search](#). In *Proceedings of the 21st International Conference on World Wide Web, WWW '12*, page 131–140, New York, NY, USA. Association for Computing Machinery.
- Jeffrey Dalton, Laura Dietz, and James Allan. 2014. Entity query feature expansion using knowledge base links. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval, SIGIR '14*, pages 365–374.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). In *North American Chapter of the Association for Computational Linguistics*.
- Laura Dietz. 2019. [ENT Rank: Retrieving Entities for Topical Information Needs through Entity-Neighbor-Text Relations](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '19*, page 215–224, New York, NY, USA. Association for Computing Machinery.
- Shuai Ding and Torsten Suel. 2011. [Faster top-k document retrieval using block-max indexes](#). In *Proceeding of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2011, Beijing, China, July 25-29, 2011*, pages 993–1002. ACM.

- Jeffrey M Dudek, Weize Kong, Cheng Li, Mingyang Zhang, and Michael Bendersky. 2023. Learning sparse lexical representations over specified vocabularies for retrieval. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, CIKM '23, pages 3865–3869.
- Faezeh Ensan and Ebrahim Bagheri. 2017. **Document retrieval model through semantic linking**. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining*, WSDM '17, page 181–190, New York, NY, USA. Association for Computing Machinery.
- Thibault Formal, Carlos Lassance, Benjamin Piwowarski, and Stéphane Clinchant. 2022. From distillation to hard negative sampling: Making sparse neural ir models more effective. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2353–2359.
- Thibault Formal, Benjamin Piwowarski, and Stéphane Clinchant. 2021. Splade: Sparse lexical and expansion model for first stage ranking. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2288–2292.
- Evgeniy Gabrilovich and Shaul Markovitch. 2009. Wikipedia-based semantic interpretation for natural language processing. *Journal of Artificial Intelligence Research*, 34:443–498.
- Darío Garigliotti and Krisztian Balog. 2017. **On type-aware entity retrieval**. In *Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval*, ICTIR '17, page 27–34, New York, NY, USA. Association for Computing Machinery.
- Emma J Gerritse, Faegheh Hasibi, and Arjen P de Vries. 2020. **Graph-Embedding Empowered Entity Retrieval**. In *Advances in Information Retrieval, Proceedings of the 42nd European Conference on Information Retrieval (ECIR 2020)*, Lecture Notes in Computer Science, pages 97–110, Cham. Springer.
- Emma J. Gerritse, Faegheh Hasibi, and Arjen P. de Vries. 2022. **Entity-aware transformers for entity search**. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '22, page 1455–1465, New York, NY, USA. Association for Computing Machinery.
- David Graus, Manos Tsagkias, Wouter Weerkamp, Edgar Meij, and Maarten de Rijke. 2016. **Dynamic collective entity representations for entity ranking**. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*, WSDM '16, page 595–604, New York, NY, USA. Association for Computing Machinery.
- Jiafeng Guo, Gu Xu, Xueqi Cheng, and Hang Li. 2009. Named entity recognition in query. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, SIGIR '09, pages 267–274.
- Nam Hai Le, Thomas Gerald, Thibault Formal, Jian-Yun Nie, Benjamin Piwowarski, and Laure Soulier. 2023. Cosplade: Contextualizing splade for conversational information retrieval. In *European Conference on Information Retrieval*, pages 537–552. Springer.
- Faegheh Hasibi, Krisztian Balog, and Svein Erik Bratsberg. 2016. **Exploiting entity linking in queries for entity retrieval**. In *Proceedings of the 2016 ACM International Conference on the Theory of Information Retrieval*, ICTIR '16, page 209–218, New York, NY, USA. Association for Computing Machinery.
- Vitor Jeronymo, Luiz Bonifacio, Hugo Abonizio, Marzieh Fadaee, Roberto Lotufo, Jakub Zavrel, and Rodrigo Nogueira. 2023. **InPars-v2: Large language models as efficient dataset generators for information retrieval**.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Rianne Kaptein, Pavel Serdyukov, Arjen De Vries, and Jaap Kamps. 2010. **Entity ranking using wikipedia as a pivot**. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*, CIKM '10, page 69–78, New York, NY, USA. Association for Computing Machinery.
- Ravi Kumar and Andrew Tomkins. 2010. A characterization of online browsing behavior. In *Proceedings of the 19th International Conference on World Wide Web*, pages 561–570.
- Md Tahmid Rahman Laskar, Cheng Chen, Aliaksandr Martsinovich, Jonathan Johnston, Xue-Yong Fu, Shashi Bhushan Tn, and Simon Corston-Oliver. 2022. **BLINK with Elasticsearch for efficient entity linking in business conversations**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track*, pages 344–352, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.
- Carlos Lassance and Stéphane Clinchant. 2022. An efficiency study for splade models. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2220–2226.
- Carlos Lassance, Hervé Déjean, Thibault Formal, and Stéphane Clinchant. 2024. Splade-v3: New baselines for splade. *arXiv preprint arXiv:2403.06789*.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef,

- Sören Auer, and Christian Bizer. 2015. [Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia](#). *Semantic Web*, 6:167–195.
- Canjia Li, Andrew Yates, Sean MacAvaney, Ben He, and Yingfei Sun. 2023. Parade: Passage representation aggregation for document reranking. *ACM Transactions on Information Systems*, 42(2):1–26.
- Jimmy Lin, Xueguang Ma, Sheng-Chieh Lin, Jheng-Hong Yang, Ronak Pradeep, and Rodrigo Nogueira. 2021. [Pyserini: A python toolkit for reproducible information retrieval research with sparse and dense representations](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '21, page 2356–2362, New York, NY, USA. Association for Computing Machinery.
- Jimmy Lin, Rodrigo Frassetto Nogueira, and Andrew Yates. 2020. [Pretrained transformers for text ranking: BERT and beyond](#). *CoRR*, abs/2010.06467.
- Xitong Liu and Hui Fang. 2015. Latent entity space: a novel retrieval approach for entity-bearing queries. *Information Retrieval Journal*, 18(6):473–503.
- Zhenghao Liu, Chenyan Xiong, Maosong Sun, and Zhiyuan Liu. 2018. [Entity-duet neural ranking: Understanding the role of knowledge graph semantics in neural information retrieval](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2395–2405, Melbourne, Australia. Association for Computational Linguistics.
- Sean MacAvaney, Craig Macdonald, and Iadh Ounis. 2022. Streamlining evaluation with ir-measures. In *European Conference on Information Retrieval*, pages 305–310. Springer.
- Sean MacAvaney, Franco Maria Nardini, Raffaele Perego, Nicola Tonellotto, Nazli Goharian, and Ophir Frieder. 2020. Expansion via prediction of importance with contextualization. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 1573–1576.
- Sean MacAvaney, Andrew Yates, Arman Cohan, and Nazli Goharian. 2019. Cedr: Contextualized embeddings for document ranking. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, pages 1101–1104.
- Iain Mackie, Shubham Chatterjee, and Jeffrey Dalton. 2023. Generative relevance feedback with large language models. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2026–2031.
- Iain Mackie, Shubham Chatterjee, Sean MacAvaney, and Jeff Dalton. 2024. Adaptive latent entity expansion for document retrieval. *The First Workshop on Knowledge-Enhanced Information Retrieval (ECIR'24)*.
- Iain Mackie, Paul Owoicho, Carlos Gemmell, Sophie Fischer, Sean MacAvaney, and Jeffrey Dalton. 2022. [Codec: Complex document and entity collection](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '22, page 3067–3077, New York, NY, USA. Association for Computing Machinery.
- Edgar Meij, Dolf Trieschnigg, Maarten de Rijke, and Wessel Kraaij. 2010. [Conceptual language models for domain-specific retrieval](#). *Inf. Process. Manage.*, 46(4):448–469.
- Donald Metzler and W. Bruce Croft. 2005. [A markov random field model for term dependencies](#). In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '05, page 472–479, New York, NY, USA. Association for Computing Machinery.
- Thong Nguyen, Mariya Hendriksen, Andrew Yates, and Maarten De Rijke. 2024. Multimodal learned sparse retrieval with probabilistic expansion control. In *Advances in Information Retrieval: 46th European Conference on Information Retrieval, ECIR 2024, Glasgow, UK*. Springer.
- Thong Nguyen, Sean MacAvaney, and Andrew Yates. 2023a. Adapting learned sparse retrieval for long documents. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1781–1785.
- Thong Nguyen, Sean MacAvaney, and Andrew Yates. 2023b. A unified framework for learned sparse retrieval. In *European Conference on Information Retrieval*, pages 101–116.
- Jianmo Ni, Gustavo Hernandez Abrego, Noah Constant, Ji Ma, Keith Hall, Daniel Cer, and Yinfei Yang. 2022a. [Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1864–1874, Dublin, Ireland. Association for Computational Linguistics.
- Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernandez Abrego, Ji Ma, Vincent Zhao, Yi Luan, Keith Hall, Ming-Wei Chang, and Yinfei Yang. 2022b. [Large dual encoders are generalizable retrievers](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9844–9855, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Fedor Nikolaev and Alexander Kotov. 2020. [Joint word and entity embeddings for entity retrieval from a](#)

- knowledge graph. In *Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part I*, page 141–155, Berlin, Heidelberg. Springer-Verlag.
- Fedor Nikolaev, Alexander Kotov, and Nikita Zhiltsov. 2016. [Parameterized fielded term dependence models for ad-hoc entity retrieval from knowledge graph](#). In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16*, page 435–444, New York, NY, USA. Association for Computing Machinery.
- Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. 2020. [Document ranking with a pre-trained sequence-to-sequence model](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 708–718, Online. Association for Computational Linguistics.
- Rodrigo Frassetto Nogueira and Kyunghyun Cho. 2019. [Passage re-ranking with BERT](#). *CoRR*, abs/1901.04085.
- Pooja Oza and Laura Dietz. 2023. Entity embeddings for entity ranking: A replicability study. In *European Conference on Information Retrieval*, pages 117–131. Springer.
- 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. 2021. [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. Association for Computational Linguistics.
- Amir Pouran Ben Veyseh, Franck Dernoncourt, and Thien Huu Nguyen. 2021. [DPR at SemEval-2021 task 8: Dynamic path reasoning for measurement relation extraction](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 397–403, Online. Association for Computational Linguistics.
- Hadas Raviv, David Carmel, and Oren Kurland. 2012. [A ranking framework for entity oriented search using markov random fields](#). In *Proceedings of the 1st Joint International Workshop on Entity-Oriented and Semantic Search, JIWES '12*, New York, NY, USA. Association for Computing Machinery.
- Hadas Raviv, Oren Kurland, and David Carmel. 2016. [Document retrieval using entity-based language models](#). In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16*, page 65–74, New York, NY, USA. Association for Computing Machinery.
- Nils Reimers and Iryna Gurevych. 2019a. [Sentencebert: Sentence embeddings using siamese bert-networks](#). *CoRR*, abs/1908.10084.
- Nils Reimers and Iryna Gurevych. 2019b. [Sentencebert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Stephen Robertson and Hugo Zaragoza. 2009. *The Probabilistic Relevance Framework: BM25 and Beyond*. Now Publishers Inc.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Michael Schuhmacher, Laura Dietz, and Simone Paolo Ponzetto. 2015. [Ranking Entities for Web Queries Through Text and Knowledge](#). In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management, CIKM '15*, page 1461–1470, New York, NY, USA. Association for Computing Machinery.
- Dahlia Shehata, Negar Arabzadeh, and Charles LA Clarke. 2022. Early stage sparse retrieval with entity linking. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 4464–4469.
- Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin, and Zhaochun Ren. 2023. [Is ChatGPT good at search? investigating large language models as re-ranking agents](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14918–14937, Singapore. Association for Computational Linguistics.
- Jiajie Tan, Jinlong Hu, and Shoubin Dong. 2023. Incorporating entity-level knowledge in pretrained language model for biomedical dense retrieval. *Computers in Biology and Medicine*, 166:107535.
- Alberto Tonon, Gianluca Demartini, and Philippe Cudré-Mauroux. 2012. [Combining inverted indices and structured search for ad-hoc object retrieval](#). In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '12*, page 125–134, New York, NY, USA. Association for Computing Machinery.
- Hai Dang Tran and Andrew Yates. 2022. Dense retrieval with entity views. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 1955–1964.
- Johannes M. van Hulst, Faegheh Hasibi, Koen Dercksen, Krisztian Balog, and Arjen P. de Vries. 2020. [Rel: An entity linker standing on the shoulders of giants](#). In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20*, page 2197–2200,

- New York, NY, USA. Association for Computing Machinery.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Ellen M Voorhees et al. 2003. Overview of the trec 2003 robust retrieval track. In *Trec*, pages 69–77.
- Chenyan Xiong and Jamie Callan. 2015. **Esdrank: Connecting query and documents through external semi-structured data**. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management, CIKM '15*, pages 951–960, New York, NY, USA. ACM.
- Chenyan Xiong, Jamie Callan, and Tie-Yan Liu. 2017a. **Word-entity duet representations for document ranking**. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '17*, page 763–772, New York, NY, USA. Association for Computing Machinery.
- Chenyan Xiong, Zhengzhong Liu, Jamie Callan, and Eduard Hovy. 2017b. **Jointsem: Combining query entity linking and entity based document ranking**. In *Proceedings of the 2017 ACM SIGIR Conference on Information and Knowledge Management, CIKM '17*, page 2391–2394, New York, NY, USA. Association for Computing Machinery.
- Chenyan Xiong, Zhengzhong Liu, Jamie Callan, and Tie-Yan Liu. 2018. **Towards better text understanding and retrieval through kernel entity salience modeling**. In *The 41st International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '18*, page 575–584, New York, NY, USA. Association for Computing Machinery.
- Chenyan Xiong, Russell Power, and Jamie Callan. 2017c. **Explicit semantic ranking for academic search via knowledge graph embedding**. In *Proceedings of the 26th International Conference on World Wide Web, WWW '17*, page 1271–1279, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- Ikuya Yamada, Akari Asai, Jin Sakuma, Hiroyuki Shindo, Hideaki Takeda, Yoshiyasu Takefuji, and Yuji Matsumoto. 2020. **Wikipedia2Vec: An efficient toolkit for learning and visualizing the embeddings of words and entities from Wikipedia**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 23–30, Online. Association for Computational Linguistics.
- Hamed Zamani, Mostafa Dehghani, W Bruce Croft, Erik Learned-Miller, and Jaap Kamps. 2018. From neural re-ranking to neural ranking: Learning a sparse representation for inverted indexing. In *Proceedings of the 27th ACM international conference on information and knowledge management*, pages 497–506.
- Pu Zhao, Can Xu, Xiubo Geng, Tao Shen, Chongyang Tao, Jing Ma, Daxin Jiang, et al. 2023. **Lexlip: Lexicon-bottlenecked language-image pre-training for large-scale image-text retrieval**. *arXiv preprint arXiv:2302.02908*.
- Tiancheng Zhao, Xiaopeng Lu, and Kyusong Lee. 2020. **Sparta: Efficient open-domain question answering via sparse transformer matching retrieval**. *arXiv preprint arXiv:2009.13013*.
- Nikita Zhiltsov, Alexander Kotov, and Fedor Nikolaev. 2015. **Fielded sequential dependence model for ad-hoc entity retrieval in the web of data**. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '15*, page 253–262, New York, NY, USA. Association for Computing Machinery.
- Shengyao Zhuang and Guido Zuccon. 2021. **Tilde: Term independent likelihood model for passage re-ranking**. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1483–1492.

A Appendix

A.1 Detailed training configuration

We train our DyVo methods using two-step distillation. In the first step, we train a base LSR model on MSMARCO without entities using standard LSR training techniques. We employ KL loss to distill knowledge from a cross-encoder with data obtained from *sentence-transformers* (Reimers and Gurevych, 2019b)². This model is trained with a batch size of 64 triplets (query, positive passage, negative passage) for 300k steps with 16-bit precision. In the second step, we start from the model pretrained on MSMARCO and further fine-tune it on the target datasets using distillation training on synthesized queries, BM25 negatives, and cross-encoder scores from MonoT5-3B (Nogueira et al., 2020). The documents in Robust04, Core 2018, and CODEC are longer than MSMARCO, so we use a smaller batch size of 16. All models are trained on one single A100 for for 100k steps.

For query generation, we re-use generated queries from InParsv2 (Jeronymo et al., 2023) available for TREC Robust04 and Core 2018. For CODEC, we generate the queries ourselves using the prompting Mixtral model. We re-use the prompt template in InParsv2 and add a small instruction at the beginning (Prompt A.1).

²<https://huggingface.co/datasets/sentence-transformers/msmarco-hard-negatives>

For sparse regularization, we apply L1 with varying regularization strengths. Entity representations are sparse themselves since we constrain the output to a small set of entity candidates and ignore other entities. Therefore, we do not apply L1 to entities.




Figure 3: Entity representation collapse during training.

A.2 Entity representation collapse

When integrating entity embeddings into DyVo, we observe that the model produces entity weights with magnitudes significantly higher than those of word piece weights. This discrepancy may arise from the lack of alignment between entity embeddings, generated by a separate model, and word piece embeddings. Initially, the model attempts to mitigate the dominance of entity weights by scaling them down. However, after a certain number of training epochs, the model overcompensates, resulting in the collapse of entity representations. This collapse is illustrated in Figure 3, where all entity weights become negative and are subsequently filtered out by the ReLU gate. Once this collapse occurs, it cannot be rectified, as there is no gradient flowing through the ReLU gate. To address this issue, we introduce a learnable scaling factor, as depicted in Equation 4, initializing it to a small value. This scaling factor is helpful to alleviate entity dominance at the beginning of training and temper the model’s aggressiveness in scaling down entity weights during training.

A.3 Qualitative comparison of different entity retrieval systems

In Table 4, we provide a qualitative comparison of the entity candidates retrieved by different systems. Within the two query samples presented, we observe that the generative approaches (i.e., Mixtral

and GPT4) consistently produce highly relevant entities. Notably, Mixtral tends to generate fewer and shorter entities compared to both GPT-4 and human annotations. Conversely, GPT4 retrieves more entities, and sometimes more entities than human-produced candidates. This discrepancy suggests an explanation for why Mixtral’s performance in generating entities to support document retrieval falls short of that achieved by GPT4.

In contrast to the consistent performance of generative entity retrieval, we observe divergent behaviors among other approaches (i.e., REL, BM25, and LaQue) across the two queries. The first query, which is less ambiguous with clearly expressed entities, allows these systems to retrieve/link simple, direct entities such as “American Revolutionary War” and “France in the American Revolutionary War”. However, they also introduce a significant amount of noise with irrelevant entities.

Conversely, the second query poses greater difficulty, with the entity “Non-fungible token” mentioned via its abbreviation “NFTs” which is further fragmented by the DistilBERT tokenizer into meaningless sub-word units. In this scenario, REL and BM25 fail entirely, while LaQue manages to retrieve only generic and distantly relevant entities. None of these systems successfully resolves “NFTs” to “Non-fungible token” as the generative approach does.

Retriever	Q: <i>“How vital was French support during the American Revolutionary War?”</i> WP : [how, vital, was, french, support, during, the, american, revolutionary, war, ?]
REL	[Vitalism, French people, Military logistics, American Revolutionary War]
BM25	[Richard Howe, 1st Earl Howe, HMS Childers (1778), Robert Howe (Continental Army officer), James Countts Crawford, Glorious First of June, George Eyre, Jacques-Antoine de Chambarlhac de Laubespain, Anthony James Pye Molloy, Nantucket during the American Revolutionary War era, Friedrich Joseph, Count of Nauendorf, Jonathan Faulknor the elder, Joseph Spear, HMS Romney (1762), HMS Roebuck (1774), France in the American Revolutionary War, Invasion of Corsica (1794), List of British fencible regiments, Northern theater of the American Revolutionary War after Saratoga, Robert Linzee, Guilin Laurent Bizanet]
LaQue	[France in the American Revolutionary War, List of French units in the American Revolutionary War, Support our troops, List of wars involving France, List of American Revolutionary War battles, American Volunteers, Colonial American military history, List of battles involving France in modern history, Military history of France, List of the lengths of United States participation in wars, 1776, France and the American Civil War, USS Confederacy (1778), Financial costs of the American Revolutionary War, List of wars involving the United States, List of American Civil War generals (Union), United States assistance to Vietnam, French Revolutionary Wars, American Revolutionary War, List of battles involving France]
Mixtral	[American Revolutionary War, France, United States, Military history, Diplomacy, Military alliance]
GPT4	[France in the American Revolutionary War, French Army, American Revolutionary War, Benjamin Franklin, Kingdom of France, Treaty of Alliance (1778), George Washington, John Adams, Treaty of Paris (1783), Continental Congress, Continental Army, Naval battles of the American Revolutionary War, Siege of Savannah, Capture of Fort Ticond]
Human	[American Revolution, France in the American Revolutionary War, Kingdom of Great Britain, United States, George Washington, Roderigue Hortalez and Company, British Empire, France, George Washington in the American Revolution, Gilbert du Motier, Marquis de Lafayette, Spain and the American Revolutionary War, American Revolutionary War, Diplomacy in the American Revolutionary War, Treaty of Paris (1783), Franco-American alliance, Naval battles of the American Revolutionary War, Treaty of Alliance (1778), Battles of Saratoga]
	Q: <i>Why are many commentators arguing NFTs are the next big investment category?</i> WP: [why, are, many, commentators, arguing, n, ##ft, ##s, are, the, next, big, investment, category]
REL	[Sports commentator, National Film and Television School, Next plc, Toronto, Investment banking, Categorization]
BM25	[Kuznets swing, The Green Bubble, Why We Get Fat, Big mama, Types of nationalism, Not for Tourists, Mark Roeder, Ernie Awards, Dramatistic pentad, Pagan Theology, RJ Balaji, Leslie Hardcastle, Why didn't you invest in Eastern Poland?, Big Data Maturity Model, Celebrity Big Brother racism controversy, The Bottom Billion, National Film and Television School, Canopy Group, The Wallypug of Why]
LaQue	[List of bond market indices, National Futures Association, NB Global, Companies listed on the New York Stock Exchange (N), Companies listed on the New York Stock Exchange (G), Companies listed on the New York Stock Exchange (F), List of exchange-traded funds, Companies listed on the New York Stock Exchange (T), Emerging and growth-leading economies, List of private equity firms, List of wealthiest organizations, Pension investment in private equity, Group of Ten (economics), Companies listed on the New York Stock Exchange (P), List of stock market indices, Lists of corporate assets, Companies listed on the New York Stock Exchange (U), List of public corporations by market capitalization, Net capital outflow, National best bid and offer]
Mixtral	[Non-fungible token, Blockchain, Cryptocurrency, Digital art, Ethereum, Value proposition, Art market, CryptoKitties, Investment strategy]
GPT4	[Non-fungible token, Cryptocurrency, Bitcoin, Ethereum, Digital art, Blockchain, CryptoKitties, Digital asset, Cryptocurrency bubble, Cryptocurrency exchange, Initial coin offering, Cryptocurrency wallet, Smart contract, Decentralized application, Digital currency]
Human	[Cryptocurrency, Public key certificate, Blockchain, Virtual economy, Bitcoin, Speculation, Non-fungible token, Ethereum]

Table 4: Example of relevant entities retrieved by different systems. List of word pieces (**WP**) returned by DistilBERT tokenizer is shown under each query.

Prompt. A.1: Prompt template for query generation with LLMs

Given an input document, your task is to generate a short and self-contained question that could be answered by the document. Three examples are given, please finish generating the query for the last example. Please generate only one short and self-contained question without numbering in a single line, and do not generate an explanation.

Example 1:

Document: We don't know a lot about the effects of caffeine during pregnancy on you and your baby. So it's best to limit the amount you get each day. If you are pregnant, limit caffeine to 200 milligrams each day. This is about the amount in 1½ 8-ounce cups of coffee or one 12-ounce cup of coffee.

Relevant Query: Is a little caffeine ok during pregnancy?

Example 2:

Document: Passiflora herbertiana. A rare passion fruit native to Australia. Fruits are green-skinned, white fleshed, with an unknown edible rating. Some sources list the fruit as edible, sweet and tasty, while others list the fruits as being bitter and inedible. Passiflora herbertiana. A rare passion fruit native to Australia. Fruits are green-skinned, white fleshed, with an unknown edible rating. Some sources list the fruit as edible, sweet and tasty, while others list the fruits as being bitter and inedible.

Relevant Query: What fruit is native to Australia?

Example 3:

Document: The Canadian Armed Forces. 1 The first large-scale Canadian peacekeeping mission started in Egypt on November 24, 1956. 2 There are approximately 65,000 Regular Force and 25,000 reservist members in the Canadian military. 3 In Canada, August 9 is designated as National Peacekeepers' Day.

Relevant Query: How large is the Canadian military?

Example 4:

Document: {input document}

Relevant Query:;

Prompt. A.2: Prompt template for few-shot generative entity retrieval

Identify Wikipedia entities that are helpful to retrieve documents relevant to a web search query. Please return a list of entity names only:

Example 1:

Query: How is the push towards electric cars impacting the demand for raw materials?

Entities: ["Cobalt", "Automotive battery", "China", "Electric car", "Electric battery", "Gigafactory 1", "Demand", "Fossil fuel", "Electric vehicle industry in China", "Electric vehicle battery", "Electric vehicle conversion", "Electric vehicle", "Supply and demand", "Mining industry of the Democratic Republic of the Congo", "Raw material", "Lithium iron phosphate", "Lithium-ion battery", "Mining", "Lithium", "Petroleum"]

Example 2:

Query: Why do many economists argue against fixed exchange rates?

Entities: ["Argentine peso", "Currency crisis", "Inflation", "Hong Kong dollar", "Exchange rate", "Gold standard", "European Exchange Rate Mechanism", "1998 Russian financial crisis", "Black Saturday (1983)", "Black Wednesday", "Optimum currency area", "Mexican peso crisis", "Milton Friedman", "Euro", "Recession", "Currency intervention", "1997 Asian financial crisis", "Devaluation", "Original sin (economics)", "Exchange-rate regime"]

Please find relevant entities for this new example:

Query: {input query}

Entities: