Leveraging Decoder Architectures for Learned Sparse Retrieval

Jingfen Qiao[1][*], Thong Nguyen[1], Evangelos Kanoulas[1], and Andrew Yates[1,2]

[1] University of Amsterdam, Netherlands
{j.qiao,t.nguyen2,e.kanoulas}@uva.nl
[2] Johns Hopkins University, United States
andrew.yates@jhu.edu

Abstract. Learned Sparse Retrieval (LSR) has traditionally focused on small-scale encoder-only transformer architectures. With the advent of large-scale pre-trained language models, their capability to generate sparse representations for retrieval tasks across different transformer-based architectures, including encoder-only, decoder-only, and encoder-decoder models, remains largely unexplored. This study investigates the effectiveness of LSR across these architectures, exploring various sparse representation heads and model scales. Our results highlight the limitations of using large language models to create effective sparse representations in zero-shot settings, identifying challenges such as inappropriate term expansions and reduced performance due to the lack of expansion. We find that the encoder-decoder architecture with multi-tokens decoding approach achieves the best performance among the three backbones. While the decoder-only model performs worse than the encoder-only model, it demonstrates the potential to outperform when scaled to a high number of parameters.

Keywords: Learned Sparse Retrieval · LLMs · Information Retrieval.

1 Introduction

Modern LLMs are knowledgeable about a wide range of topics and have achieved remarkable performance in their application to information retrieval [17,28,30]. However, much of the research has focused on using LLMs to generate dense vectors rather than sparse ones [17,18]. Learned sparse retrieval (LSR) utilizes LLMs to encode queries and documents into lexical sparse vectors, i.e. vectors where most of the elements are zero. Compared to dense retrieval, LSR's lexical representations are more interpretable as each output dimension is aligned with a term in a vocabulary. In addition, LSR relies on an inverted index for retrieval, which is smaller in size than a vector index (e.g., HNSW [21]) employed by dense retrieval. Recent studies have demonstrated that LSRs could efficiently achieve strong first-stage retrieval performance [15,4]. Its effectiveness, efficiency and transparency make LSR a compelling alternative to dense retrieval in many information retrieval applications.

---

[*] Corresponding author

Query: what is paula deen's brother (qid: 1048585)

Tokens predicted by fine-tuned FlanT5-base

MLM-MultiTokens
De, de, Dean, Leader, actor, a, son, was, Pa, ul, name, character, married, is, la, relationship, who mother, his, origin, member, boy, boyfriend, daughter, partner

Overlap
Paul, brother, pa, Brother, Dean, brothers, her, family, a, sister, father, friend, child

MLM-SingleToken
De, de, pa, P, cousin, her, D, Pe, Pa, de, first, baby, Patrick, da, Le

Tokens missed by zero-shot FlanT5-xl

sister, brothers, Pe, is, Brother, D, Paul, boy, child, son, Patrick, leader, partner, boyfriend, relationship, father, her, mother, de, la, Pa, De, cousin, actor, baby, Le, his, family, pa, first, ul, name, origin, da, friend, daughter, married, was, member, P, who, character, Dean
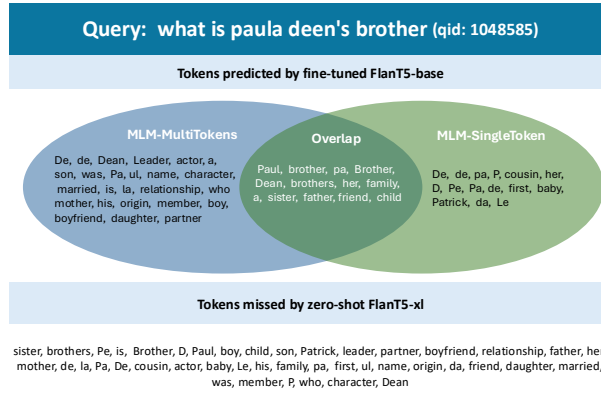
Fig. 1: Output Bags of Tokens Produced by Different Sparse Representation Heads; Zero-shot encoder (FlanT5-xl) misses many important expansion terms. The MLM-MultiTokens head captures more relevant tokens than the MLM-SingleToken head by gathering contextual information from all input token rather than a single token.

When evaluating the capabilities of LLMs for sparse retrieval, a key question is whether they can generate effective sparse representations in zero-shot settings. Our analysis shows that while LLMs excel in many generative tasks, they struggle with zero-shot sparse retrieval. As illustrated in Figure 1, which shows the tokens predicted by different sparse representation heads with a zero shot or fine-tuned model, LLMs not only fail to expand the input text to include semantically relevant terms, but they often introduce noise, leading to semantic drift and poor retrieval performance. This limitation highlights the need for supervised training to enhance LLMs' effectiveness and raises the question of whether fine-tuning could enable LLMs to outperform smaller-scale learned sparse retrieval (LSR) methods.

Previous research on LSR has predominantly focused on small-scale encoder-only transformer architectures [2,12,19,27,34]. For instance, Splade [12], a state-of-the-art LSR model, is built upon Cocondenser [13], which utilizes a BERT-base architecture with 110M parameters, specifically pre-trained for Information Retrieval (IR) tasks. The generalizability of LSR to other transformer architectures, such as decoder-only and encoder-decoder, remains unclear, as does its scalability concerning model size. Decoders are trained to sequentially generate a subsequent token based on the given input text and previously generated tokens. This is different from sparse representation learning which requires weighting existing tokens within the input and expanding them with contextually relevant terms. To learn a sparse representation from the decoder output, we investigate using a multi-tokens decoding approach. This adaptation allows the decoder to work more like an encoder, making better use of the parameters within the decoder module. Our results indicate that the multi-tokens decoding approach for encoder-decoder
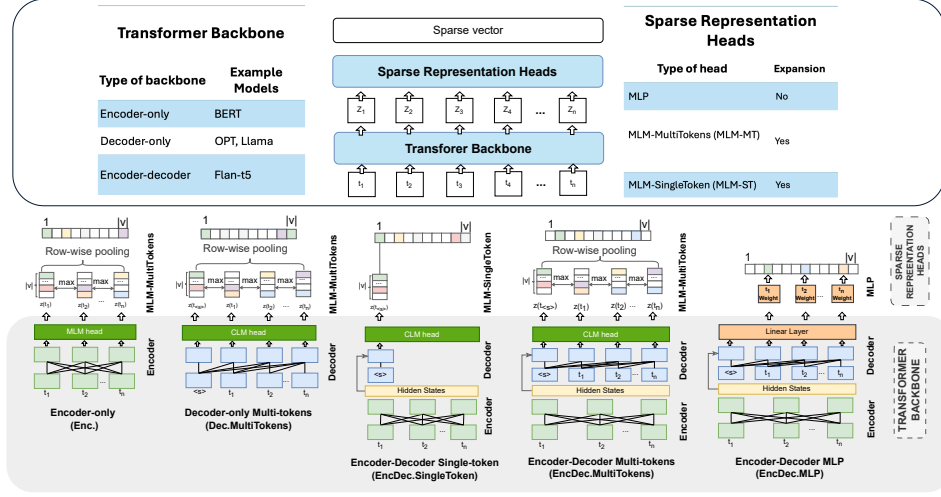
Fig. 2: Learned sparse retrieval architectures consist of (1) a transformer backbone that takes query or document text as input and outputs hidden state(s) and (2) a sparse representation head that takes the hidden state(s) as input and outputs sparse lexical representations.

models helps to effectively utilize the parameters in both encoder and decoder modules.

This paper aims to address these gaps by investigating LSR across different transformer architectures (encoder-only, decoder-only, and encoder-decoder) at several scales and exploring the capabilities of LLMs under zero-shot and fined-tuning settings. Our findings are as follows: **(1)** Without fine-tuning, LLMs struggle to generate meaningful sparse representations due to inappropriate term expansions; **(2)** the ability of decoder-only architectures to produce sparse representations is limited and comparable results to encoder-only models are only achieved when the model is scaled to a large size, such as 3 billion parameters; **(3)** by copying the encoder's input to the decoder we enable the encoder-decoder to perform multi-tokens decoding effectively aggregating lexical information across the input sequence and creating a significantly more effective sparse representation compared to the single-token approach; and **(4)** encoder-decoder architectures with multi-tokens decoding can outperform encoder-only and decoder-only.[3]

2   LSR Architectures

In LSR, queries and passages are encoded into sparse representations over terms (Figure 2). In this section, we describe how to generate sparse representations from

_____

[3] Code: https://github.com/JingfenQiao/Decoder4LSR

various transformer-based architectures, including encoder-only, decoder-only, and encoder-decoder.

2.1 Transformer Backbones

Given a sequence of n tokens $(t_1, t_2, \ldots, t_n)$ in the query or passage, the task of a LSR model is to predict term weights based on the hidden states output by the transformer-based backbone. Therefore, the transformer backbone aims to transform each input token t_i into a hidden state (latent vector) h_i. A sparse representation of the input is then produced by passing these hidden states to a sparse representation head. In this section, we first describe how the hidden states h_i are produced by encoder-only backbone in the previous research, and then illustrate how to leverage decoders to construct sparse representations.

We will use the following notation to describe each transformer-based backbone in this paper: $\oplus$ denotes the concatenation operator, $h_{1:n}$ denotes the hidden states corresponding to input tokens $t_{1:n}$ that are produced by a transformer-backbone, v_i and e_i represents the i_{th} token and transformer's initial input embedding in the model vocabulary $|V|$, where $i \in \{0, \ldots, |V| - 1\}$.

Encoder-only (Enc.): When using a transformer encoder stack as the backbone, we feed text directly into the encoder, which produces a contextualized embedding of each input term. We refer to these as hidden states h_i for consistency with the decoder-based backbone that we introduce in this paper.

$$h_{1:n} = \text{Encoder}(t_1, t_2, \ldots, t_n) \tag{1}$$

When using an encoder backbone, the hidden states are produced using bi-directional attention that considers terms both before and after a given term t_i when producing hidden states to represent the input sequence. This is an intuitive property for the transformer backbone to have. With the exception of the zero-shot PromptRep model [37] and orthogonal document expansion approaches like doc2query [26], all existing learned sparse retrieval models use backbones based on a transformer encoder [23].

Decoder-only with multi-tokens (Dec. MultiTokens): In contrast to encoder models, decoders are designed with causal attention, where the generation of a hidden state h_i depends only on the previously generated tokens, and the decoder starts with the start symbol of a model <s>. We concatenate the start symbol of the model <s> in front of the token t_1.

$$h_{1:n} = \text{Decoder}(\text{<s>}, t_1, t_2, \ldots, t_n) \tag{2}$$

When using decoder-based backbones, rather than invoking the decoder repeatedly for each token, we feed the entire input text to the decoder in a single pass to generate sparse representations. This contrasts with the typical token-by-token decoding approach, where a new hidden state is produced at each decoding step. As illustrated in Figure ??, the row-wise pooling of MLM

outputs demonstrates how the hidden state of every token contributes to the final multi-token representation.

Encoder-Decoder with single-token (EncDec. SingleToken): In encoder-decoder with single-token, similar to the Sentence-T5 embedding model [24], input text is fed into the encoder and only the start token `<s>` is fed into the decoder. Therefore, it only yields a single latent vector h_1 that is derived from the hidden state of the input token `<s>`, which is in turn derived from the hidden states of the entire input text produced by the encoder stack.

$$h_{1:n}^{enc} = \text{Encoder}(t_1, t_2, \ldots, t_n) \tag{3}$$

$$h_1 = \text{Decoder}(h_{1:n}^{enc} \oplus e_{<s>})_{n+1} \tag{4}$$

Encoder-Decoder with multi-tokens (EncDec. MultiTokens): We hypothesize that the output representation of the encoder-decoder architecture is bottlenecked by the single decoder input, the start token `<s>`, which limits the model's ability to aggregate lexical information from the entire input sequence. To overcome this bottleneck, we propose a multi-tokens decoding approach that allows the model to gather information from any input token. As shown in Figure **??**, the Encoder-Decoder with multi-tokens model receives full input text both in the encoder and decoder, and the decoder's input is prepended with the start token.

$$h_{1:n}^{enc} = \text{Encoder}(t_1, t_2, \ldots, t_n) \tag{5}$$

$$h_{1:n} = \text{Decoder}(h_{1:n}^{enc} \oplus e_{<s>} \oplus t_1, t_2, \ldots, t_n) \tag{6}$$

2.2   Sparse Representation Heads

A sparse representation head takes a hidden state produced by the transformer backbone as input and produces one or more term weights.[4] These term weights are then aggregated to produce a sparse representation of the entire input text, which can then be used for retrieval by taking the dot product of a query representation and a document representation. More formally, a sparse representation head takes a sequence of of n hidden states output by a transformer backbone $h_{1:n} = (h_1, h_2, \ldots, h_n)$ as input. The head outputs a sparse vector with $|V|$ dimensions where v_i corresponds to a term t_i in the vocabulary.

MLP head: The MLP head, used in models like DeepCT [14], uniCOIL [16] and DeepImpact [22], uses a multi-layer perceptron to process contextualized embeddings for each input term in the query or passage. It assigns non-zero weights to terms that occur only in the input text. The importance $MLP(h)_i$ of each term v_i within a model vocabulary V for a given input t_n is defined by Equation 7, where W and b are the weight and bias of the linear head, respectively.

---

[4] Nguyen et al. [23] refer to these as *sparse encoders*. We use the term *head* instead to avoid overloading the term *encoder*, which can also refer to the type of transformer backbone used.

$1(v_i = t_n)$ ensures to that only tokens in the input receive weights, the other terms are zero.

$$MLP(h_{1:n})_i = \sum_{j=1}^{n} \log\left(1(v_i = t_j) \cdot \text{ReLU}(h_j W + b) + 1\right) \qquad (7)$$

MLM head: In contrast to the MLP head, the MLM head expands the input (using the model's vocabulary) in addition to weighting tokens. MLM heads have been used in models such as SPLADE [12], Sparta [36], and TILDE [38]. The MLM head uses the logits value produced by the language model head for each input term v_i to produce a sparse representation. It can assign a weight to to any terms in the model vocabulary V regardless of whether it appears in the input text, thereby exploiting the language model's capability to expand and weight terms. The importance $MLM(h_{1:n})_i$ of each term v_i within a model's vocabulary V is aggregated from the logits matrix along the input sequence, as defined in Equation 8. The MLM-SingleToken (MLM-ST) process a single hidden state h_1 from the transformer backbone to determine the importance of each term in the vocabulary, while MLM-MultiTokens (MLM-MT) processes multiple hidden states $h_{1:n}$ and takes the max weight over the different hidden states.

$$MLM(h_{1:n})_i = \max_{j=1}^{n} \log\left(1 + \text{ReLu}\left(h_j^T e_i + b_i\right)\right) \qquad (8)$$

The *ReLu* activation function and logarithmic normalization are applied to ensure the positive of term weights [12] and prevent some terms from dominating in the sparse representation [10]. The max-pooling operator in Equation 8 is to aggregate term importance weights into a document-level sparse representation. Formal et al., [12] found that max pooling is more effective than sum pooling. Therefore, we employed max-pooling for all MLM models.

3   Experimental Setup

In this section, we outline the experimental design, including the models, training configurations, and evaluation methods, to investigate learned sparse retrieval (LSR) performance across various transformer architectures.

Retrieval and Indexing: We apply the same architecture type with shared weights for both the query and document encoders. For indexing and retrieval, we use the Anserini toolkit [32,33] to index the encoded documents and retrieve with the encoded queries.

Dataset and Evaluation: We trained our models on the MS MARCO passage dataset [3], which includes 8.8 million passages and 40 million training triplets. All models are trained on those triplets and hard negatives generated by the cross-encoder teacher model in [14]. We report the MRR@10, NDCG@10 and Recall@1000 of all models on the MS MARCO Dev and TREC Deep Learning (TREC-DL) 2019/2020 [6,7]. In addition of the effectiveness metrics, we report the FLOPs metric (the average number of term overlap between a pair of query and document) as a proxy of efficiency.

Backbones and baselines We selected three different pretrained transformer language model as the backbone of our study, including encoder-only (Distilbert-base-uncased [29]), decoder-only (OPT [35]) and encoder-decoder (T5 [5]). Distilbert-base-uncased is widely recognized and extensively employed within the LSR literature [9,12,15]. The OPT [35] and Flan-T5 [5] offers a wide range of parameter sizes. This variability enables us to explore how decoder-only architectures scale and perform in generating sparse representations, providing insights into their effectiveness across different model sizes. In terms of baselines, we select based on query expansion and term weighting strategies and compare against the unsupervised sparse model BM25 and the supervised sparse model, including EPIC [20], DeepCT [8], DocT5Query [26], DeepImpact [22], Splade-max [12] and DistilSplade-max [12]. The unsupervised sparse baseline BM25 employs neither expansion nor supervised weighting. Supervised LSR baselines include EPIC [20] and DeepCT [8], which use supervised weighting without query expansion; DocT5Query [26], which focuses solely on document term expansion without term weighting; DeepImpact [22], which applies supervised term weighting without query expansion; and Splade-max [12] and DistilSplade-max [12], which integrate both supervised query expansion and term weighting.

Training Configuration All models were fine-tuned using the ADAM optimizer with 6000 warmup steps. The batch size for all models was consistently maintained at 16. Due to model nuances, the OPT model uses a lower learning rate of 1^{e-6}, while Flan-T5 uses 5^{e-4}. We used one NVIDIA A6000 GPU with 48GB memory to train medium-sized models, such as Flan-T5-base, and OPT-3.5. Larger models (OPT-1.3B, OPT-2.7B, Flan-T5-large and Flan-T5-xl) were trained on four A6000 GPUs. Following the setting of Splade [12], we schedule the regularization weight λ to increase quadratically over the first 100k steps and we tried different values of λ ranging from 1^{-1} to 1^{-4}.

4   Optimization

Knowledge distillation has emerged as a highly effective strategy for transferring knowledge from a high-performing cross-encoder to a more efficient bi-encoder in information retrieval tasks [12,15,14]. By training the student model to replicate the teacher's score distribution, the student preserves a significant portion of the teacher's ranking ability. Compared to alternative approaches for retrieval, such as contrastive learning or reinforcement learning, the knowledge distillation inherits a well-calibrated signal from a strong teacher model, thus reducing the need for labor intensive training sample construction or careful reward engineering. Concretely, given a batch B of training triplets q_i, P_i^+, P_i^- consisting of a query q_i, a positive document P_i^+ and a negative document P_i^-, we compute the margin between the positive and negative documents for both teacher and student. We then minimize the mean squared error (MSE) of the difference between these margins, as shown in Equation 9.

$$\mathcal{L}_{\text{MarginMSE}} = \frac{1}{|B|} \sum_{i=\{0,\dots,|B|-1\}} \left(\Delta M_s^i - \Delta M_t^i \right)^2 \tag{9}$$

where $\Delta M_s^i = S(q_i, P_i^+) - S(q_i, P_i^-)$ and $\Delta M_t^i = T(q_i, P_i^+) - T(q_i, P_i^-)$ are the respective margin scores of the student S and teacher model T on the i-th triplet, and $|B|$ is the batch size. Following [14,12], we use MarginMSE as the distillation loss to ensure fair comparisons and reduce confounding factors. Although a more complex training setup (e.g. using an ensemble of multiple teachers, dual MarginMSE and KL-Div loss, continuous training as in Splade-V3 [15]) could be applied to increase the overall effectiveness, we believe that our relative findings between different architectures and approaches do not change.

Distillation Training Data To investigate how the scaling of teacher models affects performance, we experiment with two teachers: MiniLM-L-6-v2 [31] and Rankllama-13B [17]. Rankllama-13B outperforms MiniLM-L-6-v2 on the MSMARCO training set, as shown in Table 5 in Appendix. For both teacher models, we collect hard negatives from multiple retrieval systems (including dense retrieval models and BM25) to provide challenging negative samples. Each teacher then generates score pairs for the query-positive and query-hard-negative pairs using the hard negatives from Sentence-Transformers[5]. We use the smaller teacher's scores from [11] directly, whereas we generate the Rankllama-13B scores ourselves[6]. We observe that Rankllama-13B produces a relatively narrow score distribution (Appendix Table 3), which may make it more difficult for students to learn effectively. This aligns with findings by [15], who show that the MarginMSE-based distillation process is sensitive to score distributions. To address this, we apply a affine transformation to flatten Rankllama-13B's scores, ensuring that its mean and standard deviation more closely match those of MiniLM-L-6-v2.

Sparsity Regularization Finally, to encourage sparsity in the learned representations, we adopt FLOPs regularization [27] and activation functions [1] to facilitate the learning of distributed sparse embeddings. FLOPs are a differentiable relaxed approximation of the number of floating-point operations(FLOPs) to compute the score between query and document. Formally, our overall objective is defined in Equation 10.

$$\mathcal{L}_{\text{ranking}} = \mathcal{L}_{\text{MarginMSE}} + \lambda_q \mathcal{L}_{reg}^q + \lambda_d \mathcal{L}_{reg}^d \tag{10}$$

where $\mathcal{L}_{reg}^q$ and $\mathcal{L}_{reg}^d$ denote FLOPs-based regularization terms for query and document representations, respectively, and λ_q and λ_d are their weighting coefficients.

---

[5] https://huggingface.co/datasets/sentence-transformers/msmarco-hard-negatives
[6] https://huggingface.co/datasets/lsr42/rankllama-ms-marco-scores

Table 1: Comparison of sparse retrieval without fine-tuning on MS MARCO dev. $\times$ results are from [22]; A † indicates paired t-test $p < 0.05$. († indicates test between term expansion and non-expansion.

Model	Weighting	Expansion	RR@10	R@1k
Sparse retrieval				
(1) BM25$^{\times}$	Yes	No	18.8	85.8
(2) BM25+Doc2Query$^{\times}$	Yes	Yes	27.8	94.7
Learned sparse retrieval (zero-shot)				
(3) EncDec.MultiTokens$_{\text{FlanT5-xl}}$	Yes	Yes	1.30	16.5
(4) EncDec.MultiTokens$_{\text{FlanT5-xl}}$	Yes	No	3.67†	34.0†
(5) Dec.MultiTokens$_{\text{OPT-2.7B}}$	Yes	No	13.7	71.4
(6) Dec.MultiTokens$_{\text{OPT-6.7B}}$	Yes	No	12.5	70.5

5   Results and Discussion

RQ1: Can LLMs effectively generate sparse representations in a zero-shot setting when prompted?

Term weighting and expansion are key factors in the effectiveness of sparse retrieval systems [16]. We evaluate LLMs for sparse vector creation with and without term expansion. Recently, Zhuang et al. [37] proposed generating sparse and dense representations by prompting LLMs, using the logits of the last token in the input sequence. Their method only assigns weights to tokens present in the input, without term expansion. To investigate LLM's term expansion ability, we compare sparse vectors with and without expansion. In the expansion case, we keep the top 1,000 tokens, while the non-expansion approach uses only tokens from the input text. Instead of relying on the last token's logits, our method aggregates logits from all input tokens, following Equation 8.

Table 1 shows the zero-shot performance of LLMs versus traditional methods like BM25 and BM25+Doc2Query. All zero-shot models use the MLM-MultiTokens head for aggregating term weights. We find that enabling term expansion in Flan-T5-xl leads to a significant drop in Recall@1k (34.0 to 16.5), as noisy terms like 'ray,' 's,' and 'mil' receive high weights, harming performance. OPT-2.7B achieves the best results among LLMs (MRR 13.7, Recall@1k 71.4) but still falls short of BM25 and state-of-the-art LSR models, highlighting the need for fine-tuning to improve term importance assessment in LSR. Additionally, we assessed their effectiveness across various model sizes to provide a more comprehensive comparison, as shown in Table 6 in the Appendix.

RQ2: Can encoder-decoder or decoder-only backbones outperform encoder-only backbones when using multi-tokens decoding approach?

Table 2: Evaluation of different transformer backbones on MS MARCO passage (dev set) and TREC DL 2019/2020. *All models are trained with 600k steps and use the MLM-MultiTokens head for both query and document encoding. ∗ indicates reproduced results.*

Models	Backbones	FLOPs	DL2020 nDCG@10	DL2019 nDCG@10	MS MARCO dev nDCG@10	R@1k	RR@10
(1) Splade-max* [12]	Encoder-only	1.3	67.0	68.0	40.2	96.5	34.0
(2) DistilSplade-max* [12]	Encoder-only	4.0	67.9	71.0	43.3	97.9	36.8
(3) Enc.$_{\text{FlanT5-base}}$	Encoder-only	16.0	61.0	65.0	36.0	96.9	29.8
(4) Dec.MultiTokens$_{\text{FlanT5-base}}$	Decoder-only	2.3	67.4	68.7	39.9	97.1	33.6
(5) Dec.MultiTokens$_{\text{OPT-350m}}$	Decoder-only	4.9	68.6	66.4	39.8	97.2	33.7
(6) EncDec.MultiTokens$_{\text{FlanT5-base}}$	Encoder-Decoder	2.8	70.5	71.4	43.6	98.3	36.8

Encoder-only VS Encoder-decoder backbone The encoder-decoder backbone generally surpasses the encoder-only backbone, as shown in Table 2. Specifically, EncDec.MultiTokens (row 6) outperforms the encoder-only model DistilSplade-max (row 2) in NDCG@10 on the MS MARCO development set, with lower FLOPs (4 vs. 2.8) using the same training configuration. Upon inspection, we find that the encoder-decoder model is more effective at expanding the input queries and documents to relevant terms. The encoder-decoder architecture benefits from a hybrid attention mechanism, which includes bidirectional attention in the encoder and causal attention in the decoder. This allows it to identify different patterns of semantic dependency in the input. Interestingly, the encoder-only model using Flan-T5 (Enc.$_{\text{FlanT5-base}}$) performs significantly worse than the encoder-only model DistilSplade-max, especially on the MRR@10 and nDCG@10 metrics. This underperformance may be attributed to the fact that Flan-T5 was pre-trained with an encoder-decoder architecture, where the MLM head is only attached to the decoder's hidden states during training. When we attach the same MLM head to the encoder, there may be an incompatibility that makes LSR training more challenging.

Decoder-only backbone effectiveness We trained different LSR variants using the Flan-T5-base (248 million parameters) and OPT-350M (350 million parameters) decoder-only checkpoints. Among all the backbone variations (rows 3, 4, and 6) of Flan-T5-base, the decoder-only variant (row 4) was the second most effective, trailing only behind the encoder-decoder architecture (row 6). Furthermore, compared to DistilSplade-max (row 2), which is built on the BERT encoder, the decoder-only model using the OPT-350M checkpoint (row 5) shows a slight improvement in terms of NDCG@10 and MRR@10 metrics on the DL2020 dataset, but not on MS MARCO and DL2019. Additional insight could also be seen on Table 4, where we observe that the effectiveness of the Decoder-only model (row 8) can exceed Enc. (row 1) on DL2020, DL2019 and MS MARCO when the parameters are scaled to 1.3 billion.

Table 3: Performance of different Sparse Representation Heads on MS MARCO passage (dev) and TREC DL 2019/2020. *All models are trained with 600k steps. ∗ indicates reproduced results; × results are from the work in [22]; ×× results are from the work in [23]; A † indicates paired t-test $p < 0.05$. († indicates test between MLM-MT and MLM-ST); A ‡ indicates paired t-test $p < 0.01$(‡ indicates test between models where both query and passage use MLM-MT and models where either query or passage use MLM-MT)*

Models	Query	Passage	FLOPs	DL2020 nDCG@10	DL2019 nDCG@10	MS MARCO dev nDCG@10	R@1k	RR@10
(1) BM25×			-	48.7	49.7	23.5	85.8	18.8
(2) EPIC$^{××}_{top400}$ [20]	MLP	MLM-MT	-	71.8	70.9	-	97.2	37.2
(3) DeepCT×[8]	MLP	MLP	-	55.0	57.8	29.8	91.0	24.4
(4) DocT5Query×[26]	-	-	-	61.9	64.8	33.8	94.7	27.8
(5) DeepImpact×[22]	BINARY	expMLP	-	65.1	69.5	38.5	94.8	32.6
(6) Splade-max*[12]	MLM-MT	MLM-MT	1.3	67.0	68.0	40.2	96.5	34.0
(7) DistilSplade-max*[12]	MLM-MT	MLM-MT	4.0	67.9	71.0	43.3	97.9	36.8
(8) EncDec.MultiTokens$_{FlanT5\text{-}base}$	MLP	MLM-MT	1.3	53.6‡	56.6‡	41.7‡	97.6‡	35.3‡
(9) EncDec.MultiTokens$_{FlanT5\text{-}base}$	MLM-MT	MLP	3.3	50.6‡	48.1‡	31.3‡	95.9‡	25.8‡
(10) EncDec.SingleToken$_{FlanT5\text{-}base}$	MLM-ST	MLM-ST	2.5	54.5	59.7	36.1	95.0	30.4
(11) EncDec.MultiTokens$_{FlanT5\text{-}base}$	MLM-MT	MLM-MT	**2.8**	**70.5**†	**71.4**†	**43.6**†	**98.3**†	36.8†

RQ3: Which sparse representation head is better for creating a sparse representation?

MLM-MultiTokens (MLM-MT) vs. MLM-SingleToken (MLM-ST) We evaluate the MLM-MultiTokens and MLM-SingleToken heads on the best-performing encoder-decoder model using the Flan-T5 base checkpoint. As shown in Table 3, the MLM-MultiTokens head outperforms the MLM-SingleToken head across all evaluated metrics and datasets (MS MARCO, DL2019, DL2020). This suggests that the MLM-MultiTokens head, which aggregates latent vectors from multiple input tokens, facilitates more effective sparse representations compared to the single token approach from the decoder output.

MLM-MultiTokens vs. MLP We also compare the performance of the MLP and MLM-MultiTokens heads derived from the decoder output using the Flan-T5-base checkpoint. The results in Table 3 show that encoding both the query and document with the MLM-MultiTokens head achieves the highest effectiveness, significantly outperforming the MLP head, particularly in terms of nDCG@10 across different datasets. This contrasts with the results observed on an encoder-only backbone, where switching the query encoder from MLM-MultiTokens to MLP—reproduced using hard negative distillation labels in the study by [23]—did not yield a significant improvement. These findings suggest that MLM-MultiTokens head query expansion is more beneficial for encoder-

Table 4: Impact of scaling student and teacher model. *All models are trained with 600k steps and use the MLM-MultiTokens head for both querying and document encoding. A † ‡ indicates a paired t-test p < 0.05. († indicates a test between models tuned with the teacher MiniLM-L-6-v2 and Rankllama-13b. ‡ indicates between smaller student and larger student LSR model with same backbones)*

Models	DL2020 nDCG@10	DL2019 nDCG@10	MS MARCO dev RR@10	FLOPs
(1) Enc.$_{\text{Distilbert-base-uncased}}$	67.9	71.0	36.8	4.0
(2) EncDec.MultiTokens$_{\text{FlanT5-base}}$	70.5	71.4	36.8	2.8
(3) Dec.MultiTokens$_{\text{OPT-350M}}$	68.6	66.4	33.7	4.9
Scaling Teacher				
(4) Enc.$_{\text{Distilbert-base-uncased}}$	73.1†	75.1†	37.3	5.8
(5) EncDec.MultiTokens$_{\text{FlanT5-base}}$	70.1	73.4	37.8	5.2
(6) Dec.MultiTokens$_{\text{OPT-350M}}$	61	66.8	32.8	5.1
Scaling Student				
(7) EncDec.MultiTokens$_{\text{FlanT5-large}}$	70.7	73.3	37.8	3.8
(8) Dec.MultiTokens$_{\text{OPT-1.3B}}$	69.5	73.1‡	36.9‡	4.8

decoder backbones compared to encoder-only backbones.

RQ4: How is the performance of the LSR affected by scaling the teacher and student models on the different backbones?

Table 4 illustrates the impact of scaling teacher and student models across different backbones, including Encoder-only, Encoder-decoder, and Decoder-only.

Impact of scaling teacher model To assess the impact of scaling the teacher model on different backbones, we employed a more effective cross-encoder RankLLaMA-13B [17] as the teacher to supervise LSR model (student) tuning. Row 4-6 shows the effect of using a larger teacher model across different backbones. Consistent with findings by [15], the encoder-only model shows significant improvements when scaling the teacher model, increasing from 67.9 to 73.1 on DL2020 and from 71.0 to 75.1 on DL2019. The encoder-decoder backbone also exhibits improvements on DL2019 and MS MARCO, though to a lesser degree. In contrast, the decoder-only model experiences a decline in performance on MS MARCO and DL2020. This decline may be attributed to the sharp distribution of RankLlama teacher model scores, as shown in Figure 3. The sharp distribution makes it challenging for the decoder-only architecture to generate a sparse vector by relying solely on FLOPs regulation in the loss function. Further investigation is needed to fully understand this behavior. Overall, these results suggest that scaling the teacher model generally enhances the performance of encoder-only and encoder-decoder models.

Impact of scaling student model Row 7-8 of Table 4 illustrates the impact of scaling the student model on the performance of learned sparse retrieval (LSR). Both encoder-decoder and decoder-only backbones showed improvements as the student model parameters were scaled. Specifically, the nDCG@10 of the encoder-decoder backbone increased to 70.7 on DL2020 and 73.3 on DL2019, while the MRR@10 rose to 37.8 on MS MARCO. Similarly, the nDCG@10 of the decoder-only backbone increased to 60.9 on DL2020 and 73.1 on DL2019, with the MRR@10 rising to 36.9 on MS MARCO. Notably, the improvement observed for decoder-only backbones was significantly larger when scaling the student model compared to scaling the teacher model.

6   Related Work

6.1   Dense Retrieval

Dense retrieval encode queries and documents into latent dense representations. Early work on dense retrieval models used encoder transformers, but more recent efforts have been exploring other transformer architectures. Sentence T5 [24] compares different transformer architectures (encoder-decoder, encoder-only, decoder-only) for dense sentence embeddings and finds that the encoder-decoder architecture achieves the best performance on semantic textual similarity (STS) benchmarks. Ni et al. [25] later extend Sentence T5 (encoder-only) for retrieval tasks and study its effectiveness on different model scales (base, large, XL, XXL). The study suggests that scaling up the T5-encoder could improve the out-of-domain generalizability of dense retrieval. Similarly, Ma et al. [17] developed RepLlama based on Llama—a billion-scale decoder-only language model—and demonstrated the strong in-domain and out-of-domain performance of this model.

6.2   Learned Sparse Representations

Learned sparse retrieval offers an efficient alternative to dense retrieval by effectively utilizing traditional inverted indexes, which combines the capabilities of neural methods with the efficiency of classical indexing method. DeepCT [8] and uniCOIL [16] use BERT-based model to learn the importance of terms in query and passage. Although they do improve the weighting of terms compared to unsupervised approaches such as BM25, they only weight those terms that are already present in the document, hence lacking term expansions. Therefore, their effectiveness could be limited due to the term mismatch issue. To overcome this term mismatch problem, Mallia et al [22] improve effectiveness by enriching documents with predicted queries from DocT5Query and proposes a term weighting model that calculates the pairwise loss between relevant and nonrelevant texts with respect to a query. In another approach, Splade [12] directly utilizes language model logits to weight and expand terms in an end-to-end fashion. MacAvaney et al. [20] previously proposed EPIC, which has a similar architecture to Splade but with an MLP query encoder that does not perform query expansion. In a later

study, Nguyen et al. [23] found that under state-of-the-art training configurations, EPIC, without query expansion capability, can be as effective as Splade while being more efficient. This finding is also confirmed by a recent Splade-v3 [15] technical report. Unlike dense retrieval, Most of the previous LSR methods has predominantly focused on using small-scale encoder-only transformer models (e.g., DistilBERT, BERT). A concurrent work [9] uses 7 billion model Mistral as the backbone to develop a learned sparse retriever trained on extensive data. To ensure fair comparisons and avoid confounding factors, our study uses the same training data with Splade-DistillMax [15]. We address the aforementioned gap in the literature by investigating the performance of LSR across a range of transformer-based architectures, including encoder-only, decoder-only, and encoder-decoder models, as well as larger model scales.

7   Conclusion

Our work explored the utility of different transformer-based backbones for LSR. We highlight the difficulty of creating effective sparse representations using LLMs in zero-shot settings. Doing so leads to either the wrong type of term expansion or to a reduction in performance due to an inability to expand. The contribution of our work is to propose a solution for creating sparse representation from the decoder output and investigate the LSR effectiveness across different backbones for LSR. Our results indicate that incorporating the multi-tokens decoding approach helps to create a more effective sparse representation for the encoder-decoder and decoder-only backbone.

Acknowledgments

References

1. Arpit, D., Zhou, Y., Ngo, H.Q., Govindaraju, V.: Why regularized auto-encoders learn sparse representation? In: Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48. p. 136–144. ICML'16, JMLR.org (2016)

2. Bai, Y., Li, X., Wang, G., Zhang, C., Shang, L., Xu, J., Wang, Z., Wang, F., Liu, Q.: Sparterm: Learning term-based sparse representation for fast text retrieval (2020)

3. Bajaj, P., Campos, D., Craswell, N., Deng, L., Gao, J., Liu, X., Majumder, R., McNamara, A., Mitra, B., Nguyen, T., Rosenberg, M., Song, X., Stoica, A., Tiwary, S., Wang, T.: MS MARCO: A human generated machine reading comprehension dataset (2018)

4. Bruch, S., Nardini, F.M., Rulli, C., Venturini, R.: Efficient inverted indexes for approximate retrieval over learned sparse representations. In: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR). pp. 152–162. ACM (2024). `https://doi.org/10.1145/3626772.3657769`, `https://doi.org/10.1145/3626772.3657769`

5. Chung, H.W., Hou, L., Longpre, S., Zoph, B., Tai, Y., Fedus, W., Li, Y., Wang, X., Dehghani, M., Brahma, S., Webson, A., Gu, S.S., Dai, Z., Suzgun, M., Chen, X., Chowdhery, A., Castro-Ros, A., Pellat, M., Robinson, K., Valter, D., Narang, S., Mishra, G., Yu, A., Zhao, V., Huang, Y., Dai, A., Yu, H., Petrov, S., Chi, E.H., Dean, J., Devlin, J., Roberts, A., Zhou, D., Le, Q.V., Wei, J.: Scaling instruction-finetuned language models. J. Mach. Learn. Res. **25**(1) (Jan 2024)

6. Craswell, N., Mitra, B., Yilmaz, E., Campos, D.: Overview of the trec 2020 deep learning track (2021)

7. Craswell, N., Mitra, B., Yilmaz, E., Campos, D., Voorhees, E.M.: Overview of the trec 2019 deep learning track (2020)

8. Dai, Z., Callan, J.: Context-aware sentence/passage term importance estimation for first stage retrieval (2019)

9. Doshi, M., Kumar, V., Murthy, R., P, V., Sen, J.: Mistral-splade: Llms for better learned sparse retrieval (2024), `https://arxiv.org/abs/2408.11119`

10. Fang, H., Tao, T., Zhai, C.: A formal study of information retrieval heuristics. In: Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval. pp. 49–56 (2004)

11. Formal, T., Lassance, C., Piwowarski, B., Clinchant, S.: 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. p. 2353–2359. SIGIR '22, Association for Computing Machinery, New York, NY, USA (2022). `https://doi.org/10.1145/3477495.3531857`, `https://doi.org/10.1145/3477495.3531857`

12. Formal, T., Lassance, C., Piwowarski, B., Clinchant, S.: Splade v2: Sparse lexical and expansion model for information retrieval (2021). `https://doi.org/10.48550/ARXIV.2109.10086`, `https://arxiv.org/abs/2109.10086`

13. Gao, L., Callan, J.: Unsupervised corpus aware language model pre-training for dense passage retrieval. In: Muresan, S., Nakov, P., Villavicencio, A. (eds.) Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 2843–2853. Association for Computational Linguistics, Dublin, Ireland (May 2022). `https://doi.org/10.18653/v1/2022.acl-long.203`, `https://aclanthology.org/2022.acl-long.203/`

14. Hofstätter, S., Althammer, S., Schröder, M., Sertkan, M., Hanbury, A.: Improving efficient neural ranking models with cross-architecture knowledge distillation (2021)
15. Lassance, C., Déjean, H., Formal, T., Clinchant, S.: Splade-v3: New baselines for splade (2024)
16. Lin, J., Ma, X.: A few brief notes on deepimpact, coil, and a conceptual framework for information retrieval techniques (2021)
17. Ma, X., Wang, L., Yang, N., Wei, F., Lin, J.: Fine-tuning llama for multi-stage text retrieval. In: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval. p. 2421–2425. SIGIR '24, Association for Computing Machinery, New York, NY, USA (2024). https://doi.org/10.1145/3626772.3657951, https://doi.org/10.1145/3626772.3657951
18. Ma, X., Zhang, X., Pradeep, R., Lin, J.: Zero-shot listwise document reranking with a large language model (2023)
19. MacAvaney, S., Nardini, F.M., Perego, R., Tonellotto, N., Goharian, N., Frieder, O.: Expansion via prediction of importance with contextualization. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '20, ACM (Jul 2020)
20. MacAvaney, S., Nardini, F.M., Perego, R., Tonellotto, N., Goharian, N., Frieder, O.: Expansion via prediction of importance with contextualization. In: Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval. pp. 1573–1576 (2020)
21. Malkov, Y.A., Yashunin, D.A.: Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. IEEE Trans. Pattern Anal. Mach. Intell. **42**(4), 824–836 (Apr 2020). https://doi.org/10.1109/TPAMI.2018.2889473, https://doi.org/10.1109/TPAMI.2018.2889473
22. Mallia, A., Khattab, O., Suel, T., Tonellotto, N.: Learning passage impacts for inverted indexes. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. p. 1723–1727. SIGIR '21, Association for Computing Machinery, New York, NY, USA (2021). https://doi.org/10.1145/3404835.3463030, https://doi.org/10.1145/3404835.3463030
23. Nguyen, T., MacAvaney, S., Yates, A.: A unified framework for learned sparse retrieval. In: Advances in Information Retrieval: 45th European Conference on Information Retrieval, ECIR 2023, Dublin, Ireland, April 2–6, 2023, Proceedings, Part III. pp. 101–116. Springer (2023)
24. Ni, J., Hernandez Abrego, G., Constant, N., Ma, J., Hall, K., Cer, D., Yang, Y.: Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models. In: Findings of the Association for Computational Linguistics: ACL 2022. pp. 1864–1874. Association for Computational Linguistics, Dublin, Ireland (May 2022). https://doi.org/10.18653/v1/2022.findings-acl.146, https://aclanthology.org/2022.findings-acl.146/
25. Ni, J., Qu, C., Lu, J., Dai, Z., Hernandez Abrego, G., Ma, J., Zhao, V., Luan, Y., Hall, K., Chang, M.W., Yang, Y.: Large dual encoders are generalizable retrievers. In: Goldberg, Y., Kozareva, Z., Zhang, Y. (eds.) Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing. pp. 9844–9855. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates (Dec 2022). https://doi.org/10.18653/v1/2022.emnlp-main.669, https://aclanthology.org/2022.emnlp-main.669
26. Nogueira, R., Yang, W., Lin, J., Cho, K.: Document expansion by query prediction. arXiv preprint arXiv:1904.08375 (2019)

27. Paria, B., Yeh, C.K., Yen, I.E., Xu, N., Ravikumar, P., Póczos, B.: Minimizing flops to learn efficient sparse representations. In: International Conference on Learning Representations (2020), `https://openreview.net/forum?id=SygpC6Ntvr`

28. Pradeep, R., Sharifymoghaddam, S., Lin, J.: Rankvicuna: Zero-shot listwise document reranking with open-source large language models (2023)

29. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. ArXiv (2019)

30. Sun, W., Yan, L., Ma, X., Ren, P., Yin, D., Ren, Z.: Is chatgpt good at search? investigating large language models as re-ranking agent. In: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. pp. 14918–14937. Association for Computational Linguistics, Singapore (December 6–10 2023)

31. Wang, W., Wei, F., Dong, L., Bao, H., Yang, N., Zhou, M.: Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. Advances in neural information processing systems **33**, 5776–5788 (2020)

32. Yang, P., Fang, H., Lin, J.: Anserini: Enabling the use of lucene for information retrieval research. In: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. p. 1253–1256. SIGIR '17, Association for Computing Machinery, New York, NY, USA (2017)

33. Yang, P., Fang, H., Lin, J.: Anserini: Reproducible ranking baselines using lucene. J. Data and Information Quality **10**(4) (oct 2018)

34. Zamani, H., Dehghani, M., Croft, W.B., Learned-Miller, E., Kamps, J.: 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. p. 497–506. CIKM '18, Association for Computing Machinery, New York, NY, USA (2018)

35. Zhang, S., Roller, S., Goyal, N., Artetxe, M., Chen, M., Chen, S., Dewan, C., Diab, M., Li, X., Lin, X.V., Mihaylov, T., Ott, M., Shleifer, S., Shuster, K., Simig, D., Koura, P.S., Sridhar, A., Wang, T., Zettlemoyer, L.: Opt: Open pre-trained transformer language models (2022)

36. Zhao, T., Lu, X., Lee, K.: SPARTA: Efficient open-domain question answering via sparse transformer matching retrieval. In: Toutanova, K., Rumshisky, A., Zettlemoyer, L., Hakkani-Tur, D., Beltagy, I., Bethard, S., Cotterell, R., Chakraborty, T., Zhou, Y. (eds.) Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. pp. 565–575. Association for Computational Linguistics, Online (Jun 2021). `https://doi.org/10.18653/v1/2021.naacl-main.47`, `https://aclanthology.org/2021.naacl-main.47/`

37. Zhuang, S., Ma, X., Koopman, B., Lin, J., Zuccon, G.: PromptReps: Prompting large language models to generate dense and sparse representations for zero-shot document retrieval. In: Al-Onaizan, Y., Bansal, M., Chen, Y.N. (eds.) Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing. pp. 4375–4391. Association for Computational Linguistics, Miami, Florida, USA (Nov 2024). `https://doi.org/10.18653/v1/2024.emnlp-main.250`, `https://aclanthology.org/2024.emnlp-main.250/`

38. Zhuang, S., Zuccon, G.: 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. p. 1483–1492. SIGIR '21, Association for Computing Machinery, New York, NY, USA (2021)

Appendix

Table 5: Teacher performance comparison on the MS MARCO training set.

	nDCG@10	RR@10
MiniLM-L-6-v2	44.3	37.5
Rankllama-13B	47.5	40.3
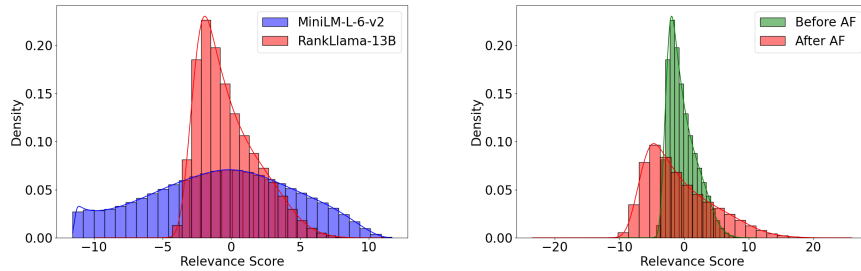
(a) Score distributions of MiniLM-L-6-v2 and RankLLama-13B.

(b) RankLLama-13B scores before and after affine transformation.

Fig. 3: Score distributions of the two teacher models on the MS MARCO training set. (a) RankLLama-13B exhibits a sharper distribution than MiniLM-L-6-v2. (b) We apply an affine transformation to align the mean and standard deviation distribution of RankLama-13B with that of MiniLM-L-6-v2.

Table 6: Comparison of LSR effectiveness across different backbones without fine-tuning on MS MARCO dev. × results are from [22]; A † indicates paired t-test $p < 0.05$. († indicates test between term expansion and non-expansion.

Model	Weighting	Expansion	RR@10	R@1k
(1) Enc.$_{\text{Distilbert-base-uncased}}$	Yes	No	0.930	40.4
(2) EncDec.MultiTokens$_{\text{FlanT5-base}}$	Yes	No	0	0
(3) EncDec.MultiTokens$_{\text{FlanT5-xl}}$	Yes	No	$3.67^{\dagger}$	$34.0^{\dagger}$
(4) Dec.MultiTokens$_{\text{OPT-350M}}$	Yes	No	13.2	50.9
(5) Dec.MultiTokens$_{\text{OPT-2.7B}}$	Yes	No	13.7	71.4
(6) Dec.MultiTokens$_{\text{OPT-6.7B}}$	Yes	No	12.5	70.5