

Efficient Multi-Vector Dense Retrieval with Bit Vectors

Franco Maria Nardini¹[0000-0003-3183-334X],
Cosimo Rulli¹[0000-0003-0194-361X], and
Rossano Venturini²[0000-0002-9830-3936]

¹ ISTI-CNR, Pisa, Italy {name.surname}@isti.cnr.it
² University of Pisa, Italy rossano.venturini@unipi.it

Abstract. Dense retrieval techniques employ pre-trained large language models to build a high-dimensional representation of queries and passages. These representations compute the relevance of a passage w.r.t. to a query using efficient similarity measures. In this line, multi-vector representations show improved effectiveness at the expense of a one-order-of-magnitude increase in memory footprint and query latency by encoding queries and documents on a per-token level. Recently, PLAID has tackled these problems by introducing a centroid-based term representation to reduce the memory impact of multi-vector systems. By exploiting a centroid interaction mechanism, PLAID filters out non-relevant documents, thus reducing the cost of the successive ranking stages. This paper proposes “Efficient Multi-Vector dense retrieval with Bit vectors” (EMVB), a novel framework for efficient query processing in multi-vector dense retrieval. First, EMVB employs a highly efficient pre-filtering step of passages using optimized bit vectors. Second, the computation of the centroid interaction happens column-wise, exploiting SIMD instructions, thus reducing its latency. Third, EMVB leverages Product Quantization (PQ) to reduce the memory footprint of storing vector representations while jointly allowing for fast late interaction. Fourth, we introduce a per-document term filtering method that further improves the efficiency of the last step. Experiments on MS MARCO and LoTTE show that EMVB is up to 2.8× faster while reducing the memory footprint by 1.8× with no loss in retrieval accuracy compared to PLAID.

Keywords: Dense Retrieval · Multi-vector · Efficiency · Bit Vectors.

1 Introduction

The introduction of pre-trained large language models (LLM) has remarkably improved the effectiveness of information retrieval systems [13,8,26,2], thanks to the well-known ability of LLMs to model semantic and context [12,1,3]. In dense retrieval, LLMs have been successfully exploited to learn high-dimensional dense representations of passages and queries. These learned representations allow answering the user query through fast similarity operations, i.e., inner product or L2 distance. In this line, multi-vector techniques [14,20] employ an LLM to build

a dense representation for each token of a passage. These approaches offer superior effectiveness compared to single-vector techniques [24,27] or sparse retrieval techniques [5]. In this context, the similarity between the query and the passage is measured by using the *late interaction* mechanism [14,20], which works by computing the sum of the maximum similarities between each term of the query and each term of a candidate passage. The improved effectiveness of multi-vector retrieval system comes at the price of its increased computational burden. First, producing a vector for each token causes the number of embeddings to be orders of magnitude larger than in a single-vector representation. Moreover, due to the large number of embeddings, identifying the candidate documents³ is time-consuming. In addition, the late interaction step requires computing the maximum similarity operator between all the candidate embeddings and the query, which is also time-consuming.

Early multi-vector retrieval systems, i.e., ColBERT [14], exploit an inverted index to store the embeddings and retrieve the candidate passages. Then, the representations of passages are retrieved and employed to compute the maximum similarity score with the query. Despite being quite efficient, this approach requires maintaining the full-precision representation of each document term in memory. On MS MARCO [17], a widely adopted benchmark dataset for passage retrieval, the entire collection of embeddings used by ColBERT requires more than 140 GB [14] to be stored. ColBERTv2 [20] introduces a centroid-based compression technique to store the passage embeddings efficiently. Each embedding is stored by saving the *id* of the closest centroid and then compressing the residual (i.e., the element-wise difference) by using 1 or 2 bits per component. ColBERTv2 saves up to 10× space compared to ColBERT while being significantly more inefficient on modern CPUs, requiring up to 3 seconds to perform query processing on CPU [19]. The reduction of query processing time is achieved by Santhanam *et al.* with PLAID [19]. PLAID takes advantage of the embedding compressor of ColBERTv2 and also uses the centroid-based representation to discard non-relevant passages (*centroid interaction* [19]), thus performing the late interaction exclusively on a carefully selected batch of passages. PLAID allows for massive speedup compared to ColBERTv2, but its average query latency can be up to 400 msec. on CPU with single-thread execution [19].

This paper presents EMVB, a novel framework for efficient query processing with multi-vector dense retrieval. First, we identify the most time-consuming steps of PLAID. These steps are i) extracting the top-*nprobe* closest centroids for each query term during the candidate passage selection, ii) computing the centroid interaction mechanism, and iii) decompression of the quantized residuals. Our method tackles the first and the second steps by introducing a highly efficient passage filtering approach based on optimized bit vectors. Our filter identifies a small set of crucial centroid scores, thus tearing down the cost of top-*nprobe* extraction. At the same time, it reduces the amount of passages for which we have to compute the centroid interaction. Moreover, we introduce a highly efficient column-wise reduction exploiting SIMD instructions to speed up

³ the terms “document” and “passage” are used interchangeably in this paper.

this step. Finally, we improve the efficiency of the late interaction by introducing Product Quantization (PQ) [9]. PQ allows to obtain in pair or superior performance compared to the bitwise compressor of PLAID while being up to $3\times$ faster. Finally, to further improve the efficiency of the last step of our pipeline, we introduce a dynamic passage-term-selection criterion for late interaction, thus reducing the cost of this step up to 30%.

We experimentally evaluate EMVB against PLAID on two datasets: MS MARCO passage [17] (for in-domain evaluation) and LoTTE [20] (for out-of-domain evaluation). Results on MS MARCO show that EMVB is up to $2.8\times$ faster while reducing the memory footprint by $1.8\times$ with no loss in retrieval accuracy compared to PLAID. On the out-of-domain evaluation, EMVB delivers up to $2.9\times$ speedup compared to PLAID, with a minimal loss in retrieval quality.

The rest of this paper is organized as follows. In Section 2, we discuss the related work. In Section 3 we describe PLAID [19], the current state-of-the-art in multi-vector dense retrieval. We introduce EMVB in Section 4 and we experimentally evaluate it against PLAID in Section 5. Finally, Section 6 concludes our work.

2 Related Work

Dense retrieval encoders can be broadly classified into single-vector and multi-vector techniques. Single-vector encoders allow the encoding of an entire passage in a single dense vector [11]. In this line, ANCE [25] and STAR/ADORE [26] employ hard negatives to improve the training of dense retrievers by teaching them to distinguish between lexically-similar positive and negative passages. Multi-vector encoders have been introduced with ColBERT. The limitations of ColBERT and the efforts done to overcome them (ColBERTv2, PLAID) have been discussed in Section 1. COIL [6] rediscover the lessons of classical retrieval systems (e.g., BM25) by limiting the token interactions to lexical matching between queries and documents. CITADEL [16] is a recently proposed approach that introduces conditional token interaction by using dynamic lexical routing. Conditional token interaction means that the relevance of the query of a specific passage is estimated by only looking at some of their tokens. These tokens are selected by the so-called lexical routing, where a module of the ranking architecture is trained to determine which of the keys, i.e., words in the vocabulary, are activated by a query/passage. CITADEL significantly reduces the execution time on GPU, but turns out to be $2\times$ slower than PLAID on CPU, at the same retrieval quality. Multi-vector dense retrieval is also exploited in conjunction with pseudo-relevance feedback both in ColBERT-PRF [23] and in CWPRF [22], showing that their combination boosts the effectiveness of the model.

Our Contribution: This work advances the state of the art of multi-vector dense retrieval by introducing EMVB, a novel framework that allows to speed up the retrieval performance of the PLAID pipeline significantly. To the best of our knowledge, this work is the first in the literature that proposes a highly efficient document filtering approach based on optimized bit vectors, a column-wise SIMD

reduction to retrieve candidate passages and a late interaction mechanism that combines product quantization with a per-document term filtering.

3 Multi-vector Dense Retrieval

Consider a passage corpus \mathcal{P} with n_P passages. In a multi-vector dense retrieval scenario, an LLM encodes each token in \mathcal{P} into a dense d -dimensional vector T_j . For each passage P , a dense representation $P = \{T_j\}$, with $j = 0, \dots, n_t$, is produced, where n_t is the number of tokens in the passage P . Employing a token-level dense representation allows for boosting the effectiveness of the retrieval systems [14,20,19]. On the other hand, it produces significantly large collections of d -dimensional vectors posing challenges to the applicability of such systems in real-world search scenarios both in terms of space (memory requirements) and time (latency of the query processor). To tackle the problem of memory requirements, ColBERTv2 [20] and successively PLAID [19] exploit a centroid-based vector compression technique. First, the K-means algorithm is employed to devise a clustering of the d -dimensional space by identifying the set of k centroids $\mathcal{C} = \{C_i\}_{i=1}^{n_c}$. Then, for each vector x , the residual r between x and its closest centroid \bar{C} is computed so that $r = x - \bar{C}$. The residual r is compressed into \tilde{r} using a b -bit encoder that represents each dimension of r using b bits, with $b \in \{1, 2\}$. The memory impact of storing a d -dimensional vector is given by $\lceil \log_2 |\mathcal{C}| \rceil$ bits for the centroid index and $d \times b$ bits for the compressed residual. This approach requires a time-expensive decompression phase to restore the approximate full-precision vector representation given the centroid id and the residual coding. For this reason, PLAID aims at decompressing as few candidate documents as possible. This is achieved by introducing a high-quality filtering step based on the centroid-approximated embedding representation, named *centroid interaction* [19]. In detail, the PLAID retrieval engine is composed of four different phases [19]. The first one regards the *retrieval* of the candidate passages. A list of candidate passages is built for each centroid. A passage belongs to a centroid C_i candidate list if one or more tokens have C_i as its closest centroid. For each query term q_i , with $i = 1, \dots, n_q$, the *top- $nprobe$* closest centroids are computed, according to the *dot product* similarity measure. The set of unique documents associated with the *top- $nprobe$* centroids then moves to a second phase that acts as a *filtering* phase. In this phase, a token embedding T_j with $j = 1, \dots, n_p$ is approximated using its closest centroid \bar{C}^{T_j} . Hence, its distance with the i -th query term q_i is approximated with

$$q_i \cdot T_j \simeq q_i \cdot \bar{C}^{T_j} = \tilde{T}_{i,j}. \quad (1)$$

Consider a candidate passage P composed of n_p tokens. The approximated score of P consists in computing the dot product $q_i \cdot \bar{C}^{T_j}$ for all the query terms q_i and all the closest centroids of each token belonging to the passage, i.e.,

$$\bar{S}_{q,P} = \sum_{i=1}^{n_q} \max_{j=1 \dots n_t} q_i \cdot \bar{C}^{T_j} \quad (2)$$

The third phase, named *decompression*, aims at reconstructing the full-precision representation of P by combining the centroids and the residuals. This is done on the top- n_{docs} passages selected according to the *filtering* phase [19]. In the fourth phase, PLAID recomputes the final score of each passage with respect to the query q using the decompressed—full-precision—representation according to *late interaction* mechanism (Equation 3). Passages are then ranked according to their similarity score and the top- k passages are selected.

$$S_{q,P} = \sum_{i=1}^{n_q} \max_{j=1\dots n_t} q_i \cdot T_j. \quad (3)$$

PLAID execution time. We provide a breakdown of PLAID execution time across its different phases, namely *retrieval*, *filtering*, *decompression*, and *late interaction*. This experiment is conducted using the experimental settings detailed in Section 5. We report the execution time for different values of k , i.e., the number of retrieved passages.




Fig. 1: Breakdown of the PLAID average query latency (in milliseconds) on CPU across its four phases.

4 EMVB

We now present EMVB, our novel framework for efficient multi-vector dense retrieval. First, EMVB introduces a highly efficient pre-filtering phase that exploits optimized bit vectors. Second, we improve the efficiency of the centroid interaction step (Equation 1) by introducing column-wise max reduction with SIMD instructions. Third, EMVB leverages Product Quantization (PQ) to reduce the memory footprint of storing the vector representations while jointly allowing for a fast late interaction phase. Fourth, PQ is applied in conjunction with a novel per-passage term filtering approach that allows for further improving the efficiency of the late interaction. In the following subsections, we detail these four contributions behind EMVB.

4.1 Retrieval of Candidate Passages

Figure 1 shows that a consistent part of the computation required by PLAID is spent on the retrieval phase. We further break down these steps to evidence its most time-consuming part. The retrieval consists of i) computing the distances between the incoming query and the set of centroids, ii) extracting the top- $nprobe$ closest centroids for each query term. The former step is efficiently carried out by leveraging high-performance matrix multiplication tools (e.g., Intel MKL [18,21]). In the latter step, PLAID extracts the top- $nprobe$ centroids using the numpy `topk` function, which implements the *quickselect* algorithm. Selecting the top- $nprobe$ within the $|C| = 2^{18}$ centroids for each of the n_q query terms costs up to $3\times$ the matrix multiplication done in the first step. In Section 4.2, we show that our pre-filtering inherently speeds up the top- $nprobe$ selection by tearing down the number of evaluated elements. In practice, we show how to efficiently filter out those centroids whose score is below a certain threshold and then execute *quickselect* exclusively on the surviving ones. As a consequence, in EMVB the cost of the top- $nprobe$ extraction becomes negligible, being two orders of magnitude faster than the top- $nprobe$ extraction on the full set of centroids.

4.2 Efficient Pre-Filtering of Candidate Passages

Figure 1 shows that the candidate filtering phase can be significantly expensive, especially for large values of k . In this section, we propose a pre-filtering approach based on a novel bit vector representation of the centroids that efficiently allows the discarding of non-relevant passages.

Given a passage P , our pre-filtering consists in determining whether $\tilde{T}_{i,j}$, for $i = 1, \dots, n_q$, $j = 1, \dots, n_t$ is large or not. Recall that $\tilde{T}_{i,j}$ represents the approximate score of the j -th token of passage P with respect to the i -th term of the query q_i , as defined in Equation 1. This can be obtained by checking whether \tilde{C}_j^T —the centroid associated with T_j —belongs to the set of the *closest centroids* of q_i . We introduce \mathbf{close}_i^{th} , the set of centroids whose scores are greater than a certain threshold th with respect to a query term q_i . Given a passage P , we define the list of centroids ids I_P , where I_P^j is the centroid id of \tilde{C}^{T_j} . The similarity of a passage with respect to a query can be estimated with our novel filtering function $F(P, q) \in [0, n_q]$ with the following equation:

$$F(P, q) = \sum_{i=1}^{n_q} \mathbf{1}(\exists j \text{ s.t. } I_P^j \in \mathbf{close}_i^{th}). \quad (4)$$

For a passage P , this counts how many query terms have at least one similar passage term in P , where “similar” describes the belonging of T_j to \mathbf{close}_i^{th} .

In Figure 2 (left), we compare the performance of our novel pre-filter working on top of the centroid interaction mechanism (orange, blue, green lines) against the performance of the centroid interaction mechanism on the entire set of candidate documents (red dashed line) on the MS MARCO dataset. The plot shows




Fig. 2: R@100 with various values of the threshold (**left**). Comparison of different algorithms to construct close_i^{th} , for different values of th (**right**).

that our pre-filtering allows to efficiently discard non-relevant passages without harming the recall of the successive centroid interaction phase. For example, we can narrow the candidate passage set to just 1000 elements using $th = 0.4$ without any loss in R@100. In the remainder of this section, we show how to implement this pre-filter efficiently.

Building the bit vectors. Given th , the problem of computing close_i^{th} is conceptually simple. Yet, an efficient implementation carefully considering modern CPUs’ features is crucial for fast computation of Equation 4.

Let $CS = q \cdot C^T$, with $CS \in [-1, 1]^{n_q \times |C|}$ be the score matrix between the query q and the set of centroids C (both matrices are L_2 normalized), where n_q is the number of query tokens, and $|C|$ is the number of centroids. In the naïve *if*-based solution, we scan the i -th row of CS and select those j s.t. $CS_{i,j} > th$. It is possible to speed up this approach by taking advantage of SIMD instructions. In particular, the `_mm512_cmp_epi32_mask` allows one to compare 16 fp32 values at a time and store the comparison result in a *mask* variable. If `mask == 0`, we can skip to the successive 16 values because the comparison has failed for all the current js . Otherwise, we extract those indexes $J = \{j \in [0, 15] \mid \text{mask}_j = 1\}$.

The efficiency of such *if*-based algorithms mainly depends on the *branch misprediction* ratio. Modern CPUs speculate on the outcome of the *if* before the condition itself is computed by recognizing patterns in the execution flow of the algorithm. When the wrong branch is predicted, a *control hazard* happens, and the pipeline is flushed with a delay of 15–20 clock cycles, i.e., about 10 ns. We tackle the inefficiency of branch misprediction by proposing a *branchless* algorithm. The branchless algorithm employs a pointer p addressing a pre-allocated buffer. While scanning $CS_{i,j}$, it writes j in the position indicated by p . Then, it sums to p the result of the comparison: 1 if $CS_{i,j} > th$, 0 otherwise. At the successive iteration, if the result of the comparison was 0, $j + 1$ will override j . Otherwise, it will be written in the successive memory location, and j will be saved in the buffer. The branchless selection does not present any *if* instruction and consequently does not contain any branch in its execution flow. The branchless algorithm can be implemented more efficiently by leveraging SIMD instructions. In particular, the above-mentioned `_mm512_cmp_epi32_mask` instruction allows to compare 16

fp32 values at the time, and the `_mm512_mask_compressstore` allows to extract J in a single instruction.

Figure 2 (right) presents a comparison of our different approaches, namely “Naïve IF”, the “Vectorized IF”, the “Branchless”, and the “VecBranchless” described above. Branchless algorithms present a constant execution time, regardless of the value of the threshold, while *if*-based approaches offer better performances as the value of th increases. With $th \geq 0.3$, “Vectorized IF” is the most efficient approach, with a speedup up to $3\times$ compared to its naïve counterpart.

Fast set membership. Once close_i^{th} is computed, we have to efficiently compute Equation 4. Here, given I_P as a list of integers, we have to test if at least one of its members I_P^j belongs to close_i^{th} , with $i = 1, \dots, n_q$. This can be efficiently done using *bit vectors* for representing close_i^{th} . A bit vector maps a set of integers up to N into an array of N bits, where the e -th bit is set to one if and only if the integer e belongs to the set. Adding and searching any integer e can be performed in constant time with bit manipulation operators. Moreover, bit vectors require N bits to be stored. In our case, since we have $|C| = 2^{18}$, a bit vector only requires $32K$ bytes to be stored.

Since we search through all the n_q bit vectors at a time, we can further exploit the bit vector representation by stacking the bit vectors vertically (Figure 3). This allows to search a centroid index through all the close_i^{th} at a time. The bits corresponding to the same centroid for different query terms are consecutive and fit a 32-bit word. This way, we can simultaneously test the membership for all the queries in constant time with a single bitwise operation. In detail, our algorithm works by initializing a mask m of $n_q = 32$ bits at zeros (Step 1, Figure 3). Then, for each term in the candidate documents, it performs a bitwise `xor` between the mask and the 32-bit word representing the membership to all the query terms (Step 2, Figure 3). Hence, Equation 4 can be obtained by counting the number of 1s in m at the end of the execution with the `popcnt` operation featured by modern CPUs (Step 3, Figure 3).

Figure 4 (up) shows that our “Vectorized” set membership implementation delivers a speedup ranging from $10\times$ to $16\times$ a “Baseline” relying on a naïve usage of bit vectors. In particular, our bit vector-based pre-filtering can be up to $30\times$ faster than the centroid-interaction proposed in PLAID [19], cf. Figure 4 (down).

4.3 Fast Filtering of Candidate Passages

Our pre-filtering approach allows us to efficiently filter out non-relevant passages and is employed upstream of PLAID’s centroid interaction (Equation 2). We now show how to improve the efficiency of the centroid interaction itself.

Consider a passage P and its associated centroid scores matrix $\tilde{P} = q_i \cdot \bar{C}^{T_j}$. Explicitly building this matrix allows to reuse it in the scoring phase, in place of the costly decompression step (Section 4.4). To build \tilde{P} , we transpose CS into CS^T of size $|C| \times n_q$. The i -th row of CS^T allows access to all the n_q query terms scores for the i -th centroids. Given the ids of the closest centroids for

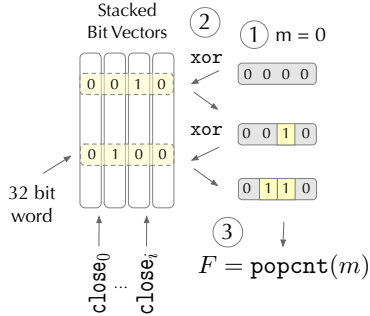


Fig. 3: Vectorized Fast Set Membership algorithm based on bit vectors.

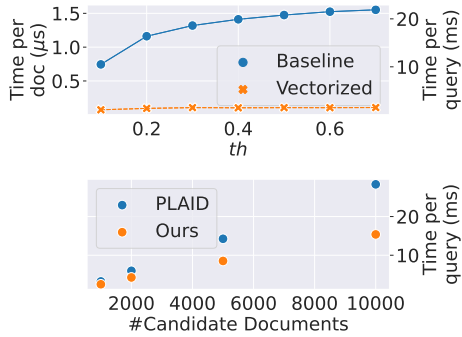


Fig. 4: Vectorized vs naïve Fast Set Membership (**up**). Ours vs PLAID filtering (**down**).

each passage term (defined as I_P in Section 4.2) we retrieve the scores for each centroid id. We build $\tilde{P}^T - \tilde{P}$ transposed—to allow the CPU to read and write contiguous memory locations. This grants more than $2\times$ speedup compared to processing \tilde{P} . We now have \tilde{P}^T of shape $n_t \times n_q$. We need to max-reduce along the columns and then sum the obtained values to implement Equation 2. This is done by iterating on the \tilde{P}^T rows and packing them into AVX512 registers. Given that $n_q = 32$, each AVX512 register can contain $512/32 = 16$ floating point values, so we need 2 registers for each row. We pack the first row into max_l and max_h . All the successive rows are packed into $current_l$ and $current_h$. At each iteration, we compare max_l with $current_l$ and max_h with $current_h$ using the `_mm512_cmp_ps_mask` AVX512 instruction described before. The output mask m is used to update the max_l and max_h by employing the `_mm512_mask_blend_ps` instruction. The `_mm512_cmp_ps_mask` has throughput 2 on IceLake Xeon CPUs, so each row of \tilde{P} is compared with max_l and max_h in the same clock cycle, on two different ports. The same holds for the `_mm512_mask_blend_ps` instruction, entailing that the max-reduce operation happens in 2 clock cycles without considering the memory loading. Finally, max_l and max_h are summed together, and the function `_mm512_reduce_add_ps` is used to ultimate the computation.

We implement PLAID’s centroid interaction in C++ and we compare its filtering time against our SIMD-based solution. The results of the comparison are reported for different values of candidate documents in Figure 4 (down). Thanks to the proficient read-write pattern and the highly efficient column-wise max-reduction, our method can be up to $1.8\times$ faster than the filtering proposed in PLAID.

4.4 Late Interaction

The b -bit residual compressor proposed in previous approaches [20,19] requires a costly decompression step before the late interaction phase. Figure 1 shows that in PLAID decompressing the vectors costs up to $5\times$ the late interaction phase.

We propose compressing the residual r by employing Product Quantization (PQ) [9]. PQ allows the computation of the dot product between an input query vector q and the compressed residual r_{pq} without decompression. Consider a query q and a candidate passage P . We decompose the computation of the max similarity operator (Equation 3) into

$$S_{q,P} = \sum_{i=1}^{n_q} \max_{j=1\dots n_t} (q_i \cdot \bar{C}^{T_j} + q_i \cdot r^{T_j}) \simeq \sum_{i=1}^{n_q} \max_{j=1\dots n_t} (q_i \cdot \bar{C}^{T_j} + q_i \cdot r_{pq}^{T_j}), \quad (5)$$

where and $r^{T_j} = T_j - \bar{C}^{T_j}$. On the one hand, this decomposition allows to exploit the pre-computed \bar{P} matrix. On the other hand, thanks to PQ, it computes the dot product between the query and the residuals without decompression.

We replace PLAID’s residual compression with PQ, particularly with JMPQ [4], which optimizes the codes of product quantization during the fine-tuning of the language model for the retrieval task. We tested $m = \{16, 32\}$, where m is the number of sub-spaces used to partition the vectors [9]. We experimentally verify that PQ reduces the latency of the late interaction phase up to $3.6\times$ compared to PLAID b -bit compressor. Moreover, it delivers the same ($m = 16$) or superior performance ($m = 32$) in terms of MRR@10 when leveraging the JMPQ version.

We propose to further improve the efficiency of the scoring phase by hinging on the properties of Equation 5. We experimentally observe that, in many cases, $q_i \cdot \bar{C}_j^{T_j} > q_i \cdot r_{pq}^{T_j}$, meaning that the \max operator on j , in many cases, is lead by the score between the query term and the centroid, rather than the score between the query term and the residual. We argue that it is possible to compute the scores on the residuals only for a reduced set of document terms \bar{J}_i , where i identifies the index of the query term. In particular, $\bar{J}_i = \{j | q_i \cdot \bar{C}_j^{T_j} > th_r\}$, where th_r is a second threshold that determines whether the score with the centroid is sufficiently large. With the introduction of this new per-term filter, Equation 5 now becomes computing the max operator on the set of passages in \bar{J}_i , i.e.,

$$S_{q,P} = \sum_{i=1}^{n_q} \max_{j \in \bar{J}_i} (q_i \cdot \bar{C}^{T_j} + q_i \cdot r_{pq}^{T_j}). \quad (6)$$

In practice, we compute the residual scores only for those document terms whose centroid score is large enough. If $\bar{J}_i = \emptyset$, we compute $S_{q,P}$ as in Equation 5. Figure 5 (left) reports the effectiveness of our approach. On the y -axis, we report the percentage of the original effectiveness, computed as the ratio between the MRR@10 computed with Equation 6 and Equation 5. Filtering document terms according to Equation 6 does not harm the retrieval quality, as it delivers substantially the same MRR@10 of Equation 5. On the right side of Figure 5, we report the percentage of scored terms compared to the number of document

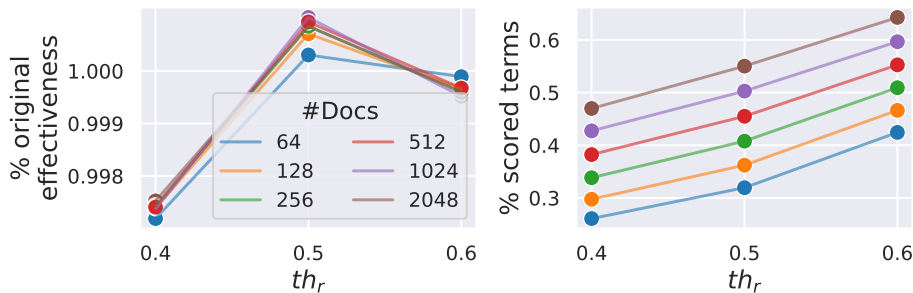


Fig. 5: Performance of our dynamic term-selection filtering for different values of th_r , in terms of percentage of original effectiveness (**left**) and in terms of percentage of original number of scored terms (**right**). The percentage of original effectiveness is computed as the ratio between the MRR@10 computed with Equation 6 and Equation 5.

terms computed using Equation 5. With $th_r = 0.5$, we are able to reduce the number of scored terms of at least 30% (right) without any performance degradation in terms of MRR@10.

5 Experimental Evaluation

Experimental Settings. This section compares our methodology against the state-of-the-art engine for multi-vector dense retrieval, namely PLAID [19]. We conduct experiments on the MS MARCO passages dataset [17] for the in-domain evaluation and on LoTTE [20] for the out-of-domain evaluation. We generate the embeddings for MS MARCO using the ColBERTv2 model. The generated dataset is composed of about 600M d -dimensional vectors, with $d = 128$. Product Quantization is implemented using the FAISS [10] library, and optimized using the JMPQ technique [4] on MS MARCO. The C++ implementation of EMVB will be made publicly available upon publication of this work. We compare EMVB against the original PLAID implementation [19], which also implements its core components in C++. Experiments are conducted on an Intel Xeon Gold 5318Y CPU clocked at 2.10 GHz, equipped with the AVX512 instruction set, with single-thread execution. Code is compiled using GCC 11.3.0 (with `-O3` compilation options) on a Linux 5.15.0-72 machine. When running experiments with AVX512 instruction on 512-bit registers, we ensure not to incur in the frequency scaling down event reported for Intel CPUs [15].

Evaluation. Table 1 compares EMVB against PLAID on the MS MARCO dataset, in terms of memory requirements (num. of bytes per embedding), average query latency (in milliseconds), MRR@10, and Recall@100, and 1000.

Results show that EMVB delivers superior performance along both the evaluated trade-offs. With $m = 16$, EMVB almost halves the per-vector memory

k	Method	Latency (<i>msec.</i>)	Bytes	MRR@10	R@100	R@1000
10	PLAID	131	36	39.4	-	-
	EMVB ($m=16$)	62 (2.1 \times)	20	39.4	-	-
	EMVB ($m=32$)	61 (2.1 \times)	36	39.7	-	-
100	PLAID	180	36	39.8	90.6	-
	EMVB ($m=16$)	68 (2.6 \times)	20	39.5	90.7	-
	EMVB ($m=32$)	80 (2.3 \times)	36	39.9	90.7	-
1000	PLAID	260	36	39.8	91.3	97.5
	EMVB ($m=16$)	93 (2.8 \times)	20	39.5	91.4	97.5
	EMVB ($m=32$)	104 (2.5 \times)	36	39.9	91.4	97.5

Table 1: Comparison between EMVB and PLAID in terms of average query latency, number of bytes per vector embeddings, MRR, and Recall on MS MARCO.

burden compared to PLAID, while being up to 2.8 \times faster with almost no performance degradation regarding retrieval effectiveness. By doubling the number of sub-partitions per vector, i.e., $m = 32$, EMVB outperforms the performance of PLAID in terms of MRR and Recall with the same memory footprint with up to 2.5 \times speed up.

Table 2 compares EMVB and PLAID in the out-of-domain evaluation on the LoTTE dataset. As in PLAID [19], we employ Success@5 and Success@100 as retrieval quality metrics. On this dataset, EMVB offers slightly inferior performance in terms of retrieval quality. Recall that JMPQ [4] cannot be applied in the out-of-domain evaluation due to the lack of training queries. Instead, we employ Optimized Product Quantization (OPQ) [7], which searches for an optimal rotation of the dataset vectors to reduce the quality degradation that comes with PQ. To mitigate the retrieval quality loss, we only experiment PQ with $m = 32$, given that an increased number of partitions offers a better representation of the original vector. On the other hand, EMVB can offer up to 2.9 \times speedup compared to PLAID. This larger speedup compared to MS MARCO is due to the larger average document lengths in LoTTE. In this context, filtering nonrelevant documents using our bit vector-based approach has a remarkable impact on efficiency. Observe that for the out-of-domain evaluation, our pre-filtering method could be ingested into PLAID. This would allow to maintain the PLAID accuracy together with EMVB efficiency. Combinations of PLAID and EMVB are left for future work.

6 Conclusion

We presented EMVB, a novel framework for efficient multi-vector dense retrieval. EMVB advances PLAID, the current state-of-the-art approach, by introducing four novel contributions. First, EMVB employs a highly efficient pre-filtering step of passages using optimized bit vectors for speeding up the candidate passage

k	Method	Latency (<i>msec.</i>)	Bytes	Success@5	Success@100
10	PLAID	131	36	69.1	-
	EMVB (m=32)	82 (1.6×)	36	69.0	-
100	PLAID	202	36	69.4	89.9
	EMVB (m=32)	129 (1.6×)	36	69.0	89.9
1000	PLAID	411	36	69.6	90.5
	EMVB (m=32)	142 (2.9×)	36	69.0	90.1

Table 2: Comparison between EMVB and PLAID in terms of average query latency, number of bytes per vector embeddings, Success@5, and Success@100 on LoTTE.

filtering phase. Second, the computation of the centroid interaction is carried out with reduced precision. Third, EMVB leverages Product Quantization to reduce the memory footprint of storing vector representations while jointly allowing for fast late interaction. Fourth, we introduce a per-passage term filter for late interaction, thus reducing the cost of this step of up to 30%. We experimentally evaluate EMVB against PLAID on two publicly available datasets, i.e., MS MARCO and LoTTE. Results show that, in the in-domain evaluation, EMVB is up to 2.8× faster, and it reduces by 1.8× the memory footprint with no loss in retrieval quality compared to PLAID. In the out-of-domain evaluation, EMVB is up to 2.9× faster with little or no retrieval quality degradation.

7 Acknowledgements

This work was partially supported by the EU - NGEU, by the PNRR - M4C2 - Investimento 1.3, Partenariato Esteso PE00000013 - “FAIR - Future Artificial Intelligence Research” - Spoke 1 “Human-centered AI” funded by the European Commission under the NextGeneration EU program, by the PNRR ECS00000017 Tuscany Health Ecosystem Spoke 6 “Precision medicine & personalized healthcare”, by the European Commission under the NextGeneration EU programme, by the Horizon Europe RIA “Extreme Food Risk Analytics” (EFRA), grant agreement n. 101093026, by the “Algorithms, Data Structures and Combinatorics for Machine Learning” (MIUR-PRIN 2017), and by the “Algorithmic Problems and Machine Learning” (MIUR-PRIN 2022).

References

1. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in Neural Information Processing Systems (NIPS)* (2020)
2. Bruch, S., Lucchese, C., Nardini, F.M.: Efficient and effective tree-based and neural learning to rank. *Found. Trends Inf. Retr.* (2023)

3. Dai, D., Sun, Y., Dong, L., Hao, Y., Sui, Z., Wei, F.: Why can gpt learn in-context? language models secretly perform gradient descent as meta optimizers. arXiv preprint arXiv:2212.10559 (2022)
4. Fang, Y., Zhan, J., Liu, Y., Mao, J., Zhang, M., Ma, S.: Joint optimization of multi-vector representation with product quantization. In: Natural Language Processing and Chinese Computing (2022)
5. Formal, T., Piwowarski, B., Clinchant, S.: 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 (2021)
6. Gao, L., Dai, Z., Callan, J.: Coil: Revisit exact lexical match in information retrieval with contextualized inverted list. In: Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (2021)
7. Ge, T., He, K., Ke, Q., Sun, J.: Optimized product quantization. IEEE Transactions on Pattern Analysis and Machine Intelligence (2013)
8. Guu, K., Lee, K., Tung, Z., Pasupat, P., Chang, M.: Retrieval augmented language model pre-training. In: Proceedings of the International Conference on Machine Learning (ICML) (2020)
9. Jegou, H., Douze, M., Schmid, C.: Product quantization for nearest neighbor search. IEEE Transactions on Pattern Analysis and Machine Intelligence (2010)
10. Johnson, J., Douze, M., Jegou, H.: Billion-scale similarity search with gpus. IEEE Transactions on Big Data (2021)
11. Karpukhin, V., Oguz, B., Min, S., Lewis, P., Wu, L., Edunov, S., Chen, D., Yih, W.t.: Dense passage retrieval for open-domain question answering. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). Association for Computational Linguistics (2020)
12. Kenton, J.D.M.W.C., Toutanova, L.K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of NAACL-HLT (2019)
13. Khattab, O., Potts, C., Zaharia, M.: Baleen: Robust multi-hop reasoning at scale via condensed retrieval. Advances in Neural Information Processing Systems (NIPS) (2021)
14. Khattab, O., Zaharia, M.: Colbert: Efficient and effective passage search via contextualized late interaction over bert. In: Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval. pp. 39–48 (2020)
15. Lemire, D., Downs, T.: Avx-512: when and how to use these new instructions (2023), <https://lemire.me/blog/2018/09/07/avx-512-when-and-how-to-use-these-new-instructions/>
16. Li, M., Lin, S.C., Oguz, B., Ghoshal, A., Lin, J., Mehdad, Y., Yih, W.t., Chen, X.: Citadel: Conditional token interaction via dynamic lexical routing for efficient and effective multi-vector retrieval. arXiv e-prints (2022)
17. Nguyen, T., Rosenberg, M., Song, X., Gao, J., Tiwary, S., Majumder, R., Deng, L.: Ms marco: A human-generated machine reading comprehension dataset
18. Qian, G., Sural, S., Gu, Y., Pramanik, S.: Similarity between euclidean and cosine angle distance for nearest neighbor queries. In: Proceedings of the 2004 ACM symposium on Applied computing (2004)
19. Santhanam, K., Khattab, O., Potts, C., Zaharia, M.: Plaid: an efficient engine for late interaction retrieval. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management (2022)

20. Santhanam, K., Khattab, O., Saad-Falcon, J., Potts, C., Zaharia, M.: Colbertv2: Effective and efficient retrieval via lightweight late interaction. In: Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (2022)
21. Wang, E., Zhang, Q., Shen, B., Zhang, G., Lu, X., Wu, Q., Wang, Y.: Intel math kernel library. In: High-Performance Computing on the Intel® Xeon Phi™ (2014)
22. Wang, X., MacAvaney, S., Macdonald, C., Ounis, I.: Effective contrastive weighting for dense query expansion. In: Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (2023)
23. Wang, X., Macdonald, C., Tonello, N., Ounis, I.: Colbert-prf: Semantic pseudo-relevance feedback for dense passage and document retrieval. *ACM Transactions on the Web* **17**(1), 1–39 (2023)
24. Xiong, L., Xiong, C., Li, Y., Tang, K.F., Liu, J., Bennett, P.N., Ahmed, J., Overwijk, A.: Approximate nearest neighbor negative contrastive learning for dense text retrieval. In: International Conference on Learning Representations
25. Xiong, L., Xiong, C., Li, Y., Tang, K.F., Liu, J., Bennett, P.N., Ahmed, J., Overwijk, A.: Approximate nearest neighbor negative contrastive learning for dense text retrieval. In: International Conference on Learning Representations (2020)
26. Zhan, J., Mao, J., Liu, Y., Guo, J., Zhang, M., Ma, S.: Optimizing dense retrieval model training with hard negatives. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (2021)
27. Zhan, J., Mao, J., Liu, Y., Guo, J., Zhang, M., Ma, S.: Learning discrete representations via constrained clustering for effective and efficient dense retrieval. In: Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining. pp. 1328–1336 (2022)