



Constructing Tree-based Index for Efficient and Effective Dense Retrieval

Haitao Li
DCST, Tsinghua University
Zhongguancun Laboratory
liht22@mails.tsinghua.edu.cn

Qingyao Ai
DCST, Tsinghua University
Zhongguancun Laboratory
aiqy@tsinghua.edu.cn

Jingtao Zhan
DCST, Tsinghua University
Zhongguancun Laboratory
chenjia0831@gmail.com

Jiaxin Mao
GSAI, Renmin University of China
maojiaxin@gmail.com

Yiqun Liu*
DCST, Tsinghua University
Zhongguancun Laboratory
yiqunliu@tsinghua.edu.cn

Zheng Liu
Huawei Poisson Lab
liuzheng107@huawei.com

Zhao Cao
Huawei Poisson Lab
caozhao1@huawei.com

ABSTRACT

Recent studies have shown that Dense Retrieval (DR) techniques can significantly improve the performance of first-stage retrieval in IR systems. Despite its empirical effectiveness, the application of DR is still limited. In contrast to statistic retrieval models that rely on highly efficient inverted index solutions, DR models build dense embeddings that are difficult to be pre-processed with most existing search indexing systems. To avoid the expensive cost of brute-force search, the Approximate Nearest Neighbor (ANN) algorithm and corresponding indexes are widely applied to speed up the inference process of DR models. Unfortunately, while ANN can improve the efficiency of DR models, it usually comes with a significant price on retrieval performance.

To solve this issue, we propose JTR, which stands for Joint optimization of TRee-based index and query encoding. Specifically, we design a new unified contrastive learning loss to train tree-based index and query encoder in an end-to-end manner. The tree-based negative sampling strategy is applied to make the tree have the maximum heap property, which supports the effectiveness of beam search well. Moreover, we treat the cluster assignment as an optimization problem to update the tree-based index that allows overlapped clustering. We evaluate JTR on numerous popular retrieval benchmarks. Experimental results show that JTR achieves better retrieval performance while retaining high system efficiency compared with widely-adopted baselines. It provides a potential solution to balance efficiency and effectiveness in neural retrieval system designs.

*Corresponding author



This work is licensed under a Creative Commons Attribution International 4.0 License.

SIGIR '23, July 23–27, 2023, Taipei, Taiwan
© 2023 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9408-6/23/07.
<https://doi.org/10.1145/3539618.3591651>

CCS CONCEPTS

• Information systems → Information retrieval query processing; Retrieval models and ranking.

KEYWORDS

Information Retrieval, Approximate Nearest Neighbor, Tree-based Index

ACM Reference Format:

Haitao Li, Qingyao Ai, Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Zheng Liu, and Zhao Cao. 2023. Constructing Tree-based Index for Efficient and Effective Dense Retrieval. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*, July 23–27, 2023, Taipei, Taiwan. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3539618.3591651>

1 INTRODUCTION

Information Retrieval (IR) system has become one of the most important tools for people to find useful information online. In practice, an industry IR system usually needs to retrieve a small set of relevant documents from millions or even billions of documents. The computational cost of brute-force search is usually unacceptable due to the system latency requirements. Therefore, indexes are often employed to achieve a fast response. One of the most well-known examples is the inverted index [24]. Unfortunately, such traditional indexing techniques cannot be applied to recently proposed Dense Retrieval (DR) models. As DR models have shown promising performance in multi-stage retrieval systems, especially in the first stage retrieval [16, 17, 20, 29, 33, 39], how to develop effective and efficient indexing technique for dense retrieval has become an important question for the research community.

Efficient DR solutions such as Approximate Nearest Neighbor (ANN) search algorithms have been widely applied in the first stage retrieval process. Popular methods include tree-based indexes [2], locality sensitive hashing (LSH) [15], product quantization (PQ) [13, 14], hierarchical navigable small-world network (HNSW) [23], etc. Since tree-based indexes can achieve sub-linear time complexity by pruning low-quality candidates, it has drawn significant attention.

Despite success in improving retrieval efficiency, most existing tree-based indexes often come with the degradation of retrieval performance. There are two main reasons: (1) The majority of existing indexes cannot benefit from supervised data because they use task-independent reconstruction error as the loss function. (2) The training objectives of the index structure and the encoder are inconsistent. In general, the optimization goal of index structure is to minimize the approximation error, while the optimization goal of the encoder is to get better retrieval performance. This inconsistency may lead to sub-optimal results.

To balance the effectiveness and efficiency of the tree-based indexes, we propose JTR, which stands for **J**oint optimization of **T**ree-based index and query encoding. To jointly optimize index structure and query encoder in an end-to-end manner, JTR drops the original “encoding-indexing” training paradigm and designs a unified contrastive learning loss. However, training tree-based indexes using contrastive learning loss is non-trivial due to the problem of differentiability. To overcome this obstacle, the tree-based index is divided into two parts: cluster node embeddings and cluster assignment. For differentiable cluster node embeddings, which are small but very critical, we design tree-based negative sampling to optimize them. For cluster assignment, an overlapped cluster method is applied to iteratively optimize it.

To verify the effectiveness of our method, we compared JTR with a wide range of ANN methods on two large publicly available datasets. The empirical experimental results show that JTR achieves better performance compared to baselines while keeping high efficiency. Through the ablation study, we have confirmed the effectiveness of the proposed strategies¹.

In short, our main contributions are as follows:

- (1) We propose a novel tree-based index for Dense Retrieval. Benefiting from the tree structure, it achieves sub-linear time complexity while still yielding promising results.
- (2) We propose a new joint optimization framework to learn both the tree-based index and query encoder. To the best of our knowledge, JTR is the first joint-optimized retrieval approach with tree-based index. This framework improves end-to-end retrieval performance through unified contrastive learning loss and tree-based negative sampling.
- (3) We relax the constraint that documents are mutually exclusive between clusters and propose an efficient clustering solution in which a document can be assigned into multiple clusters to optimize the cluster assignment. We demonstrate that overlapped cluster is beneficial for improving retrieval performance.

2 RELATED WORK

2.1 Dual Encoders

With the development of Pre-trained Language Models (PLMs), Dense Retrieval (DR) has developed rapidly in the last decade. DR typically uses dual encoders to obtain embeddings of queries and documents and utilizes inner products as similarity factors. For a

query-document pair $\langle q, d \rangle$, the primary goal is to learn a presentation function $f(\cdot)$ that maps text to higher-dimensional embeddings. The general form of a dual encoder is as follows:

$$s(q, d) = \langle f(q), f(d) \rangle \quad (1)$$

Where, \langle, \rangle refers to the inner product. $s(q, d)$ is the relevance score. To further improve the performance of dual encoders, many researchers have improved the loss function, sampling method, and so on to make them more suitable for information retrieval tasks [3, 8, 12, 18, 19, 34, 35, 37]. For example, Lee et al. [18] suggest the use of random negative sampling in large batches, which increases the difficulty of training. Zhan et al. [37] propose the dynamic hard negative sampling method, which effectively improves the performance of the DR model.

2.2 Approximate Nearest Neighbor Search

When queries and documents are encoded into embeddings, the retrieval problem can be considered as a Nearest Neighbor (NN) search problem. The simplest nearest neighbor search method is brute-force search, which becomes impractical when the corpus size explodes. Therefore, most studies use Approximate Nearest Neighbor (ANN) search. In general, there are four common methods, including tree-based methods [2], hashing [15, 28], quantization-based methods [13, 14], and graph-based methods [23]. Tree-based and graph-based methods divide embeddings into different spaces and query embeddings are retrieved only in similar spaces. Hashing and quantization-based methods compress vector representations in different ways. Both are important and are often used in combination in practice.

2.3 Joint Optimization

Existing DR systems follow the “encoding-indexing” paradigm, Inconsistent optimization objectives of the two steps during training will lead to sub-optimal performance. Previous studies have shown that the combination of optimized index construction and retrieval model training can make the index directly benefit from the annotation information and achieve better retrieval performance [9–11, 36, 38, 41]. Zhan et al [36] explore joint optimization of query encoding and product quantization, achieving state-of-the-art results. Furthermore, RepCONC [38] treats quantization as a constrained clustering process, requiring vectors to be clustered around quantized centroids. These approaches achieve good vector compression results, but the retrieval time complexity is still linear with respect to the corpus size. Tree-based methods have shown promising performance in recommender systems [40–42]. For example, TDM [41] proposes a tree-based deep recommendation model, which can be incorporated into any high-level model for retrieval. However, these methods are designed for recommendation systems and cannot be directly applied to retrieval systems.

Recently, Tay et al. [30] regarded transformer memory as the Differentiable Search Index (DSI). DSI maps documents into document IDs when indexing, and generates the corresponding docid for query when retrieving, so as to unify the training and indexing process. On this basis, Wang et al. [31] propose the Neural Corpus Indexer (NCI), which supports end-to-end document

¹Code are available at <https://github.com/CSHaitao/JTR>.

Table 1: Important notations present in this paper.

<u>Common Notation</u>	
q	a specific query
d	a specific document
n	a specific cluster node
D	the document corpus, $D = \{d_1, d_2, \dots, d_n\}$
\mathbb{D}	dimension of dense embeddings
\tilde{e}_c	embedding of cluster nodes
\tilde{e}_d	embedding of documents
S^i	documents set assigned to leaf node i
$\Phi(\cdot)$	the query encoder
<u>Tree Parameters</u>	
β	branch balance factor, representing the number of child nodes of non-leaf nodes
γ	leaf balance factor, representing the maximum number of documents contained by one leaf node
b	beam size, indicating that the top b leaf nodes will be retrieved
<u>Overlapped Cluster Parameters</u>	
L	the number of queries in the training set
K	the number of leaf nodes
N	the number of documents
M	matrix $M \in \{0, 1\}^{L \times K}$, which represents the relationship between the training queries and the leaf nodes
Y	matrix $Y \in \{0, 1\}^{L \times N}$, which indicates the relevance of the queries and documents
C	matrix $C \in \{0, 1\}^{N \times K}$, which represents the relationship between documents and leaf nodes
λ	number of repeated documents, which means that one document appears at most λ times in the leaf nodes after overlapped cluster

retrieval through a sequence-to-sequence network. These model-based indexes provide a new paradigm for learning end-to-end retrieval systems. However, they are not suitable for larger-scale web searches. And when adding or deleting documents to the system, it is difficult to update the index. Compared with them, this paper focuses on the joint optimization of index and encoders, which preserves the advantages of the original index structure.

3 THE JTR MODEL

In this section, we first introduce the preliminary and tree structure in JTR. Secondly, the end-to-end joint optimization process is described in detail. Finally, we show how to optimize the cluster assignments of documents. Table 1 shows the important notations that are present in this paper.

3.1 Preliminary

Dense Retrieval can be formalized as follows: given a query q and the corpus $D = \{d_1, d_2, \dots, d_n\}$, the model needs to retrieve the top k most relevant documents from D with q . The training set is given in the form of $\{(q_i, d_i) \dots\}$, which means that q_i and d_i are relevant.

Tree-based index clusters all documents and prunes irrelevant cluster nodes at retrieval time to improve retrieval efficiency. As depicted in Figure 1, components in JTR are as follows:

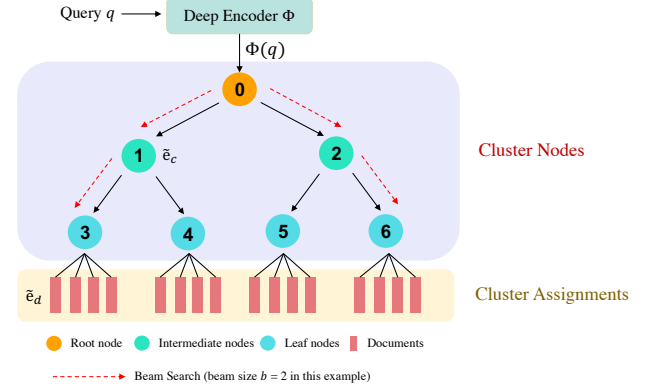


Figure 1: Illustration of the JTR tree structures. The integer represents the sequence number of the node. In this case, The tree has a depth of 3, number of clusters 4, branch balance factor $\beta = 2$, and leaf balance factor $\gamma = 4$. The beam size b is set to 2.




Figure 2: Initialization of the tree structure. The integer in nodes indicates the number of documents the node contains. In this case, the total number of documents is 10, the branch balance factor $\beta = 2$, and the leaf balance factor $\gamma = 5$. If the node contains more documents than γ , then k-means will be performed on the document embedding in the node until all nodes contain less than γ documents. The embedding of each node is initialized as the cluster centroid embedding.

- **Deep Encoder Φ** , which encodes the query q into a \mathbb{D} -dimensional dense embedding. Following the previous work, BERT [7] is employed as the encoder.
- **Cluster Nodes**, which consist of root node, intermediate nodes, and leaf nodes. The leaf nodes correspond to the fine-grained clusters, while the intermediate nodes correspond to the coarse-grained clusters. All clustering nodes are represented as trainable dense embeddings $\tilde{e}_{c_k} \in \mathbb{R}^{\mathbb{D}}$, where k denotes the k -th clustering node. More specifically, given the dense embedding of a query $\Phi(q) \in \mathbb{R}^{\mathbb{D}}$, the relevance scores of the query and the clustered nodes are calculated by $s = \tilde{e}_c^T \cdot \Phi(q)$ at each level of the tree. Based on these scores, top b leaf nodes are returned and the documents within them are further ranked. The parameter b represents the beam size. Cluster node embeddings are crucial for a tree. We will describe how to optimize them in the following section.
- **Cluster Assignment**, which represents the distribution of documents to leaf nodes. Assume there are K leaf nodes, we use

$S^i = \{d_1^i, d_2^i, \dots\}$ to denote the documents assigned to leaf node i . The initial cluster assignment is constructed by the k -means algorithm on document embeddings $\tilde{e}_{d_k} \in \mathbb{R}^D$, where k denotes the k -th document. The score $s = \tilde{e}_d^T \cdot \Phi(q)$ is applied to indicate the relevance of the query to the document. As mentioned above, we only calculate the relevance score for the documents in the top b leaf nodes. Since k -means is an unsupervised clustering method, which divides documents into mutually exclusive clusters where documents in each cluster share the same semantics. However, this does not correspond to reality. Therefore, we relax this constraint and apply an overlapped cluster approach to optimize cluster assignment. The details will be described in section 3.3.


In this paper, we define β as the branch balance factor, representing the number of child nodes of non-leaf nodes, and γ as the leaf balance factor, representing the maximum number of documents contained by one leaf node.

JTR builds the tree index by recursively using the k -means algorithm. Specifically, given the corpus to be indexed, all documents are encoded into embeddings with the trained document encoder. Then, all embeddings are clustered into β clusters by the k -means algorithm. For each node that contains more than γ documents, the k -means is applied recursively until all nodes contain less than γ documents. The embedding of each node is initialized as the cluster centroid embedding. Figure 2 illustrates the initialization of the tree structure. We can observe that the tree index has the following properties: (1) Each non-leaf node of the tree has β child nodes. The depth of the tree is influenced by branch balance factor β and leaf balance factor γ . (2) The tree may be unbalanced. In the large corpus, this unbalance is usually insignificant and does not significantly affect the effectiveness of retrieval.

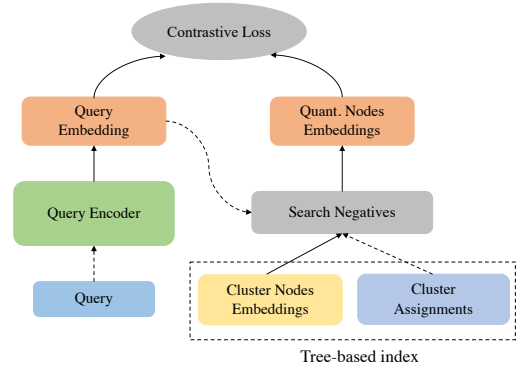
3.2 End-to-End Optimization

In this section, we introduce the end-to-end joint optimization process. In Figure 3, we compare the workflow of JTR and existing works. As shown in Figure 3(a), the existing methods follow a two-step “encoding-indexing” process. They first train the query encoder and document encoder with the ranking loss. After that, the well-trained document embeddings are used to train the tree-based index under the guidance of MSE loss. The training process of tree-based index is independent and cannot benefit from supervised data. In contrast, Figure 3(b) shows the optimization process of JTR. The joint optimization is mainly implemented by two strategies: unified contrastive learning loss and tree-based negative sampling. Next, we describe our motivation and the specific design in detail.

3.2.1 Motivation. For tree-based indexes, ensuring retrieval effectiveness and efficiency relies on two main components: pruning low-quality nodes and cluster assignment. For pruning low-quality nodes, tree-based indexes usually utilize beam search [25] to achieve efficient top b leaf nodes. Pruning the correct node at the upper levels can lead to a serious accumulation of errors. From this perspective, we argue that tree-based indexes should have the maximum heap property, which can well support the effectiveness of beam search. More specifically, the maximum heap property can be expressed as follows:



(a) Workflow of existing tree-based methods



(b) Workflow of JTR

Figure 3: Comparison of the workflow of JTR and existing methods. The solid arrows indicate that the gradient propagates backward, while the dashed arrows indicate that the gradient does not propagate.

$$p^j(n|q) = \frac{\max_{n_c \in \{n' \text{ s children node in level } j+1\}} p^{j+1}(n_c|q)}{\alpha^{(j)}} \quad (2)$$

Where $p^j(n|q)$ represents the relevant probability between query q and cluster node n in level j and $\alpha^{(j)}$ is the normalized term in level j . This formula indicates that the relevant probability of a parent node relies on the maximum relevant probability of its children nodes. In other words, given a specific query q , the parent of the optimal top node also belongs to the top node of the upper level. The maximum heap property is the basis of beam search [25].

In practice, we do not need to know the exact relevant probabilities of the cluster nodes. The relevance rank order of the nodes at each level is sufficient for accurate beam search. Therefore, we design a new contrastive learning loss to optimize cluster nodes. Moreover, the tree-based negative sample sampling technique is applied to improve the performance of JTR.

3.2.2 Unified Contrastive Learning Loss. Existing DR models usually use the contrastive learning loss function. Specifically, given a query q , let d^+ and d^- be relevant documents and negative documents. The loss function is formulated as follows:

$$L(q, d^+, d_1^-, \dots, d_n^-) = -\log \frac{\exp(s(q, d^+))}{\exp(s(q, d^+)) + \sum_{j=1}^n \exp(s(q, d_j^-))} \quad (3)$$

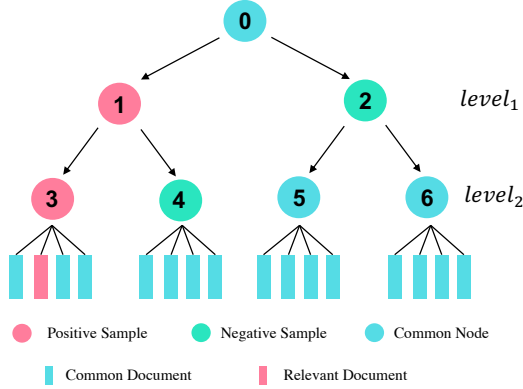


Figure 4: The sampling process of JTR. The integer represents the sequence number of the node. We select the brother nodes of positive samples as negative samples.

Where $s(\cdot)$ is the relevance score. The purpose of contrastive learning loss is to make the query closer to related documents in the embedding space compared to the irrelevant ones. However, this loss function optimizes the embeddings of queries and documents and is not applicable to existing ANN methods. In JTR, we adapt this loss to optimize the tree-based index. Specifically, given a training data (q_k, d_k) , n_k is the leaf node that contains d_k . Therefore, n_k is the positive sample of the current level, and the leaf nodes not containing d_k are negative samples. On this basis, the ancestor nodes of n_k are also treated as positive samples of the level in which they are located. To make the tree-based index meet the properties of the maximum heap, we optimize it with negative sampling at each level. Let n^+ denotes the node embedding of positive samples and n^- denotes the node embedding of negative samples. The unified contrastive learning loss is formalized as follows:

$$L(q, n^+, n_1^-, \dots, n_n^-) = -\log \frac{\exp(s(q, n^+))}{\exp(s(q, n^+)) + \sum_{j=1}^n \exp(s(q, n_j^-))} \quad (4)$$

where $s(\cdot)$ is the inner product. The unified contrastive learning loss function jointly optimizes the query encoder parameters and the embedding of clustering nodes. Since the ancestor node of the positive sample remains the positive sample during training, the tree-based index can meet the maximum heap property under the guidance of this loss function.

Training the tree-based index with the contrastive learning loss is non-trivial because the clustering assignment is not differentiable. To solve this problem, we initialize the cluster assignments using Figure 3(a) and only train the cluster node embeddings, which can benefit from supervised data directly.

3.2.3 Tree-based Negative Sampling. To further improve the performance of JTR, we design the tree-based negative sampling strategy. As shown in Figure 4, given a specific query, the leaf node with number three is a positive sample since it contains the relevant document. As mentioned above, the leaf node with number one is a positive sample at $level_1$ because it serves as the father node of the positive sample. Since the tree is constructed with the k-means algorithm, the brother nodes of each positive sample are

Algorithm 1: Tree Retrieval algorithm

Input: the trained tree \mathcal{T} , query q , beam size b , query encoder $\Phi(\cdot)$
Output: k nearest approximate candidates

- 1 Result set $A = \emptyset$, candidate set $Z = \{\text{the root node } n_1\}$.
- 2 **while** $Z \neq \emptyset$ **do**
- 3 Remove all leaf nodes from Z and insert them into A .
- 4 Calculate $s = \tilde{e}_{c_n}^T \cdot \Phi(q)$ for each remaining node $n \in Z$.
 e_{c_n} is the embedding of current node.
- 5 According to the s , top $(b - \text{len}(A))$ nodes in Z are selected to form the set I . There are no leaf nodes in I .
- 6 Update the Z : $Z = \{\text{children nodes of } n | n \in I\}$.
- 7 **end**
- 8 Compute $s = \tilde{e}_d^T \cdot \Phi(q)$ for documents d contained by nodes in A to get top k candidates.

considered to be closer to the positive sample in the embedding space. Hence, we select them as negative samples for training. As shown in Figure 3(b), the query embedding is fed into the ‘‘Search Negatives’’ module and brother nodes of the positive sample are returned under the current parameter.

3.3 Optimized Overlapped Cluster

As mentioned above, cluster assignment plays an essential role in the tree-based index. However, k-means is an unsupervised clustering method, which divides documents into mutually exclusive clusters, and documents in each cluster should share the same semantics. We observe in practice that the semantics of a document are complex and multi-topic. Putting each document in a cluster can limit the performance of the tree-based index. To solve this problem, we propose overlapped clusters to further optimize cluster assignment.

The overlapped cluster has been studied in the field of unsupervised learning [4, 21, 32]. Inspired by Liu et al. [21], we formulate the cluster assignment problem in information retrieval as an optimization problem. Suppose that there are L queries and N documents in the training corpus and all documents have been clustered and put into the K leaf nodes of the tree index.

First, we define the ground truth matrix $\mathbb{Y} \in \{0, 1\}^{L \times N}$ as:

$$\mathbb{Y}_{i,j} = \begin{cases} 1 & \text{if } doc_j \text{ is relevant to } query_i \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

\mathbb{Y} describes the relevance of the training queries and the documents and is the best result the model is trying to achieve.

Then, we input the training queries into the tree-based index to get the leaf nodes associated with them. We define the relationship between queries and leaf nodes as matrix $\mathbb{M} \in \{0, 1\}^{L \times K}$ and the original cluster assignment as matrix $\mathbb{C} \in \{0, 1\}^{N \times K}$. It is worth noting that the initialized cluster assignment is obtained from Figure 3(a).

$$\mathbb{M}_{i,j} = \begin{cases} 1 & \text{if } query_i \text{ matches } leaf_j \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

$$\mathbb{C}_{i,j} = \begin{cases} 1 & \text{if } doc_i \text{ matches } leaf_j \\ 0 & \text{otherwise} \end{cases} \quad (7)$$




Figure 5: Illustration of the optimized overlapped cluster. In this case, there are 3 queries, 4 leaf nodes, and 5 documents. The $q_i \setminus l_i \setminus d_i$ represent the i -th query \leaf node \document respectively. We set the number of overlapped clustering $\lambda = 2$. The values in the red boxes are identified by the $Proj(\cdot)$ function. In practice, if a document has the same value for two leaves in C^* , the $Proj(\cdot)$ function prefers to keep the document in its original leaf.

Then we can represent the relationship between queries and documents of JTR as:

$$\hat{Y} = \text{Binary}(\mathbb{M} \times \mathbb{C}^T) \quad (8)$$

where $\text{Binary}(A) = I_A$ is the element-wise indicator function. When the element $a > 0$, $I_a = 1$, which ensures $\hat{Y} \in \{0, 1\}^{L \times N}$. \times stands for matrix cross product. Then, the performance of JTR can be expressed as the intersection of \hat{Y} and Y . The formula is as follows:

$$\text{Recall} = |\hat{Y} \cap Y| \quad (9)$$

where $|\cdot|$ returns the number of non-zero elements in the matrix. Since \hat{Y} and Y are binary matrices, We can obtain the following formula:

$$\text{Recall} = |\hat{Y} \cap Y| = \text{Tr}(\hat{Y}^T \times Y) \quad (10)$$

where $\text{Tr}(\cdot)$ returns the trace of the matrix. It is found that the performance of JTR is proportional to the trace of the matrix. In order to maximize the recall rate, we formulate the cluster assignment problem as follows:

$$\underset{S_1, S_2, \dots, S_k}{\text{maximize}} \text{Tr}(\hat{Y}^T \times \text{Binary}(\mathbb{M} \times \mathbb{C}^T)) \quad (11)$$

$$\text{s.t.} \sum_{i=1}^K I_{d \in S_i} \leq \lambda, \forall d \in \{d_1, \dots, d_N\} \quad (12)$$

$$\bigcup_{i=1}^K S_i = \{d_1, \dots, d_N\} \quad (13)$$

where S_i represents the i th cluster assignment. I is the indicator function. Equation 12 guarantees that any document d appears at most λ times in all clusters. Equation 13 ensures that the clusters include all documents.

This optimization problem is an NP-complete problem. Following Liu et al. [21], we approximate the objective function with a continuous, ReLU-like function.

$$\text{Binary}(\mathbb{M} \times \mathbb{C}^T) \approx \max(\mathbb{M} \times \mathbb{C}^T, 0) = \mathbb{M} \times \mathbb{C}^T \quad (14)$$

In the optimization function, the cluster information is hidden in the matrix \mathbb{C} . $\text{Tr}(\hat{Y}^T \times \mathbb{M} \times \mathbb{C}^T)$ is linear with matrix \mathbb{C} , so the optimal solution of \mathbb{C} is the projection of $\hat{Y}^T \times \mathbb{M}$ onto the constraint set. Therefore, the optimization problem has a closed-form solution:

$$C^* = \text{Proj}(\hat{Y}^T \times \mathbb{M}) \quad (15)$$

where the $\text{Proj}(\cdot)$ operator selects the top λ elements for each line of the matrix. In the field of Information Retrieval, there exists the problem of data sparsity. Only a small set of the documents have the corresponding training queries, resulting in a lot of rows in \hat{Y}^T with all entities being zeros. To solve this problem, we use the trained DR model to retrieve the top- k documents of each training query to construct matrix $\bar{Y} \in \{0, 1\}^{L \times N}$. In other words, if doc_j is the top k document recalled by $query_i$ using the DR model, then $\bar{Y}_{i,j} = 1$. Hence, every line of \bar{Y} has k nonzero entries. In short, the final solution is:

$$C^* = \text{Proj}(\bar{Y}^T \times \mathbb{M}) \quad (16)$$

The proposed solution is based on the intuition that two documents are more likely to cluster into the same class if they are related to the same query. It is worth noting that only the cluster assignment is changed in this process, the structure of the tree and the cluster node embeddings do not change.

After the C^* is determined, We re-optimize the JTR as described in Section 3.2 to accommodate the new cluster assignment. Figure 5 shows the process of optimized overlapped cluster. We acknowledge that there exists a few documents that are not relevant to any training query. These documents are retained in the original clusters.

3.4 Tree Retrieval

For a specific query, the tree retrieval process is described in algorithm 1. At each level of the tree, we select the top nodes and their child nodes as candidates for the next level. It saves retrieval time by pruning less relevant nodes. In the tree structure of JTR, the leaf nodes are not always at the same level. Therefore, the A set is used to preserve the encountered leaf nodes. In line 5 of Algorithm 1, we select top $b - \text{len}(A)$ nodes in Z each time to create I . Since the leaf nodes in Z are removed and inserted into A in line 3, there are no leaf nodes in I .

In the JTR, the retrieval process is hierarchical and top-to-down. If the tree has K leaf nodes, N documents and beam size is b and branch balance factor is β , the time complexity of retrieving leaf nodes is $O(\beta * b * \log K)$. For retrieval in set A , the maximum number of documents in A is $b * \gamma$, where γ is leaf balance factor. Since γ can be approximated as N/K , the time complexity of this part is $O(b * N/K)$. In summary, the retrieval time complexity of JTR is

Table 2: Results on the MS MARCO dataset. AQT stands for Average Query processing Time, which is measured by averaging time over each query of the MS MARCO Dev set with a single thread and a single batch on the CPU. */ denotes that JTR performs significantly better than baselines at $p < 0.05/0.01$ level using the two-tailed pairwise t-test. The best method in each column is marked in bold.**

Model	MARCO Passage		DL19 Passage	DL20 Passage	AQT	MARCO Doc		DL19 Doc	DL20 Doc	AQT
	MRR@100	R@100	NDCG@10	NDCG@10	ms	MRR@100	R@100	NDCG@10	NDCG@10	ms
IVFFlat	0.311*	0.778	0.580*	0.615	75	0.349**	0.839	0.572	0.527*	63
PQ	0.289**	0.717**	0.448**	0.546**	149	0.290**	0.788**	0.498**	0.458**	58
IVFPQ	0.252**	0.653**	0.532**	0.540**	15	0.279**	0.695**	0.465**	0.423**	10
JPQ	0.306**	0.832	0.611	0.607	152	0.347**	0.889	0.575	0.536	55
Annoy	0.144**	0.263**	0.361**	0.385**	132	0.148**	0.253**	0.504**	0.463**	57
FALCONN	0.295**	0.719**	0.554**	0.532**	23	0.321**	0.746**	0.496**	0.460**	9
FLANN	0.271**	0.629**	0.551**	0.578**	18	0.294**	0.649**	0.418**	0.368**	5
IMI	0.314	0.697**	0.542**	0.565**	37	0.348**	0.788**	0.535*	0.468**	26
HNSW	0.289**	0.732**	0.546**	0.559**	11	0.334**	0.783**	0.503**	0.507**	5
JTR	0.318	0.778	0.610	0.632	30	0.364	0.848	0.590	0.565	18

$O(\beta * b * \log K) + O(b * N/K)$. The overall time complexity is below the linear time complexity, which greatly improves the retrieval efficiency.

4 EXPERIMENT SETTINGS

In this section, we introduce our experimental settings, including implementation details, datasets and metrics, baselines.

4.1 Datasets and Metrics

The experiments are conducted on the dataset MS MARCO [27], which is a large-scale ad-hoc retrieval benchmark. It contains two large-scale tasks: Document Retrieval and Passage Retrieval. Passage Retrieval has a corpus of 8.8M passages, 0.5M training queries, and 7k development queries. Document Retrieval has a corpus of 3.3M documents, 0.4M training queries, and 5k development queries.

We conduct our experiments on two tasks from MS MARCO Dev, TREC2019 DL [5] and TREC2020 DL [6]. The MS MARCO Dev is extracted from Bing’s search logs, and each query is marked as relevant to a few documents. TREC2019 DL and TREC2020 DL are collections that contain extensively annotated documents for each query, i.e., using four-level annotation criteria. We use $R@100$ to evaluate the recall performance of different methods. $MRR@100$ and $NDCG@10$ are applied to measure ranking performance.

4.2 Baselines

We select several state-of-the-art ANN indexes as our baselines. Here are more details:

IVFFlat [24]: IVFFlat is the classic inverted index. When applied to dense retrieval, IVF defines $nlist$ clusters in the vector space, and only the top $nprobe$ clusters close to the query embedding are retrieved. We set $nprobe$ to be the same as beam size b .

PQ [14]: PQ is an ANN index based on product quantization. We set the number of embedding segments to 32 and the number of coding bits to 8.

IVFPQ [14]: IVFPQ is one of the fastest indexes available, which combines “inverted index + product quantization”. The parameter settings are the same as IVFFlat and PQ.

JPQ [36]: JPQ implements the joint optimization of query encoder and product quantization. We set the number of segments into which each embedding will be divided as 24.

Annoy [2]: Annoy improves retrieval efficiency by building a binary tree. We set the number of trees to 100.

FALCONN [28]: FALCONN is an optimized LSH method, which supports hyperplane LSH and cross polyhedron LSH. We use the recommended parameter settings as adopted in the corresponding paper.

FLANN [26]: FLANN is one of the most complete ANN open-source libraries, including liner, kdtree, kmeans tree, and other index methods. We use automatic tuning to get the best parameters.

IMI [1]: IMI is a multilevel inverted index, which combines product quantization and inverted index to bring very good search performance. The number of bits is set to 12.

HNSW [23]: HNSW is a typical and widely-used graph-based index. We set the number of links to 8 and ef-construction to 100.

For a fair comparison, all ANN indexes are operated on the document embeddings formed by STAR [37]. For IVFFlat, PQ, IVFPQ, HNSW and IMI, we implement them with the Faiss². For Annoy³, FALCONN⁴ and FLANN⁵, we use the official library to implement them.

4.3 Implementation Details

We implemented JTR using PyTorch and Transformers⁶. Initial embeddings were obtained using STAR [37] for all documents. We also load the checkpoint of STAR to warm up the query encoder. In our experiment, The default γ is 1000 and β is 10. The dimension of embeddings is 768. For the training setup, we use the AdamW [22] optimizer with a batch size of 32. The learning rate is set to 5×10^{-6} . In the overlapped cluster of the tree, we use the top 100 documents obtained from ADORE-STAR [37] to build matrix \tilde{Y} . All experiments are evaluated on the workstation with Xeon Gold 5218 CPUs and RTX 3090 GPUs.


²<https://github.com/facebookresearch/faiss>

³<https://github.com/spotify/annoy>


⁴<https://github.com/FALCONN-LIB/FALCONN>

⁵<https://github.com/flann-lib/flann>

⁶<https://huggingface.co/docs/transformers/index>



(a) MRR curves



(b) Recall curves

Figure 6: Trade-off curves for different ANN methods. AQT stands for Average Query processing Time. Bottom and right is better.

5 EXPERIMENT RESULTS

5.1 Comparison with ANN Methods

The performance comparisons between SAILER and baselines are shown in Table 2. We derive the following observations from the experiment results.

- PQ and JPQ are ANN methods based on product quantization, which have a linear time complexity with respect to the corpus size. As expected, their average query processing time is the longest among all baselines. IVFPQ is an inverted index based on product quantization, which improves retrieval efficiency but damages retrieval effectiveness.
- Benefiting from a different index structure, the average query processing time of IVFPQ, FALCONN, FLANN, and HNSW is lower than that of JTR. However, they suffered a loss in effectiveness.
- Annoy and FLANN are existing tree-based indexes. Compared with them, JTR achieved the best results. To the best of our knowledge, JTR is the best tree-based index available, which shows the great potential of tree-bases indexes for dense retrieval.
- Overall, JTR performs significantly better than the baseline method on most measures. In terms of effectiveness, JTR achieves the best MRR/NDCG and the second best recall performance (i.e., R@100) on all datasets. Compared to the baseline with the best recall (i.e., JPQ), JTR achieves 3 to 5 times latency speedup. We can conclude that JTR has a very competitive effectiveness compared to other ANN indexes.

To further analyze the ability of different indexes to balance efficiency and effectiveness, we plot AQT-MRR curves with varying parameters. As shown in Figure 6, we have the following findings:




Figure 7: R@100 different λ ranging from 1 to 5 and b ranging from 10 to 50.




Figure 8: MRR@100 versus different λ ranging from 1 to 5 and b ranging from 10 to 50.




Figure 9: AQT versus different λ ranging from 1 to 5 and b ranging from 10 to 50. AQT stands for Average Query processing Time.

- IVFPQ and FALCONN are limited by resources, which makes it difficult to achieve high retrieval performance.
- When our requirement on recall is extremely high (e.g., larger than 0.8), brute-force search could be more efficient than any indexing technique, which makes brute-force search-based algorithms like JPQ more efficient than JTR. After all, any indexing process would add more time complexity if we eventually need to check all documents in the corpus. However, such high requirements on recall are not common in web search or other similar retrieval tasks.
- To sum up, JTR has a very outstanding effectiveness-efficiency trade-off ability, which can always achieve better results with a short retrieval latency. JTR provides a potential solution to balance efficiency and effectiveness for dense retrieval.

5.2 Hyperparameter Sensitivity

5.2.1 Hyperparameters for Tree Retrieval. In this section, we investigate the impact of the overlap number λ and the beam size b . All experiments were conducted on the MS MARCO Doc Dev dataset.

Table 3: Ablation study on MS MARCO Dev Doc dataset.

Model	MARCO Dev Doc		
	MRR@100	R@100	AQT(ms)
IVFFlat	0.310	0.714	24
Tree	0.256	0.556	5
+Joint Optimization	0.296	0.640	5
+Reorganize clusters	0.303	0.678	5
+Overlapped clustering	0.327	0.743	8

We set $\lambda = 1, 2, 3, 4, 5$ and $b = 10, 20, 30, 40, 50$ to observe the retrieval performance of JTR. Figures 7, 8 and 9 show the changes in retrieval quality and retrieval efficiency with hyperparameters. We can get the following observations:

- Both retrieval quality and latency increase with λ . When λ increases, the leaf node contains more documents. The number of candidate documents goes up which leads to better results and higher latency.
- The slope of MRR and R curve decreases with the increase of λ , but the slope of AQT remains unchanged. That means the gain from overlapped cluster is getting smaller.
- When we fix λ , it is found that the gains of MRR and R decreased with the increase of b . The reason may be that the most relevant documents have already clustered in the top leaf nodes.
- In Figure 9, we note that the gap between curves increases when λ is larger, which shows that a large λ could lead to larger computation costs on each leaf node.
- Overall, the increase of b and λ can lead to better performance and longer latency. Therefore, we recommend the moderate value for λ and b to maximize the trade-off between effectiveness and efficiency.


5.2.2 Hyperparameter for Overlapped Cluster. In this section, we fix the model as ADORE-STAR and change k from 25 to 150, where k denotes the number of documents recalled with the dense retrieval model. Figure 10 presents the retrieval performance on each k and b value. Obviously, the larger k , the more information contained in \bar{Y} , and the better the final performance. When k is larger than 100, the gain becomes smaller. Therefore, we set $k = 100$ in our experiments, which provides sufficient information.

5.3 Ablation Study


In this section, we conduct ablation studies on JTR to explore the importance of different components. We use IVFFlat as the baseline. IVFFlat forms as many clusters as the leaf nodes of the tree. The n_{probe} of IVFFlat and the beam size of JTR are both set to 10. We use the following four model variants:

- **Tree:** The STAR model is used to obtain the document embeddings, and then we build the tree index with MSE loss.
- **+Joint Optimization:** Joint optimization of query encoders and tree-based indexes using the unified contrastive learning loss function.
- **+Reorganize Clusters:** Update the clustering of documents, which means the number of overlapped clustering $\lambda = 1$.
- **+Overlapped Clustering:** Documents can be re-occurring in the clustering nodes. For convenience, the number of overlapped clustering $\lambda = 2$.

Table 3 shows the $MRR@100$ and $R@100$ on the development set of the MS MARCO Document Ranking task. As the results show,



(a) MRR curves



(b) Recall curves

Figure 10: MRR@100 and R@100 versus different k ranging from 25 to 150 and b ranging from 10 to 50.

the tree structure significantly reduces the retrieval latency. All of Joint Optimization, Reorganize Clusters and Overlapped Clustering improve the performance of JTR. Tree-based indexes benefit from supervised data directly by Joint Optimization. Reorganize clusters makes the distribution of embeddings more reasonable. Overlapped cluster mines different semantic information in documents. This result demonstrates the effectiveness of our method.

6 CONCLUSION

To improve the efficiency of the DR models while ensuring the effectiveness, we propose JTR, which jointly optimizes tree-based index and query encoder in an end-to-end manner. To achieve this goal, we carefully design a unified contrastive learning loss and tree-based negative sampling strategy. Through the above strategies, the constructed index tree possess the maximum heap property which easily supports beam search. Moreover, for cluster assignment, which is not differentiable w.r.t. contrastive learning loss, we introduce overlapped cluster optimization. We further conducted extensive experiments on several popular retrieval benchmarks. Experimental results show that our approach achieves competitive results compared with widely-adopted baselines. Ablation studies demonstrate the effectiveness of our strategies. Unfortunately, since different indexes have varying degrees of code optimization, we do not report the memory case in our paper. In the future, we will try to jointly optimize PQ and tree-based index to achieve the “effectiveness-efficiency-memory” tradeoff.

ACKNOWLEDGMENTS

This work is supported by the Natural Science Foundation of China (62002194), Tsinghua University Guoqiang Research Institute, Tsinghua-Tencent Tiangong Institute for Intelligent Computing and the Quan Cheng Laboratory.

REFERENCES

- [1] Artem Babenko and Victor Lempitsky. 2014. The inverted multi-index. *IEEE transactions on pattern analysis and machine intelligence* 37, 6 (2014), 1247–1260.
- [2] Erik Bernhardsson. 2017. Annoy: approximate nearest neighbors in C++/Python optimized for memory usage and loading/saving to disk. *GitHub* <https://github.com/spotify/annoy> (2017).
- [3] Jia Chen, Haitao Li, Weihang Su, Qingyao Ai, and Yiqun Liu. 2023. THUIR at WSDM Cup 2023 Task 1: Unbiased Learning to Rank. *arXiv preprint arXiv:2304.12650* (2023).
- [4] Guillaume Cleuziou. 2008. An extended version of the k-means method for overlapping clustering. In *2008 19th International Conference on Pattern Recognition*. IEEE, 1–4.
- [5] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M Voorhees. 2020. Overview of the TREC 2019 deep learning track. *arXiv preprint arXiv:2003.07820* (2020).
- [6] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Fernando Campos, and Ellen M. Voorhees. 2020. Overview of the TREC 2020 Deep Learning Track. *ArXiv abs/2003.07820* (2020).
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [8] Qian Dong, Yiding Liu, Suqi Cheng, Shuaiqiang Wang, Zhicong Cheng, Shuzi Niu, and Dawei Yin. 2022. Incorporating Explicit Knowledge in Pre-trained Language Models for Passage Re-ranking. *arXiv preprint arXiv:2204.11673* (2022).
- [9] Yan Fang, Jingtao Zhan, Yiqun Liu, Jiaxin Mao, Min Zhang, and Shaoping Ma. 2022. Joint Optimization of Multi-vector Representation with Product Quantization. In *Natural Language Processing and Chinese Computing: 11th CCF International Conference, NLPCC 2022, Guilin, China, September 24–25, 2022, Proceedings, Part I*. Springer, 631–642.
- [10] Chao Feng, Wuchao Li, Defu Lian, Zheng Liu, and Enhong Chen. 2022. Recommender Forest for Efficient Retrieval. *Advances in Neural Information Processing Systems* 35 (2022), 38912–38924.
- [11] Chao Feng, Defu Lian, Zheng Liu, Xing Xie, Le Wu, and Enhong Chen. 2022. Forest-based Deep Recommender. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 523–532.
- [12] Luyu Gao, Zhuyun Dai, Tongfei Chen, Zhen Fan, Benjamin Van Durme, and Jamie Callan. 2021. Complementary lexical retrieval model with semantic residual embeddings. In *European Conference on Information Retrieval*. Springer, 146–160.
- [13] Tiezheng Ge, Kaiming He, Qifa Ke, and Jian Sun. 2013. Optimized product quantization. *IEEE transactions on pattern analysis and machine intelligence* 36, 4 (2013), 744–755.
- [14] Herve Jegou, Matthijs Douze, and Cordelia Schmid. 2010. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence* 33, 1 (2010), 117–128.
- [15] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data* 7, 3 (2019), 535–547.
- [16] Vladimir Karpukhin, Barlas Ögüz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2004.04906* (2020).
- [17] Omar Khattab and Matei Zaharia. 2020. 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*. 39–48.
- [18] Jinhyuk Lee, Minjoon Seo, Hannaneh Hajishirzi, and Jaewoo Kang. 2019. Contextualized sparse representations for real-time open-domain question answering. *arXiv preprint arXiv:1911.02896* (2019).
- [19] Haitao Li, Jia Chen, Weihang Su, Qingyao Ai, and Yiqun Liu. 2023. Towards Better Web Search Performance: Pre-training, Fine-tuning and Learning to Rank. *arXiv preprint arXiv:2303.04710* (2023).
- [20] Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2020. Distilling dense representations for ranking using tightly-coupled teachers. *arXiv preprint arXiv:2010.11386* (2020).
- [21] Xuanqing Liu, Wei-Cheng Chang, Hsiang-Fu Yu, Cho-Jui Hsieh, and Inderjit Dhillon. 2021. Label disentanglement in partition-based extreme multilabel classification. *Advances in Neural Information Processing Systems* 34 (2021), 15359–15369.
- [22] Ilya Loshchilov and Frank Hutter. 2018. Fixing weight decay regularization in adam. (2018).
- [23] Yu A Malkov and Dmitry A Yashunin. 2018. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence* 42, 4 (2018), 824–836.
- [24] Antonio Mallia, Omar Khattab, Torsten Suel, and Nicola Tonello. 2021. Learning passage impacts for inverted indexes. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1723–1727.
- [25] Clara Meister, Tim Vieira, and Ryan Cotterell. 2020. Best-first beam search. *Transactions of the Association for Computational Linguistics* 8 (2020), 795–809.
- [26] Marius Muja and David Lowe. 2009. Flann-fast library for approximate nearest neighbors user manual. *Computer Science Department, University of British Columbia, Vancouver, BC, Canada* 5 (2009).
- [27] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A human generated machine reading comprehension dataset. In *CoCo@ NIPS*.
- [28] Ninh Pham and Tao Liu. 2022. Falconn++: A Locality-sensitive Filtering Approach for Approximate Nearest Neighbor Search. *arXiv preprint arXiv:2206.01382* (2022).
- [29] Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2020. RocketQA: An optimized training approach to dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2010.08191* (2020).
- [30] Yi Tay, Vinh Q Tran, Mostafa Dehghani, Jianmo Ni, Dara Bahri, Harsh Mehta, Zhen Qin, Kai Hui, Zhe Zhao, Jai Gupta, et al. 2022. Transformer memory as a differentiable search index. *arXiv preprint arXiv:2202.06991* (2022).
- [31] Yujing Wang, Yingyan Hou, Haonan Wang, Ziming Miao, Shibin Wu, Hao Sun, Qi Chen, Yuqing Xia, Chengmin Chi, Guoshuai Zhao, et al. 2022. A Neural Corpus Indexer for Document Retrieval. *arXiv preprint arXiv:2206.02743* (2022).
- [32] Joyce Jiyoung Whang, Inderjit S Dhillon, and David F Gleich. 2015. Non-exhaustive, overlapping k-means. In *Proceedings of the 2015 SIAM international conference on data mining*. SIAM, 936–944.
- [33] Xiaohui Xie, Qian Dong, Bingning Wang, Feiyang Lv, Ting Yao, Weinan Gan, Zhijing Wu, Xiangsheng Li, Haitao Li, Yiqun Liu, et al. 2023. T2Ranking: A large-scale Chinese Benchmark for Passage Ranking. *arXiv preprint arXiv:2304.03679* (2023).
- [34] Shenghao Yang, Haitao Li, Zhumin Chu, Jingtao Zhan, Yiqun Liu, Min Zhang, and Shaoping Ma. 2022. THUIR at the NTCIR-16 WWW-4 Task. *Proceedings of NTCIR-16, to appear* (2022).
- [35] Shenghao Yang, Yiqun Liu, Xiaohui Xie, Min Zhang, and Shaoping Ma. 2023. Enhance Performance of Ad-hoc Search via Prompt Learning. In *Information Retrieval: 28th China Conference, CCIR 2022, Chongqing, China, September 16–18, 2022, Revised Selected Papers*. Springer, 28–39.
- [36] Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Jiafeng Guo, Min Zhang, and Shaoping Ma. 2021. Jointly optimizing query encoder and product quantization to improve retrieval performance. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 2487–2496.
- [37] Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Jiafeng Guo, Min Zhang, and Shaoping Ma. 2021. Optimizing dense retrieval model training with hard negatives. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1503–1512.
- [38] Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Jiafeng Guo, Min Zhang, and Shaoping Ma. 2022. 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*. 1328–1336.
- [39] Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Min Zhang, and Shaoping Ma. 2020. ReBERT: Contextualized text embeddings for first-stage retrieval. *arXiv preprint arXiv:2006.15498* (2020).
- [40] Han Zhu, Daqing Chang, Ziru Xu, Pengye Zhang, Xiang Li, Jie He, Han Li, Jian Xu, and Kun Gai. 2019. Joint optimization of tree-based index and deep model for recommender systems. *Advances in Neural Information Processing Systems* 32 (2019).
- [41] Han Zhu, Xiang Li, Pengye Zhang, Guozheng Li, Jie He, Han Li, and Kun Gai. 2018. Learning tree-based deep model for recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1079–1088.
- [42] Jingwei Zhuo, Ziru Xu, Wei Dai, Han Zhu, Han Li, Jian Xu, and Kun Gai. 2020. Learning optimal tree models under beam search. In *International Conference on Machine Learning*. PMLR, 11650–11659.