

DOGR: Leveraging Document-Oriented Contrastive Learning in Generative Retrieval

Penghao Lu, Xin Dong*, Yuansheng Zhou, Lei Cheng, Chuan Yuan, Linjian Mo

Ant Group

{lupenghao.lph,zhaoxin.dx,zhouyuansheng.zys,lei.chenglei,yuanzheng.xy,linyi01}@antgroup.com

Abstract

Generative retrieval constitutes an innovative approach in information retrieval, leveraging generative language models (LM) to generate a ranked list of document identifiers (docid) for a given query. It simplifies the retrieval pipeline by replacing the large external index with model parameters. However, existing works merely learned the relationship between queries and document identifiers, which is unable to directly represent the relevance between queries and documents. To address the above problem, we propose a novel and general generative retrieval framework, namely *Leveraging Document-Oriented Contrastive Learning in Generative Retrieval* (DOGR), which leverages contrastive learning to improve generative retrieval tasks. It adopts a two-stage learning strategy that captures the relationship between queries and documents comprehensively through direct interactions. Furthermore, negative sampling methods and corresponding contrastive learning objectives are implemented to enhance the learning of semantic representations, thereby promoting a thorough comprehension of the relationship between queries and documents. Experimental results demonstrate that DOGR achieves state-of-the-art performance compared to existing generative retrieval methods on two public benchmark datasets. Further experiments have shown that our framework is generally effective for common identifier construction techniques.

Introduction

Information retrieval aims to satisfy user queries by retrieving documents. With the rapid increase in the volume of information, the importance of retrieval technology is becoming gradually prominent. In recent years, a novel retrieval method known as generative retrieval has emerged. Unlike traditional retrieval methods (sparse retrieval and dense retrieval), generative retrieval assigns an identifier to each document and generates relevant document identifiers end-to-end through language models. In addition, generative retrieval can better utilize the capabilities of the large language models (LLMs), especially the knowledge learned during the pre-training phase.

Most generative retrieval methods (Cao et al. 2021; Tay et al. 2022; Wang et al. 2022a) first define an identifier for

each document and then treated document retrieval as a standard sequence-to-sequence task. Specifically, this task learns the generation of identifiers through the interaction between queries and identifiers in an end-to-end manner. However, the identifier length is often much smaller than documents, identifiers may experience a certain degree of information loss compared to the complete document. Existing generative retrieval methods mainly concentrate on learning the relationship between the queries and the identifiers, thus they are unable to capture the complete semantic relationships between queries and documents.

There are some studies introducing additional optimization objectives to comprehensively characterize the relationship between queries and documents. A previous approach (Li et al. 2024) proposed a document-oriented margin-based loss based on generation probability to optimize the document ranking. Another work (Lee et al. 2023) modeled query-document relevance by separately learning the generation probabilities of the query and document to the target identifier. However, these work only focused on the intermediate target of learning identifiers and lack the interaction between query and documents. This could potentially affect the retrieval performance. Since the model can only learn the nuanced interactions through the identifier, it is insufficient to model the relevance between queries and documents by solely relying on the generation probability of identifiers.

To address the above problems, we introduce a novel and general framework, namely *Leveraging Document-Oriented Contrastive Learning in Generative Retrieval* (DOGR), which captures the relationship between queries and documents comprehensively through their direct interactions. DOGR proposed a two-stage learning strategy consisting of an identifier generation stage and a document ranking stage. First, in the identifier generation stage, we train an encoder-decoder architecture model based on the generation loss to learn the connections between the document identifier and the query. Next, in the document ranking stage, we fine-tune the model with contrastive learning objectives to directly learn the correlation between queries and documents. We enhance document representation through the elaborate design of two negative sampling methods. Specifically, DOGR explicitly leverages *prefix-oriented negative sampling* to learn semantic representations of documents with the same identifier prefix and *retrieval-augmented negative sampling* to

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

comprehend the representation of documents with distinct identifier prefixes. Drawing on the above mentioned negative samples, two corresponding contrastive losses are utilized to enhance the semantic similarity between query and its positive document, aiming to achieve analogous representations for both by increasing direct interaction. During the inference phase, we not only consider the generation probability of document identifiers, but also leverage the semantic score to jointly model the relevance between the query and documents. By means of ranking the documents corresponding to the generated identifiers, the performance of generative retrieval is further enhanced.

In summary, our main contributions are as follows:

- We propose **DOGR**, a novel and general framework which adopts contrastive learning to enhance generative retrieval. It incorporates a two-stage learning strategy that jointly models the relevance of query-document by combining the generation probability of identifiers with document semantics.
- According to the characteristics of generative retrieval, we introduced two negative sampling methods and corresponding contrastive learning objectives to help the model fully capture the correlation between query and documents through direct interaction.
- Comprehensive experiments show that **DOGR** outperforms existing state-of-the-art generative retrieval methods on two benchmark datasets. Further experiments have demonstrated that our method is effective for several common identifier construction approaches.

Related Work

Document Retrieval

Traditional document retrieval methods mainly include sparse retrieval and dense retrieval. Sparse retrieval (classic methods such as BM25 (Robertson, Zaragoza et al. 2009) and TF-IDF (Robertson and Walker 1997)) has been widely used in industry due to its high efficiency. However, sparse retrieval relies on lexical overlap and finds it challenging to match semantically related but literally different documents. Dense retrieval involves representing queries and documents as dense vectors (Lee, Chang, and Toutanova 2019), and calculating the similarity between the query and the document using inner product or cosine distance. Benefiting from the continuous development of pre-trained language models, the performance of dense retrieval is also constantly improving. Contrastive learning (Khosla et al. 2020; Gao, Yao, and Chen 2021) and hard negative sampling techniques (Xiong et al. 2021; Karpukhin et al. 2020; Ni et al. 2021, 2022) have become another important driving force in the development of dense retrieval. Therefore, it has gradually become another mainstream retrieval method in industrial implementation. Although dense retrieval has achieved great success, it suffers from issues such as embedding space bottleneck and a lack of fine-grained interactions between embeddings.

Generative Retrieval

Different from traditional document retrieval methods, generative retrieval uses only the language model to directly

generate the identifiers of relevant documents for a given query. GENRE (Cao et al. 2021) is the first step in performing information retrieval tasks using an autoregressive language model. It implemented constrained beam search based on a prefix tree to generate entity titles within a given candidate set. Since document identifiers largely determines the effectiveness of generative retrieval, many studies have explored the construction methods of identifiers. In general, identifiers can be categorized into lexical and numerical types. Lexical identifiers allow for explicit expression including n-grams (Bevilacqua et al. 2022; Wang et al. 2023), titles (Cao et al. 2021; Li et al. 2023), etc, while numerical identifiers (Tay et al. 2022; Wang et al. 2022a,b) mainly capture similarities and differences among documents through clustering methods. However, these works focused on the learning process of mapping queries to identifiers, which cannot directly represent the relevance between queries and documents.

Recently, there have been studies that introduce additional optimization objectives to help characterize the relationship between queries and documents. LTRGR (Li et al. 2024) proposed a margin-based ranking loss based on generation probability from the document level to improve document ranking. GLEN (Lee et al. 2023) utilized identifiers as a bridge to model query-document relevance by separately learning the generation probabilities of both the query and the document in relation to the target identifier. Although they open up new ideas for generative retrieval, they still only utilized the generation probability of document identifiers, which has a lack of direct interaction between the query and the document. To address the aforementioned issues, we propose a two-stage learning strategy to learn the generation probabilities at the identifier level and the semantic representation at the document level. Through direct interaction between queries and documents, we enhance generative retrieval through contrastive learning objectives based on negative sampling methods.

The Proposed Scheme

In this work, we propose a two-stage learning strategy that enhances query representations to better align with document representations for generative retrieval. Figure 1 shows an illustration of the core procedure of DOGR, including an identifier-level generation stage and a document-level ranking stage. First, in the *Identifier Generation Stage*, DOGR trains an encoder-decoder architecture model to learn the connections between the keyword-based document identifier with query.

Then, in the *Document Ranking Stage*, DOGR fine-tunes the trained model by introducing contrastive learning to depict the relationship between the query and the complete document. We improve document representation by elaborately designing two negative sampling methods. In particular, we use prefix-oriented negative sampling to capture the semantic representations of documents sharing the same identifier prefix. On the other hand, we propose retrieval-augmented negative sampling to better understand the representation of documents with different identifier prefixes.

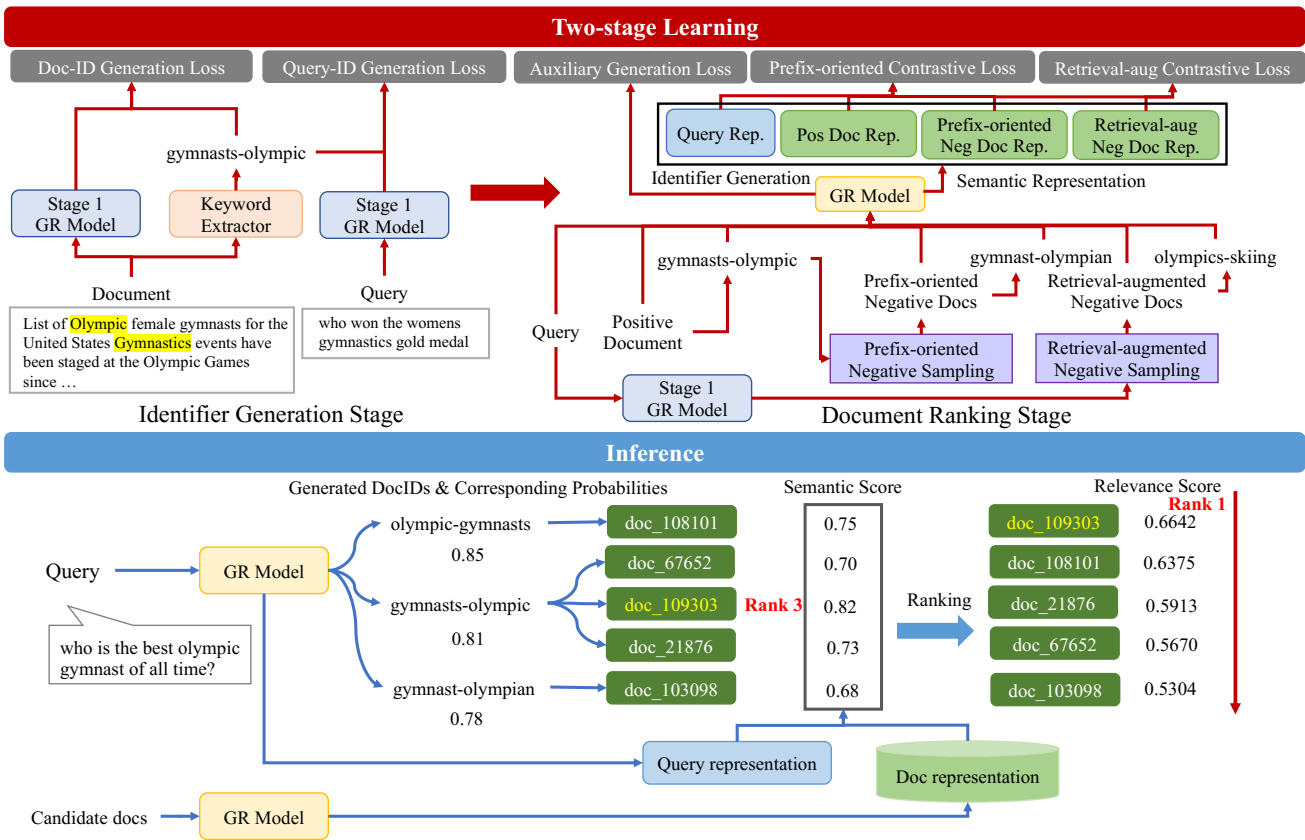


Figure 1: An overview of the framework of DOGR. The training phase can be divided into two stages. The first is the identifier generation stage, which is used to learn the relationship between text and document identifiers. The second is the document ranking stage, which employs contrastive learning to learn semantic representation for document ranking. In the inference phase, DOGR first generates identifiers for the given query, then the documents are ranked according to the relevance scores derived from generation probabilities of identifiers and the semantic scores.

Informed by the aforementioned negative samples, two related contrastive losses are applied to maximize the semantic score between the query and the positive document. This approach strives to attain comparable representations for both.

In the inference phase, both the generation probability of document identifiers and the semantic score are taking into account to collaboratively model the relevance between query and documents. Through ranking the documents corresponding to the generated identifiers, the performance of generative retrieval is further enhanced.

Identifier Generation Stage

Keyword-Based Lexical Identifier. Existing generative retrieval methods require constructing identifier for each document. Since keywords are important carriers of document information, we extract keywords as the document identifier. Specifically, for each document d , pre-trained KeyBert (Grootendorst 2020) is used to extract the top keywords and sort the keywords in descending order of importance. After encoding the keyword sequence with the tokenizer, the first n tokens are selected as the document identifier $z^d = \{z_0^d, z_1^d, \dots, z_{n-1}^d\}$. As a result, each doc-

ument d has its corresponding identifier z^d . While a document identifier z can correspond to a collection of documents $\{d_i | z^{d_i} = z\}$, since similar documents may have the same identifier. In addition to keyword-based identifier, our framework is applicable to common identifier types including both lexical and numerical type, and its generalization will be validated in the experiment section.

Pointwise Generation Task. Similar to the previous identifier generation tasks, the training objective in this stage is a standard sequence-to-sequence task. This task aims at learning the relationship between query and document identifiers.

For each sample (q, d) in training data, we adopt sequence-to-sequence cross-entropy loss with teacher forcing to maximize the log-likelihoods of $q \rightarrow z^d$. It is far from sufficient to simply model the generation task of query to target document identifier. This is because the query cannot represent the entirety of the document and it cannot assist the model in fully learning the relationship between the document and the target identifier, which may lead to poor robustness. Thus we utilize the same loss function to maximize the loglikelihoods of $d \rightarrow z^d$. Different from the query input, document segments are randomly sampled as input

to learn the relationship between document and its identifier more comprehensively. It was observed that training the document and query tasks together in a multi-task setting leads to better performance in generating identifiers than training these tasks separately. Accordingly, the generation loss is formulated as follows,

$$L_1 = - \sum_{t=1}^n \log P(z_t^d | q, z_{<t}^d) - \sum_{t=1}^n \log P(z_t^d | d, z_{<t}^d), \quad (1)$$

where z_t^d is the token at t -th step of z^d .

Document Ranking Stage

Since the identifier generation task mainly focuses on learning the relationship between queries and identifiers, which has the limited ability to comprehensively understand the relationship between queries and documents. Therefore, the purpose of the document ranking stage is to learn the semantic representation of document through its direct interaction with the query. In light of the characteristics of generative retrieval, two negative sampling strategies and corresponding contrastive learning objectives are implemented to optimize document representation.

Prefix-oriented Negative Sampling. We introduce contrastive loss to help the model to represent queries as close to positive documents d^+ and far from negative documents d^- . Then, we propose prefix-oriented negative sampling to sample negative documents with the same identifier prefix within the batch. Specifically, we first select documents with the same first n tokens as z^{d^+} (excluding d^+ itself). Next, documents that have the same first $n - 1$ tokens as d^+ are fetched. We repeatedly reduce the prefix length until $|\mathcal{N}_z|$ negative samples are obtained. Since identifiers with the same prefix tend to have similar document information, this negative sampling method can help the model better distinguish difficult samples. Based on the semantic representation, we define the semantic score between the query and document as follows,

$$s(q, d) = h_q \cdot h_d^T, \quad (2)$$

where $h_x \in \mathbb{R}^m$ is the average pooling of the hidden representations at each position of the encoder for a given input x . According to the relationship between the positive and negative samples mentioned above, the prefix-oriented contrastive learning loss function can be formulated:

$$L_{c_i} = - \log \frac{\exp(s(q, d^+)/\tau)}{\exp(s(q, d^+)/\tau) + \sum_{d^- \in \mathcal{N}_z} \exp(s(q, d^-)/\tau)},$$

where τ is the temperature and \mathcal{N}_z is the set of negative documents obtained via prefix-oriented negative sampling.

Retrieval-augmented Negative Sampling. Then, we introduce retrieval-augmented negative sampling to sample negative documents with distinct identifier prefixes. To be specific, it samples negative documents by generating high-ranking identifiers for the given query. For each query q during training, the top k document identifiers $\{z_1, \dots, z_k\}$ are retrieved based on the model trained in the first stage. The document collection corresponding to these identifiers

is $\bigcup_{j=1}^k \{d_i | z^{d_i} = z_j\}$. In these documents, the documents that have the same identifier as the positive documents are removed, as they have already been sampled in the prefix-oriented negative sampling. Then the negative documents \mathcal{N}_q are sampled from the remain documents, i.e. $\bigcup_{j=1}^k \{d_i | z^{d_i} = z_j\} / \{d_i | z^{d_i} = z_j\}$, where $|\mathcal{N}_q| < k$. Thanks to the negative samples provided by the first stage learning, a retrieval-augmented contrastive learning loss function is formulated as follows,

$$L_{c_q} = - \log \frac{\exp(s(q, d^+)/\tau)}{\exp(s(q, d^+)/\tau) + \sum_{d^- \in \mathcal{N}_q} \exp(s(q, d^-)/\tau)}.$$

Training. We also incorporate the generation loss into the document ranking stage. It ensures that the model’s ability to generate document identifiers does not deteriorate due to the introduction of ranking learning objectives. Moreover, we use an additional loss to learn the corresponding document identifiers for the negative documents obtained by the retrieval-augmented negative sampling. We add this loss to the generation loss L_1 as the auxiliary generation loss L_g , which is calculated as follows,

$$L_g = L_1 + \left(\sum_{d^- \in \mathcal{N}_q} \sum_{t=1}^n - \log P(z_t^{d^-} | d^-, z_{<t}^{d^-}) \right). \quad (3)$$

Therefore, the final loss function is defined as,

$$L = L_{c_i} + L_{c_q} + \lambda_g \cdot L_g, \quad (4)$$

where λ_g is the hyperparameter used to control the importance of generation loss, which balances the generation objective and contrastive learning objectives.

Inference

During the inference phase, we integrate the identifier generation probabilities with the semantic representation to retrieve and rank documents.

First, for a given query q , we apply constrained beam search (Cao et al. 2021) to generate a ranked list of top k document identifiers $\{z_1, z_2, \dots, z_k\}$ with the corresponding generation probabilities $P(z_1|q), P(z_2|q), \dots, P(z_k|q)$. Next, we calculate the query representation h_q and document representations $\{h_d | d \in \mathcal{D}_q\}$, where \mathcal{D}_q is the corresponding document collection to the top k document identifiers. Based on the generation probability of the identifier and overall semantic information, the fusion relevance score between query and document is defined as follows:

$$rel(q, d) = P(z^d|q) \cdot s(q, d). \quad (5)$$

Finally, we rank the obtained document collection \mathcal{D}_q based on their relevance scores and the top k documents with the highest relevance scores are selected as the final retrieval results for the given query. Meanwhile, the document ranking issue under identifier conflicts can be simply addressed.

Experiments

Datasets

Natural Questions(NQ320k) (Kwiatkowski et al. 2019) contains 320k training data (relevant query-document pairs),

Model	NQ320K(7830)			Seen test(6075)			Unseen test(1755)		
	R@1	R@10	MRR@100	R@1	R@10	MRR@100	R@1	R@10	MRR@100
Sparse & Dense Retrieval									
BM25(Robertson, Zaragoza et al. 2009)	29.7	60.3	40.2	29.1	59.8	39.5	32.3	61.9	42.7
DocT5Query(Nogueira, Lin, and Epistemic 2019)	38.0	69.3	48.9	35.1	68.3	46.7	48.5	72.9	57.0
DPR(Karpukhin et al. 2020)	50.2	77.7	59.9	50.2	78.7	60.2	50.0	74.2	58.8
ANCE(Xiong et al. 2021)	50.2	78.5	60.2	49.7	79.2	60.1	52.0	75.9	60.5
SentenceT5(Ni et al. 2021)	53.6	83.0	64.1	53.4	83.9	63.8	56.5	79.5	64.9
GTR-base(Ni et al. 2022)	56.0	84.4	66.2	54.4	84.7	65.3	61.9	83.2	69.6
Generative Retrieval									
GENRE(Cao et al. 2021)	55.2	67.3	59.9	69.5	83.7	75.0	6.0	10.4	7.8
DSI(Tay et al. 2022)	55.2	67.4	59.6	69.7	83.6	74.7	1.3	7.2	3.5
SEAL(Bevilacqua et al. 2022)	59.9	81.2	67.7	-	-	-	-	-	-
DSI-QG(Zhuang et al. 2022)	63.1	80.7	69.5	68.0	85.0	39.5	32.3	61.9	42.7
NCI(Wang et al. 2022a)	66.4	85.7	73.6	69.8	88.5	76.8	54.5	75.9	62.4
GENRET(Wang et al. 2022b)	68.1	88.8	<u>75.9</u>	70.2	<u>90.3</u>	77.7	62.5	83.6	70.4
LTRGR(Li et al. 2024)	67.5	86.2	74.8	70.2	88.7	77.3	58.1	77.3	66.2
GLEN(Lee et al. 2023)	<u>69.1</u>	86.0	75.4	<u>72.5</u>	88.9	<u>78.5</u>	57.6	75.9	63.9
DOGR(Ours)	70.2	89.1	76.8	72.6	90.5	78.9	<u>61.9</u>	84.3	<u>69.5</u>

Table 1: Performance comparison for the proposed framework and baseline methods for NQ320k. Bolded numbers are the best performance of each column and the second best method is underlined. The number in parentheses indicates the number of queries. We refer to the results of baselines reported by (Lee et al. 2023). All the numbers in the table are percentage numbers with ‘%’ omitted. Results not available are denoted as ‘-’.

100k documents, and 7,830 test queries, and is widely used in existing generative retrieval methods. We follow the same setup as previous work (Lee et al. 2023) and split the test set into two subsets: seen test and unseen test. The seen test consists of queries with target documents that are covered in the training data, while the documents in the unseen test are not covered.

MS MARCO passage ranking (MS MARCO) (Nguyen et al. 2016) is a large-scale benchmark dataset that includes 8.8 million passages collected from Bing search results and 1 million real-world queries, with the test set containing 6980 queries.

Metrics

According to the evaluation methods of previous works, we adopt the widely accepted metrics, namely Recall@1, Recall@10, and Mean Reciprocal Rank (MRR), to comprehensively assess retrieval performance.

Baselines

We compare with several types of baseline models, including traditional document retrieval methods and generative retrieval models.

Traditional Document Retrieval Methods In traditional document retrieval methods, we consider several commonly used highlighting models, covering both sparse retrieval (BM25 (Robertson, Zaragoza et al. 2009), DocT5Query (Nogueira, Lin, and Epistemic 2019)) and dense retrieval techniques (ANCE (Xiong et al. 2021), DPR (Karpukhin et al. 2020), SentenceT5 (Ni et al. 2021), GTR-base (Ni et al. 2022)).

Generative Retrieval Methods Furthermore, we compare DOGR with several generative retrieval baselines including GENRE (Cao et al. 2021), DSI (Tay et al. 2022), SEAL (Bevilacqua et al. 2022), DSI-QG (Zhuang et al. 2022), NCI (Wang et al. 2022a), GENRET (Wang et al. 2022b), and two methods (GLEN (Lee et al. 2023) and LTRGR (Li et al. 2024)) introducing additional objectives to model query-document relationship. For the sake of experimental fairness, we replaced the backbone of LTRGR from BART (Lewis et al. 2020) to T5-base (Raffel et al. 2020) for training.

Implementation Details

Keyword-based Lexical Identifier For each document, we first extract the top 20 diverse keywords based on KeyBert using maximal marginal relevance, while the diversity parameter is set to 0.5 and n-gram is set to 1. After experimenting with different identifier token length, we set the token length of document identifiers to 8 and 12 for NQ320k and MS MARCO respectively. We will analyze in detail the impact of identifier length on retrieval performance below.

Query Generation It has been proved that augmented queries related to the document can effectively help the model capture the relationship between query and documents. Therefore, we follow the previous work (Wang et al. 2022a) and augment the query through Document as Query and DocT5Query.

Training and Inference To ensure a fair comparison with previous work, we use T5-base (Raffel et al. 2020) as the generative language model backbone. In the training phase, batch size is set to 256 and 32, and the model is optimized for up to 3M and 1M steps using the Adam optimizer

Model	MS MARCO Dev(6980) MRR@10
BM25	18.4
DocT5Query	27.2
GTR-base(Ni et al. 2022)	34.8
DSI(Tay et al. 2022)	3.1
DSI-QG(Zhuang et al. 2022)	11.8
NCI(Wang et al. 2022a)	17.4
LTRGR(Li et al. 2024)	<u>21.6</u>
GLEN(Lee et al. 2023)	20.1
DOGR	22.5

Table 2: Performance comparison for the proposed framework and baseline methods for MS MARCO. Bolded numbers are the best performance of each column and the second best method is underlined. All the numbers are percentage numbers with ‘%’ omitted.

with learning rates $5e-5$ for identifier generation stage and document-level ranking stage, respectively. The number of negatives from retrieval-augmented negative sampling is set to 4 per query, while the prefix-oriented negative sampling employs in-batch negatives. In the document ranking stage, τ is set to 0.5 as the temperature parameter for contrastive learning, and λ_g is set to 0.1 to balance the generative task and the contrastive learning task. In the inference phase, we use the beam search with constrained decoding during inference and set the beam size to 100. Our experimental code is implemented on Python 3.8 using transformers 4.37.0, while experiments are conducted on 6 NVIDIA A100 GPUs with 80 GB of memory.

Experimental Results

Main Results

Evaluation on NQ320k. Table 1 shows the retrieval performance on NQ320k. The main observations are as follows:

- DOGR surpasses traditional document retrieval methods and outperforms existing generative retrieval methods. Specifically, DOGR outperforms the best competitive generative retrieval model on the full test set of NQ320k by +1.6%, +0.3%, and +1.2% on Recall@1, Recall@10 and MRR@100, respectively.
- DOGR not only outperforms existing baseline models on the full and seen test sets, but also achieves competitive performance with the previous best model on the unseen test set, which shows that our method has good robustness.
- Our method exhibits higher performance than other methods adopting additional objectives to model query-document relationship in generative retrieval. This demonstrates the effectiveness of enhancing generative retrieval via document-oriented contrastive learning.

Evaluation on MS MARCO. Table 2 presents the retrieval performance on MS MARCO. DOGR improves MRR@10 by 4.2% compared to the best competitive generative retrieval method and by 22.3% compared to BM25(Robertson, Zaragoza et al. 2009). Existing generative retrieval methods

DocID Type	Enhance Type	NQ320K(7830)		
		R@1	R@10	MRR@100
keyword	DOGR	70.2	89.1	76.8
	LTRGR	69.2	88.3	76.3
	GLEN	68.7	88.0	75.9
	-	68.2	87.7	75.5
n -gram	DOGR	67.2	86.9	74.3
	LTRGR	66.2	86.2	73.5
	GLEN	65.4	85.8	72.8
	-	65.2	85.4	72.8
first n -tokens	DOGR	66.3	86.7	73.9
	LTRGR	65.4	85.7	73.0
	GLEN	64.5	85.1	72.3
	-	64.3	84.8	71.8
numerical	DOGR	67.8	86.8	74.9
	LTRGR	67.1	86.1	74.3
	GLEN	66.2	85.9	73.5
	-	66.4	85.7	73.6

Table 3: Retrieval performance of different methods on various identifier types. All the numbers are percentage numbers with ‘%’ omitted. Models with only generation task are denoted as ‘-’.

still struggle to memory knowledge of the document corpus solely through identifiers, making it difficult to operate effectively in large-scale corpora. In contrast, DOGR enhances generative retrieval through contrastive learning and document semantics, thus it is successfully performed in large-scale corpora.

In-depth Analysis

Generalization of DOGR. We verify the generalization of DOGR, i.e. whether it can provide a certain improvement for common identifier construction methods including both lexical and numerical type in generative retrieval. Table 3 presents the improvements of DOGR for common identifiers on NQ320k, along with a comparison to two other competitive methods. We find that DOGR significantly improves the performance of lexical and numerical identifiers. For keyword DocID, we improve Recall@1 from 68.2 to 70.2 (a relative increase of 2.9%). Meanwhile, for other DocID types, DOGR achieves at least a 2.1%, 1.2% and 1.7% enhancement in Recall@1, Recall@10, and MMR@100 respectively, and demonstrates the generalization of the framework proposed. Additionally, DOGR outperforms two other methods that introduce additional objectives to capture query-document correlation in generative retrieval, achieving at least a 1.0% improvement in R@1. This may be attributed to integrating the generation probability of identifiers with semantic representations, allowing for a more comprehensive modeling of the query-document relevance.

Impact of Identifier length. Since our framework allows different documents to share the same identifier, it may lead to identifier conflicts. Table 4 shows the impact of identifier length on performance of NQ320k, where the conflict rate is defined as the ratio of the number of documents with identifier conflicts to the total number of documents. We find

Identifier length	Conflict Rate	NQ320K(7830)		
		R@1	R@10	MRR@100
$l = 3$	46.9	66.7	87.2	75.3
$l = 4$	28.6	67.9	88.0	75.7
$l = 6$	8.9	68.9	88.7	76.2
$l = 8$	3.0	70.2	89.1	76.8
$l = 10$	1.1	69.8	88.6	75.8
$l = 12$	0.4	69.6	88.3	75.5

Table 4: Impact of Identifier length on NQ320k. All the numbers are percentage numbers with ‘%’ omitted.

that when the identifier length is short, there are significant identifier conflict issues, which reduces the learning ability of the identifier generation task while increasing the ranking difficulty at the document level. Additionally, the shorter the identifier, the more limited its ability to convey information about the document. These two reasons may lead to a decline in model performance when the identifiers is short. Specifically, when the identifier length decreases from 8 to 3, Recall@1 drops by 5.0%. On the other hand, as the length of the identifiers increases, the issue of identifier conflicts gradually diminishes, and the performance improves. Once the length reaches 8, the Recall@1 does not improve further. One possible reason is that the model needs to memorize the order of keywords as the length increases, this may lead to a decrease in robustness. Moreover, the training and inference speeds also decrease with increasing identifier length.

Ablation Study. We also conducted detailed experiments to investigate the roles of different strategies in the training and inference phase:

- “w/o doc segment sampling”: The document segment sampling is removed in identifier generation stage.
- “w/o prefix-ori contrastive / retrieval-aug contrastive / auxiliary generation loss”: The prefix-oriented contrastive / retrieval-augmented contrastive / auxiliary generation loss is removed in document ranking stage, respectively.
- “w/o semantic score”: Semantic score is removed during the inference phase. This means that ranking is performed only based on generation probability, and in particular, random ranking is used when encountering identifier conflicts.
- “w/o generation probability”: The generation probability is removed during the inference phase. This means that after generating identifiers, the corresponding documents are ranked only based on the semantic score.
- “w/o fusion rank”: We adopt the same strategy as (Lee et al. 2023) instead of ranking by the fusion relevance score, i.e. overall ranking by generation probability and only rank by semantic score when encountering identifier conflicts.

The effectiveness of various strategies in **DOGR** on NQ320k is demonstrated in Table 5, which provides the following insights:

- Sampling of document segments can help the model capture the relationship between documents and identifiers

Model	NQ320K(7830)		
	R@1	R@10	MRR@100
DOGR	70.2	89.1	76.8
w/o doc segment sampling	68.6	88.2	75.9
w/o prefix-ori contrastive loss	69.0	88.4	76.0
w/o retrieval-aug contrastive loss	69.3	88.5	76.2
w/o auxiliary generation loss	65.8	84.3	72.9
w/o semantic score	68.2	87.7	75.5
w/o generation probability	63.3	87.6	72.0
w/o fusion rank	69.9	88.8	76.5

Table 5: Ablation study of **DOGR** with different strategy on NQ320k. All the numbers are percentage numbers with ‘%’ omitted.

more comprehensively, enhancing robustness and yielding a 2.3% improvement in Recall@1.

- Two contrastive learning tasks extract negative samples from different perspectives, resulting in gains of 1.7% and 1.3% in Recall@1, respectively. This indicates that contrastive learning can effectively enhance generative retrieval by document ranking. Additionally, we found that without the assistance of generative loss, the generative task experiences a decline, with Recall@1 dropping by 6.3%, highlighting the necessity of generative loss.
- During the inference phase, relying solely on the generation probability cannot address identifier collision issues, which results in a 2.8% drop in Recall@1. On the other hand, discarding the generation probability in ranking causes a 9.8% drop in Recall@1, which is mainly due to the lack of interaction between query and documents. Moreover, integrating both the generation probability and semantic information yields a further 0.4% improvement in Recall@1 compared to ranking solely based on semantic information under identifier collisions.

Conclusion

In this paper, we propose **DOGR**, a novel and general framework that enhancing generative retrieval via document-oriented contrastive learning. It employs a two-stage learning strategy that simultaneously models the relevance between the query and documents by integrating the generation probability of identifiers with document semantic representation. Moreover, two negative sampling methods and contrastive learning objectives are employed to comprehensively learn the relationship between the query and the document through their direct interaction. Experimental results demonstrate that **DOGR** achieves state-of-the-art among generative retrieval methods on two benchmark datasets and brings improvements on common identifier construction methods. In the future, we will continue to explore how to end-to-end integrate information from documents to enhance the retrieval capabilities of generative retrieval in large-scale corpora.

References

- Bevilacqua, M.; Ottaviano, G.; Lewis, P. S. H.; Yih, S.; Riedel, S.; and Petroni, F. 2022. Autoregressive Search Engines: Generating Substrings as Document Identifiers. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Cao, N. D.; Izacard, G.; Riedel, S.; and Petroni, F. 2021. Autoregressive Entity Retrieval. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Gao, T.; Yao, X.; and Chen, D. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In Moens, M.; Huang, X.; Specia, L.; and Yih, S. W., eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, 6894–6910. Association for Computational Linguistics.
- Grootendorst, M. 2020. KeyBERT: Minimal keyword extraction with BERT.
- Karpukhin, V.; Oguz, B.; Min, S.; Lewis, P. S. H.; Wu, L.; Edunov, S.; Chen, D.; and Yih, W. 2020. Dense Passage Retrieval for Open-Domain Question Answering. In Webber, B.; Cohn, T.; He, Y.; and Liu, Y., eds., *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, 6769–6781. Association for Computational Linguistics.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised Contrastive Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Kwiatkowski, T.; Palomaki, J.; Redfield, O.; Collins, M.; Parikh, A. P.; Alberti, C.; Epstein, D.; Polosukhin, I.; Devlin, J.; Lee, K.; Toutanova, K.; Jones, L.; Kelcey, M.; Chang, M.; Dai, A. M.; Uszkoreit, J.; Le, Q.; and Petrov, S. 2019. Natural Questions: a Benchmark for Question Answering Research. *Trans. Assoc. Comput. Linguistics*, 7: 452–466.
- Lee, K.; Chang, M.; and Toutanova, K. 2019. Latent Retrieval for Weakly Supervised Open Domain Question Answering. In Korhonen, A.; Traum, D. R.; and Màrquez, L., eds., *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, 6086–6096. Association for Computational Linguistics.
- Lee, S.; Choi, M.; Lee, J.; and X. 2023. GLEN: Generative Retrieval via Lexical Index Learning. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, 7693–7704. Association for Computational Linguistics.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J. R., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, 7871–7880. Association for Computational Linguistics.
- Li, Y.; Yang, N.; Wang, L.; Wei, F.; and Li, W. 2023. Multi-view Identifiers Enhanced Generative Retrieval. In Rogers, A.; Boyd-Graber, J. L.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, 6636–6648. Association for Computational Linguistics.
- Li, Y.; Yang, N.; Wang, L.; Wei, F.; and Li, W. 2024. Learning to Rank in Generative Retrieval. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2024, February 20-27, 2024, Vancouver, Canada*, 8716–8723. AAAI Press.
- Nguyen, T.; Rosenberg, M.; Song, X.; Gao, J.; Tiwary, S.; Majumder, R.; and Deng, L. 2016. MS MARCO: A Human Generated MACHINE READING COMPREHENSION DATASET. *CoRR*, abs/1611.09268.
- Ni, J.; Ábrego, G. H.; Constant, N.; Ma, J.; Hall, K. B.; Cer, D.; and Yang, Y. 2021. Sentence-T5: Scalable Sentence Encoders from Pre-trained Text-to-Text Models. *CoRR*, abs/2108.08877.
- Ni, J.; Qu, C.; Lu, J.; Dai, Z.; Ábrego, G. H.; Ma, J.; Zhao, V. Y.; Luan, Y.; Hall, K. B.; Chang, M.; and Yang, Y. 2022. Large Dual Encoders Are Generalizable Retrievers. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, 9844–9855. Association for Computational Linguistics.
- Nogueira, R.; Lin, J.; and Epistemic, A. 2019. From doc2query to docTTTTTquery. *Online preprint*, 6(2).
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *J. Mach. Learn. Res.*, 21: 140:1–140:67.
- Robertson, S.; Zaragoza, H.; et al. 2009. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4): 333–389.
- Robertson, S. E.; and Walker, S. 1997. On relevance weights with little relevance information. In *Proceedings of the 20th annual international ACM SIGIR conference on Research and development in information retrieval*, 16–24.
- Tay, Y.; Tran, V.; Dehghani, M.; Ni, J.; Bahri, D.; Mehta, H.; Qin, Z.; Hui, K.; Zhao, Z.; Gupta, J. P.; Schuster, T.; Cohen, W. W.; and Metzler, D. 2022. Transformer Memory as a Differentiable Search Index. In Koyejo, S.; Mohamed, S.;

Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.

Wang, Y.; Hou, Y.; Wang, H.; Miao, Z.; Wu, S.; Chen, Q.; Xia, Y.; Chi, C.; Zhao, G.; Liu, Z.; Xie, X.; Sun, H.; Deng, W.; Zhang, Q.; and Yang, M. 2022a. A Neural Corpus Indexer for Document Retrieval. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.

Wang, Y.; Hou, Y.; Wang, H.; Miao, Z.; Wu, S.; Chen, Q.; Xia, Y.; Chi, C.; Zhao, G.; Liu, Z.; Xie, X.; Sun, H.; Deng, W.; Zhang, Q.; and Yang, M. 2022b. A Neural Corpus Indexer for Document Retrieval. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.

Wang, Z.; Zhou, Y.; Tu, Y.; and Dou, Z. 2023. NOVO: Learnable and Interpretable Document Identifiers for Model-Based IR. In Frommholz, I.; Hopfgartner, F.; Lee, M.; Oakes, M.; Lalmas, M.; Zhang, M.; and Santos, R. L. T., eds., *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM 2023, Birmingham, United Kingdom, October 21-25, 2023*, 2656–2665. ACM.

Xiong, L.; Xiong, C.; Li, Y.; Tang, K.; Liu, J.; Bennett, P. N.; Ahmed, J.; and Overwijk, A. 2021. Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

Zhuang, S.; Ren, H.; Shou, L.; Pei, J.; Gong, M.; Zuccon, G.; and Jiang, D. 2022. Bridging the Gap Between Indexing and Retrieval for Differentiable Search Index with Query Generation. *CoRR*, abs/2206.10128.