

One-Step Retrieval Framework for Real-Time Sponsored Search Ads Using Hierarchical Text Representations

Tongtong Liu, Renyu Zhang, Jiayu Ding, Hongchao Guo,
Xintao Yang, He Wei, Zhaoyu Li, Haiyang Wu

Tencent Inc.

{uniqueliu,yurenzhang,garryding}@tencent.com
{peterrguo,xintaoyang,whywei,joezyli,gavinwu}@tencent.com

Abstract

Traditional retrieval systems typically use multi-stage cascading architectures (MCA), where each module is optimized independently, leading to inconsistent objectives and the premature elimination of high-potential candidates. Recent LLM-based generation methods offer end-to-end solutions but use discrete semantic identifiers (SIDs) to retrieve ads, which are not learned by the base LLM and require memorization of numerous SID-to-ad mappings during SFT, suffering from limited generalization to unseen ads, high maintenance and update costs. The one-to-one mapping between SIDs and advertisements leads to inefficient decoding. Moreover, these methods rely on a small reward model (e.g. pctr) for relevance and ranking, limiting the LLM’s ability to fully assess ads’ commercial value. To address these challenges, we propose **A** uNified **G**eneration-discriminative-ranking **r**eaL time **r**Etrieval (**ANGLE**) framework. ANGLE uses LLM-generated hierarchical textual representations, which consist of commercial intent that provide high-level overviews and ad abstract that deliver fine-grained details. Additionally, ANGLE integrates retrieval, relevance, and ranking directly within a single LLM, enabling precise and efficient ranking of ads by leveraging the full capabilities of the LLM. We applied ANGLE to the real-world search scenarios, achieving a 1.81% increase in consumption and a 2.16% increase in gross merchandise volume (GMV). We also conducted offline evaluations of ANGLE and seven baselines, with ANGLE outperforming all across key metrics such as HR and ACR.

1 Introduction

Ad retrieval systems efficiently retrieve and display high-value ads that match users’ search intent from large-scale corpora. Retrieval, relevance, pre-ranking and ranking constitute four essential components of such a system, addressing user search intents, advertiser delivery requirements, and plat-

form revenue objectives. Traditional retrieval systems typically use multi-stage cascading architectures (MCA), in which candidates are sequentially passed through retrieval, relevance, pre-ranking, and ranking component, with each stage taking the previous stage’s output as input, and finally selecting the top-k ads for display. However, this approach faces several drawbacks. Firstly, errors or omissions in earlier stages are propagated to subsequent ones, limiting the overall quality of the final results. Secondly, the modules are often trained independently, lacking joint optimization, which makes it difficult to achieve a globally optimal solution. Thirdly, processing a large number of candidates at each stage leads to high computational costs and slower response times. In recent years, several studies (Fei et al., 2021; Huang et al., 2023; Zhang et al., 2023) have been conducted to address the aforementioned challenges of MCA. However, these methods still fail to break through the framework of MCA.

With the rapid development of LLM, generative approaches have recently been explored for directly producing candidates for users. Most existing methods (Chen et al., 2025a; Deng et al., 2025) employ discrete semantic identifiers (SIDs), typically generated via RQ-VAE (Lee et al., 2022) to represent candidate docs/ads. Upon receiving a user query, the LLM generates a set of SIDs conditioned on the query, which are subsequently used to retrieve relevant ads that are presented to the user. To further refine the recommendation process, lightweight reward models—such as click-through rate (CTR) predictors—provide feedback on the relevance and effectiveness of the displayed ads.

However, existing work suffers from two main **limitations**. Firstly, ads are represented by discrete SIDs (Tay et al., 2022; Fu et al., 2025) that rely on unique discrete IDs to characterize each ad. Consequently, the LLM is required to learn the mapping between SIDs and ads from scratch, which limits its

ability to fully leverage its inherent text generation capabilities and reduces the overall interpretability of the model. Furthermore, these mappings exhibit poor generalization, necessitating the creation of new mappings for previously unseen docs/ads. Secondly, current methods depend on one or more lightweight models(e.g., predicted click-through rate predictors) as reward models to guide the optimization of the LLM, which constrains the full potential and capabilities of the LLM. Therefore, developing more effective advertising representations and reward mechanisms is essential for fully leveraging the capabilities of LLMs within the one-step generation framework.

To address these challenges, we propose **A** **u**nified **G**eneration-discriminative-ranking **r**eal-time **r**etrieval framework named **ANGLE**, employs a hierarchical semantic text —intent and abstract pair—to represent advertisements (ads), integrating relevance and ranking objectives directly into the LLM training process to achieve end-to-end retrieval. We incorporate both discriminative and ranking capabilities into the base LLM, resulting in the generation-discriminative-ranking LLM (GDR-LLM). Specifically, ANGLE leverages base LLM to generate intents (broad, intent-expressive text) and abstracts (precise, ad-derived text) for ads, subsequently constructing an inverted index mapping intent-abstract (int-abs) pairs to the associated ads. Upon receiving a user query, the GDR-LLM first infers the query intent, then generates the corresponding ad abstract aligned with this intent, and retrieves the relevant ads via the aforementioned inverted index. These retrieved ads are subsequently sent directly to the ranking stage for processing, after which they are displayed to users.

The primary **innovation** of this work lies in utilizing hierarchical semantic text as an intermediate representation and designing a GDR-LLM to enable end-to-end retrieval from users to ads. Specifically, we represent user search intents and advertising needs as int-abs pairs, where both components are concise texts conveying rich semantic information. This semantic text representation allows ANGLE to fully utilize the world knowledge embedded in LLM, as it aligns well with their generative mechanisms. Furthermore, our approach demonstrates strong generalization capabilities: for new ads, the int-abs pair can be efficiently updated by generating it directly from ad textual content. This representation also offers high interpretability, which facilitates clear understanding and analysis.

In addition, GDR-LLM incorporates raw relevance evaluation and commercial ranking data into its supervised fine-tuning process, thereby generating highly commercial advertisements that align with relevance criteria. Unlike recommendation scenarios, search scenarios require much higher precision and relevance. With enhanced relevance and ranking capabilities, GDR-LLM enables the simplification of the MCA pipeline: retrieved ads are sent directly to the re-ranking stage and subsequently displayed to users, achieving end-to-end retrieval.

The main contributions of our work are summarized as follows:

- **Framework:** We propose ANGLE, an end-to-end real-time framework validated at tens-of-millions scale in production.
- **Representation:** We introduce hierarchical semantic text for ad representation, addressing interpretability and generalization issues of SIDs.
- **Model:** We design a generation-discriminative-ranking LLM that integrates discrimination and ranking capabilities within a unified framework, fully leveraging the capabilities of LLM.
- **Evaluation:** We conducted both online A/B testing and offline experiments to evaluate the effectiveness of the ANGLE. The A/B testing yielded a 1.81% increase in consumption, a 2.16% increase in GMV, and a 1.5% increase in clicks. Additionally, offline comparisons against seven baselines show that ANGLE achieved the best performance across HR, MAP, and ACR metrics.

2 Related Work

We review related work from three perspectives: traditional multi-stage cascading architectures, generative retrieval approaches, and document representation methods in LLM-based retrieval systems.

2.1 Multi-stage Cascading Architectures

Multi-stage cascading architectures (MCA) (Covington et al., 2016; Wang et al., 2011) are widely used in industrial advertising systems to balance computational efficiency and ranking accuracy. This approach divides retrieval pipeline into sequential stages—typically retrieval(Freymuth et al., 2025), relevance(Dey et al., 2025), pre-ranking(Zhao et al., 2025), and ranking(Huang et al., 2025)—where the output candidates from each stage serve as the input for the next. However, our proposed ANGLE framework eliminates relevance and pre-ranking, allowing generated content




Figure 1: The overall architecture of ANGLE framework, comprising retrieval, ad indexing and ranking. After the user inputs a query, the retrieval employs a GDR-LLM to produce valid intent-abstract (int-abs) pairs through dynamic constrained decoding. The Ad Indexing module then retrieves relevant ads based on the generated int-abs pairs. The retrieved ads are directly sent to the ranking stage for sorting before being presented to the user.

to proceed directly to ranking and thus significantly simplifying the retrieval pipeline.

2.2 Generative Retrieval

Generative retrieval (Zhou et al., 2025) reformulates the retrieval problem as a generation task, whereby the model directly generates the content (Li et al., 2023) or SIDs (Zhang et al., 2026) of relevant docs/ads to accomplish retrieval. For example, UniSearch (Chen et al., 2025b) generates SIDs for video, retrieving relevant videos by producing SIDs based on user queries. RARE (Liu et al., 2025a) utilizes commercial intent (CIs) as the representation of ads to enable one-step query-to-ad retrieval.

2.3 Doc Representation

LLM transforms the retrieval into a generation task by generating semantic representations of doc/ads. Existing semantic representations are mainly divided into three categories: semantic identifiers (SIDs), semantic term sets (STerms), and semantic compressed text (SText). SIDs are discrete semantic representations generated by encoding advertising content using models such as RQ-VAE (Lee et al., 2022). For instance, these works (Liu et al., 2025b; Zhang et al., 2025) employs SIDs to effec-

tively represent ads. STerms refer to a set of representative terms extracted directly from the content of candidates. For example, SEAL (Bevilacqua et al., 2022) represents documents using n-grams extracted from the text. SText uses semantically compressed text to represent doc/ads. For example, RARE (Liu et al., 2025a) uses the commercial intent to represent the advertisement, and GRAM (Pang et al., 2025) performs product retrieval from queries using code-based implementation. Unlike the aforementioned work, we propose a novel hierarchical text representation approach to enable generative retrieval for ads.

3 Method

In this section, we introduce in detail the proposed ANGLE framework, and the two components of ANGLE: ad indexing and retrieval.

3.1 End-to-end Generative Framework

In this paper, we propose a **u**nified **G**eneration-discriminative-ranking **r**eaL-time **r**etrieval (**ANGLE**) framework that directly generates high-quality ads based on user queries and sends these candidate ads directly to the ranking. In the ANGLE framework, an integrated generation-




Figure 2: Comparison of ANGLE and MCA: ANGLE uses a LLM to directly generate hundreds of candidate ads for ranking, streamlining the retrieval pipeline.

discriminative-ranking LLM (GDR-LLM) is designed to produce ads that balance relevance and commercial value, enabling direct access to the final ranking stage as shown in the Figure 1. Additionally, the framework employs semantic representations at varying granularities, ranging from coarse-grained intent overviews to fine-grained abstracts, to enhance retrieval accuracy and indexing efficiency. More detailed information about int-abs pairs will be provided in Section 3.3.

The retrieval process is shown in the Figure 1, when ANGLE receives a user query, the GDR-LLM produces int-abs pairs with high commercial value and relevance. This phase is called **retrieval**. After obtaining int-abs pairs, ANGLE queries the ad index to retrieve the corresponding ads, which corresponds to the **ad indexing** stage. Once the ads are acquired, ANGLE bypasses the traditional relevance and pre-ranking modules and directly sends the ads to the ranking stage. The design of the GDR-LLM, as well as the details of the retrieval, ad indexing, and ranking components, will be thoroughly discussed in the following sections.

3.2 Generation-Discriminative-Ranking LLM

This section focuses on the training process of the generation-discriminative-ranking LLM (GDR-LLM), which combines generation, discrimination, and ranking. As shown in the Figure 1, the GDR-LLM relies primarily on three loss functions for training: L_{gen} , L_{dis} , and L_{dpo} . The following sections describe the model design and data composition of these three components.

Learning to Generate. The objective of generation is to produce natural language representations of user search intents and advertising abstracts. The training data primarily originate from online user click data consisting of query-ad pairs. The organization of SFT data will be described in detail in the “Training Data” in section 3.2 and 4.1. During

training, we employ the cross-entropy loss function to measure the discrepancy between the predicted probability distribution and the ground-truth labels, as formally defined as follows,

$$L_{gen} = - \sum_{t=1}^T \log P(y_t | y_{<t}, x) \quad (1)$$

where T is the number of generated tokens, y_t is the t -th token in the target sequence, $y_{<t}$ represents all tokens before the t -th token, x is the model input and $P(y_t | y_{<t}, x)$ is the probability that the model correctly predicts token y_t given the previously generated tokens and the input x .

Learning to Discriminate. The goal of discrimination is to judge the relevance between the user’s query and the int-abs generated by LLM to improve the accuracy of the model generation results. Specifically, we append a special token [CLS] at the end of the input prompt. Let the output representation of the model corresponding to this token be denoted as $h_{[CLS]} \in R^d$. We then apply a multi-layer perceptron (MLP) function $f: R^d \rightarrow R$ to $h_{[CLS]}$ to obtain a scalar relevance score:

$$s = f(h_{[CLS]}). \quad (2)$$

The ground-truth relevance labels, denoted by $y \in \{0, 1\}$, are obtained from real online click-exposure logs. During training, we interpret the scalar score s as the logit for the predicted relevance probability via the sigmoid function:

$$\hat{p} = \sigma(s) = \frac{1}{1 + e^{-s}}, \quad (3)$$

where $\hat{p}$ represents the predicted probability that the input corresponds to a relevant sample. The training objective is to minimize the binary cross-entropy loss between the predicted relevance probability $\hat{p}$ and the ground-truth label y , which is expressed as:

$$L_{dis} = -[y \log(\hat{p}) + (1 - y) \log(1 - \hat{p})]. \quad (4)$$

Learning to Rank. The goal of the ranking stage is to prioritize ads with higher commercial value. Specifically, we adopt direct preference optimization (DPO) (Rafailov et al., 2023) to learn the relative ranking of ads. The training data consists of tuples (x, y_w, y_l) , where x denotes the context information, and y_w and y_l represent the winning (higher-performing) and losing (lower-performing)

Ad Information	Commercial Intents	Ad Abstract
Title: In Endless Winter, as long as the fire burns well, there are no worries about life in the apocalypse!	Role-playing Games	Endless Winter
	Post-apocalyptic Mobile Games	Endless Winter
	Survival Challenge Game	Endless Winter

Table 1: The example of an advertisement int-abs pair: broad intents capture potential needs, precise abstract conveys product information.

ads, respectively. The model assigns a generation probability to each ad y conditioned on context x via the policy π_θ , while π_{ref} is a reference policy. To measure the commercial difference and guide the model to learn the correct ranking, we define a loss function based on the ratio of probabilities as follows:

$$L_{dpo}(\pi_\theta; \pi_{ref}) = -[\log \sigma(\beta \log \frac{\pi_\theta(y_w|x)}{\pi_{ref}(y_w|x)}) - \beta \log \frac{\pi_\theta(y_l|x)}{\pi_{ref}(y_l|x)}] \quad (5)$$

where $\sigma(\bullet)$ is the sigmoid function that maps the difference of log-probability ratios into the range (0,1), reflecting the relative relevance. $\beta > 0$ is a scaling hyperparameter controlling the sharpness of this comparison. By minimizing this loss, the model is encouraged to assign significantly higher relative probabilities to winning ads than to losing ads, thereby achieving a ranking that prioritizes advertisements with greater commercial value.

Finally, the overall loss of GDR-LLM is the sum of multiple loss functions, as shown in Formula 6,

$$L_{model} = \alpha L_{gen} + \beta L_{dis} + \gamma L_{dpo} \quad (6)$$

where α, β, γ are hyperparameters.

Training Data. Given a user query q , we first collect online logs to obtain two sets of advertisements: a set of related ads with clicks $A = \{a_1 \dots a_n\}$ and a set of irrelevant ads $B = \{b_1 \dots b_n\}$ associated with q . We then train the generation model using the relevant set A by optimizing the loss L_{gen} . Meanwhile, the discriminator loss L_{dis} and the ranking loss L_{dpo} are trained using both A and B .

To generate candidate ads for ranking, we first take the union of sets A and B and sort them in descending order of commercial value, producing the candidate list C . This ordering reflects the ranking as predicted by the model trained with the L_{dpo} .

For each candidate ad $c_i \in C$, we further assign relevance labels and determine the generation loss

mask based on its set membership. Specifically, if $c_i \in A$, we set the generation loss mask to 1 and assign a relevance label of 1. If $c_i \in B$, the generation loss mask is set to 0 and the relevance label is 0. These relevance labels correspond to the correlations predicted by the model trained using L_{dis} . The final generation loss is calculated as the element-wise product of the mask and L_{gen} .

3.3 Ad Indexing

We employ LLM to generate commercial intents (SText) and ad abstracts (STerm). An inverted index mapping SText and STerm to ads is then built for efficient retrieval.

Hierarchical Identifiers: Intent x Abstract. In this work, each advertisement is represented by combining two complementary textual sources—broad intent and precise abstract. The semantic intent is a compressed text (SText) that provides a high-level overview of the commercial intent of the ad. In contrast, the text abstract (STerm) is n-grams extracted from the ad display title, preserving specific factual details and salient phrases. By integrating these two components (SText x STerm), the joint representation leverages both the broad conceptual understanding offered by the semantic intent and the precise, grounded information from the extracted abstract.

Table 1 presents an example of an ad’s commercial intent and abstract. For a game title of Endless Winter, the intent refers to its broad commercial category or gameplay type, such as "Role-playing Game" or "Survival Challenge Game." In contrast, the abstract is limited to keywords that appear in the advertisement and are directly related to the title, such as "Endless Winter". Thus, "intent" captures the general purpose or genre, while "abstract" reflects the specific terms highlighted in ad.

Ad Indexing The ad indexing module aims to build an index that maps hierarchical representations to ads. We first use a large language model (LLM) to generate semantic intents for all ads in




Figure 3: Overview of dynamic constrained decoding: generate intent, filter invalid intent, and dynamical reconstruct constrained trees for ad abstract generation.

the database. These intents are deduplicated and cleaned at the industry level, producing approximately 700,000 coarse-grained commercial intent categories. A constrained beam search is then performed across the ad corpus over these intent categories, as shown in ad index construction module of Figure 1. Next, the LLM generates a free (unconstrained) text abstracts for each ad, forming hierarchical representations of intent $\times$ abstract (int-abs) pair, which are used to construct an inverted index. For new ads, intents are generated under constraints, followed by unconstrained abstract generation. This enables efficient index updates without requiring the model to retrain or relearn the mapping from ad representations to ads.

3.4 Retrieval

Dynamical Constrained Beam Search (DCBS). Dynamic constrained beam search leverages a dynamic tree-building module to ensure that all generated advertisements (ads) belong to a predefined valid set. As illustrated in Figure 3 with a beam size of 2, the process begins by constructing a prefix trie of intent-abstract (int-abs) pairs. Generation starts from the root token ([BOS]) by producing intent tokens. During decoding, the ad library index module continuously verifies the validity of ads associated

with each generated abstract linked to the intents. Only abstracts corresponding to valid ads are retained as candidates for further expansion, while those without valid ads are immediately pruned from the search space. If all abstracts under a given intent are invalid, that intent is also discarded, terminating its generation path. For the same query, different users have different sets of valid ads, resulting in different sets of valid int-abs pairs. Therefore, valid int-abs pairs are dynamically generated based on user features. This dynamic filtering effectively restricts decoding to valid int-abs pairs, ensuring that the generation process produces only valid ads in real time.

Typically, organizing the complete set of ads into a constrained tree ensures that every ad can potentially be retrieved during inference. However, many of these retrieved ads may be irrelevant or invalid for the specific user context. For instance, an ad targeted exclusively to female users would be ineffective if the current searcher is male. In practice, the retrieval system relies on target filtering within the MCA module (Figure 2) to exclude ads that do not meet delivery requirements. As a result, a large portion of generated ad candidates, despite satisfying relevance and commercial criteria, may ultimately prove ineffective. By contrast, DCBS directly integrates validity constraints into the generation process, dynamically pruning invalid int-abs pairs in real time. This approach not only guarantees that all generated candidates are valid for the target user but also substantially improves generation efficiency and practical effectiveness by minimizing wasted retrieval and filtering efforts downstream.

Retrieval. This section mainly introduces the retrieval part of the ANGLE. As shown in the retrieval of Figure 3, upon receiving a user query, the GDR-LLM first performs a constrained generation to produce a beam of candidate user search intents. The dynamical constrained beam search then filters out ad abstracts that do not correspond to any valid ads, retaining only those linked to available ads. These valid ad abstracts are reorganized into a constrained tree structure, on which the LLM continues generation. This ensures that every hierarchical output aligns with ads that are valid and ready for display. By dynamically focusing decoding on this refined set of candidates, the framework significantly improves efficiency. After generating all valid int-abs pairs, the discriminator head pro-

Method		HR@10	HR@50	HR@100	MAP	ACR
Semantic-based	Bert	0.1313	0.2384	0.3075	0.3910	47.70%
	SimBert-v2-R	0.1339	0.2265	0.2830	0.3966	47.71%
Generative Retrieval	SimBert-v2-G	0.1294	0.2374	0.3095	0.3852	36.89%
	T5-base	0.0981	0.1880	0.2495	0.2910	29.15%
LLM-based Generative Retrieval	Qwen-1.8B	0.0960	0.1887	0.2524	0.3138	47.27%
	Hunyuan-2B	0.1153	0.2184	0.2889	0.3519	42.51%
	DSI-1B	0.0320	0.0830	0.1058	0.2831	55.49%
Ours	ANGLE-1B	0.1961	0.4834	0.5273	0.4834	69.17%

Table 2: Comparison of ANGLE and Baseline Models in Offline Scenarios.

Online Scenarios	Consumption	GMV	Clicks	Conversions	Exposure
WTS(anonymous)	+1.81%	+2.16%	+1.50%	+1.44%	+2.49%
QBS(anonymous)	+6.57%	+5.01%	+6.11%	+0.72%	+7.82%

Table 3: Application of ANGLE to Real-world Search Systems.

duces relevance scores between the query and each int-abs pair, which are used to assess the quality of the generated outputs.

4 Experiments

In this section, we first present the experimental settings, including the training dataset, evaluation dataset, baselines, and implementation details. We then report the experimental results of ANGLE in both online and offline scenarios, followed by an ablation study to evaluate the contribution and effectiveness of each component within ANGLE.

4.1 Experimental Settings

Training Dataset. In this work, The LLMs used for queries and advertising, as well as their corresponding supervised fine-tuning (SFT) data, differ significantly. The ad-side model employs a base LLM, while the query-side model utilizes the GDR-LLM. The advertising-side data primarily originates from open-source models (e.g., ChatGPT, Achiam et al., 2023). Online advertising information, such as ad title and landing pages, is input into the open-source model, which generates the int-abs pair of each ad. After cleaning, approximately 5,000 training samples were obtained. On the query side, the GDR-LLM is employed to regulate the relevance and commercial value of the generated outputs. Specifically, the SFT data is derived from online click logs. We collected click data from the past month and constructed a training dataset com-

prising 430,000 samples by incorporating negative examples. Subsequently, advertisements associated with each query were commercially ranked based on actual online expenditure to train the DPO model. Specific examples of the training data are presented in Table 5 of Appendix A, and the organization rules for the training data are described in the “Training Data” in section 3.2.

Evaluation dataset. We collected online impression and click data spanning one month and selected 10k queries along with their corresponding clicked ads to form the evaluation dataset. For each query, up to 1k candidate ads were included. To increase the evaluation’s realism and challenge, a set of irrelevant ads was randomly sampled from the ad library and added to the candidate pool. As a result, the final evaluation dataset comprised approximately 220k ad candidates, ensuring both completeness and diversity for a thorough assessment.

Baselines. We compare ANGLE with seven benchmark models to validate its effectiveness.

Semantic-based methods use BERT (Devlin et al., 2019) or SimBert-v2 (RoFormer-Sim, 2021) to generate semantic embeddings for queries and keywords, respectively, and apply approximate nearest neighbor (ANN) to retrieve keywords relevant to the query. Both BERT and SimBert-v2 models are based on a 12-layer transformer architecture with a hidden size of 768 and 12 attention heads, comprising approximately 110 million pa-

Method	HR@10	HR@50	HR@100	MAP	ACR
w/o. Intent	0.1867	0.4291	0.4645	0.4618	68.77%
w/o. Abstract	0.1738	0.3384	0.3519	0.4361	68.88%
w/o. DCBS	0.1442	0.2803	0.2975	0.3130	55.14%
w/o. GDR-LLM	0.1278	0.4329	0.4412	0.4275	62.62%
ANGLE	0.1961	0.4834	0.5273	0.4834	69.17%

Table 4: Ablation Study of ANGLE.

rameters. During training, we used online click logs as positive samples, and randomly sampled within the batch as negative examples.

Generative retrieval (T5, Raffel et al., 2020) and LLM-based generative retrieval (Qwen (Bai et al., 2023), HunYuan (Tencent-Hunyuan, 2026), DSI (Tay et al., 2022)) directly produce business keywords from the query. However, unconstrained decoding may generate keywords that do not correspond to any real ads. To address this issue, we adopt the hierarchical navigable small world (HNSW) (Malkov and Yashunin, 2018) to identify the most similar keyword within the existing database, which is then used as the final output instead of the original LLM-generated keywords.

Implementation Details. We use the Hunyuan-13B to generate hierarchical text for ads. Because the online system has strict time constraints, the query uses the Hunyuan-1B, with a beam size set to 64. To ensure the stability of online system performance, the maximum length of the intent is limited to 3 tokens, while the abstract is restricted to 4 tokens. Hunyuan-1B requires approximately 8 milliseconds to generate each token; therefore, the total time for generating an int-abs pair is limited to within 90 milliseconds, meeting the latency requirements for online system. The ad index is updated hourly, allowing newly added ads to quickly generate int-abs pairs and update the index accordingly. Meanwhile, expired ads are promptly removed from the index.

4.2 Experimental Results

Online Result. We conducted an A/B testing experiment over five days using 20% of the traffic, which resulted in a significant increase of 1.18% in consumption and 2.16% in GMV. At the same time, both the number of conversions and clicks showed varying degrees of improvement, as presented in Table 3. These results demonstrate that ANGLE effectively alleviates the issue of inconsis-

tent target optimization in MCA. Currently, it has been fully applied to WTS (anonymous) and QBS (anonymous) as a supplement to the ranking queue.

Offline Evaluation. To assess the effectiveness of ANGLE, we constructed an evaluation dataset consisting of 10,000 queries and compared its performance with seven strong baselines. The results, summarized in Table 2, demonstrate that ANGLE achieves state-of-the-art (SOTA) performance across all evaluation metrics, including HR, MAP, and ACR.

Ablation Study. We validated the effectiveness of ANGLE through ablation experiments, as shown in Table 4, where "w/o. Intent" refers to representing ads solely using abstracts; "w/o. Abstract" refers to representing ads solely using commercial intents. "w/o. DCBS" refers to removing the DCBS—allowing the LLM to freely generate the int-abs pairs for ads and subsequently using ANN to obtain similar int-abs pairs for the purpose of retrieving ads. "w/o. GDC-LLM" refers to utilize a base LLM to perform constrained generation. The ablation results demonstrate that removing any module negatively impacts the overall performance.

5 Conclusion

This work introduces ANGLE, a unified end-to-end generative framework for real-time ad retrieval. ANGLE integrates ad generation, discrimination, and ranking into a single LLM, enabling direct access to final ranking for ads. This effectively addresses the common issues of logical inconsistency and error propagation in MCA. Hierarchical text representations further alleviate generalization challenges related to SID and reduce reliance on memorizing SID-Ad pairs, improving adaptability and robustness. Extensive evaluations, including online A/B tests, offline experiments, validate ANGLE’s effectiveness and highlight its strong potential for advancing LLM-based advertising solutions.

Limitations

This section discusses the limitations of the current work. Unlike recommendation systems, retrieval systems involve complex components such as dynamic creative selection and intelligent bidding, all of which rely heavily on ranking. Consequently, our proposed model does not eliminate the use of ranking. Future research will investigate the feasibility of removing precise ranking from search systems, aiming to further leverage the potential of large-scale models.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Areej Alsini, Du Q Huynh, and Amitava Datta. 2020. Hit ratio: An evaluation metric for hashtag recommendation. *arXiv preprint arXiv:2010.01258*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, and 1 others. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Michele Bevilacqua, Giuseppe Ottaviano, Patrick Lewis, Scott Yih, Sebastian Riedel, and Fabio Petroni. 2022. Autoregressive search engines: Generating substrings as document identifiers. *Advances in Neural Information Processing Systems*, 35:31668–31683.
- Ben Chen, Xian Guo, Siyuan Wang, Zihan Liang, Yue Lv, Yufei Ma, Xinlong Xiao, Bowen Xue, Xuxin Zhang, Ying Yang, and 1 others. 2025a. Onesearch: A preliminary exploration of the unified end-to-end generative framework for e-commerce search. *arXiv preprint arXiv:2509.03236*.
- Jiahui Chen, Xiaozhe Jiang, Zhibo Wang, Quanzhi Zhu, Junyao Zhao, Feng Hu, Kang Pan, Ao Xie, Maohua Pei, Zhiheng Qin, and 1 others. 2025b. Unisearch: Rethinking search system with a unified generative architecture. *arXiv preprint arXiv:2509.06887*.
- Gordon V Cormack and Thomas R Lynam. 2006. Statistical precision of information retrieval evaluation. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 533–540.
- Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*, pages 191–198.
- Jiaxin Deng, Shiyao Wang, Kuo Cai, Lejian Ren, Qigen Hu, Weifeng Ding, Qiang Luo, and Guorui Zhou. 2025. Onerec: Unifying retrieve and rank with generative recommender and iterative preference alignment. *arXiv preprint arXiv:2502.18965*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Soumik Dey, Wei Zhang, Hansi Wu, Bingfeng Dong, and Binbin Li. 2025. Middleman bias in advertising: Aligning relevance of keyphrase recommendations with search. In *Companion Proceedings of the ACM on Web Conference 2025*, pages 2701–2705.
- Miao Fan, Jiacheng Guo, Shuai Zhu, Shuo Miao, Mingming Sun, and Ping Li. 2019. Mobius: towards the next generation of query-ad matching in baidu’s sponsored search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2509–2517.
- Hongliang Fei, Jingyuan Zhang, Xingxuan Zhou, Junhao Zhao, Xinyang Qi, and Ping Li. 2021. Gemnn: Gating-enhanced multi-task neural networks with feature interaction learning for ctr prediction. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pages 2166–2171.
- Niklas Freymuth, Dong Liu, Thomas Ricatte, and Saab Mansour. 2025. Hierarchical multi-field representations for two-stage e-commerce retrieval. *arXiv preprint arXiv:2501.18707*.
- Kairui Fu, Tao Zhang, Shuwen Xiao, Ziyang Wang, Xinming Zhang, Chenchi Zhang, Yuliang Yan, Junjun Zheng, Yu Li, Zhihong Chen, and 1 others. 2025. Forge: Forming semantic identifiers for generative retrieval in industrial datasets. *arXiv preprint arXiv:2509.20904*.
- Xu Huang, Defu Lian, Jin Chen, Liu Zheng, Xing Xie, and Enhong Chen. 2023. Cooperative retriever and ranker in deep recommenders. In *Proceedings of the ACM Web Conference 2023*, pages 1150–1161.
- Yanhua Huang, Yuqi Chen, Xiong Cao, Rui Yang, Mingliang Qi, Yinghao Zhu, Qingchang Han, Yaowei Liu, Zhaoyu Liu, Xuefeng Yao, and 1 others. 2025. Towards large-scale generative ranking. *arXiv preprint arXiv:2505.04180*.
- Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. 2022. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11523–11532.
- Yongqi Li, Nan Yang, Liang Wang, Furu Wei, and Wenjie Li. 2023. Multiview identifiers enhanced generative retrieval. *arXiv preprint arXiv:2305.16675*.

- Tongtong Liu, Zhaohui Wang, Meiyue Qin, Zenghui Lu, Xudong Chen, Yuekui Yang, and Peng Shu. 2025a. Real-time ad retrieval via llm-generative commercial intention for sponsored search advertising. *arXiv preprint arXiv:2504.01304*.
- Zhanyu Liu, Shiyao Wang, Xingmei Wang, Rongzhou Zhang, Jiaxin Deng, Honghui Bao, Jinghao Zhang, Wuchao Li, Pengfei Zheng, Xiangyu Wu, and 1 others. 2025b. Onerec-think: In-text reasoning for generative recommendation. *arXiv preprint arXiv:2510.11639*.
- 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):824–836.
- Ming Pang, Chunyuan Yuan, Xiaoyu He, Zheng Fang, Donghao Xie, Fanyi Qu, Xue Jiang, Changping Peng, Zhangang Lin, Zheng Luo, and 1 others. 2025. Generative retrieval and alignment model: A new paradigm for e-commerce retrieval. In *Companion Proceedings of the ACM on Web Conference 2025*, pages 413–421.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36:53728–53741.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- RoFormer-Sim. 2021. Roformer-sim. <https://github.com/ZhuiyiTechnology/roformer-sim>.
- Yi Tay, Vinh Tran, Mostafa Dehghani, Jianmo Ni, Dara Bahri, Harsh Mehta, Zhen Qin, Kai Hui, Zhe Zhao, Jai Gupta, and 1 others. 2022. Transformer memory as a differentiable search index. *Advances in Neural Information Processing Systems*, 35:21831–21843.
- Tencent-Hunyuan. 2026. Tencent-hunyuan. <https://github.com/Tencent-Hunyuan>.
- Lidan Wang, Jimmy Lin, and Donald Metzler. 2011. A cascade ranking model for efficient ranked retrieval. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, pages 105–114.
- Jun Zhang, Yi Li, Yue Liu, Changping Wang, Yuan Wang, Yuling Xiong, Xun Liu, Haiyang Wu, Qian Li, Enming Zhang, and 1 others. 2025. Gpr: Towards a generative pre-trained one-model paradigm for large-scale advertising recommendation. *arXiv preprint arXiv:2511.10138*.
- Kun Zhang, Jingming Zhang, Wei Cheng, Yansong Cheng, Jiaqi Zhang, Hao Lu, Xu Zhang, Haixiang Gan, Jiangxia Cao, Tenglong Wang, and 1 others. 2026. Onemall: One model, more scenarios—end-to-end generative recommender family at kuaishou e-commerce. *arXiv preprint arXiv:2601.21770*.
- Zhixuan Zhang, Yuheng Huang, Dan Ou, Sen Li, Longbin Li, Qingwen Liu, and Xiaoyi Zeng. 2023. Rethinking the role of pre-ranking in large-scale e-commerce searching system. *arXiv preprint arXiv:2305.13647*.
- Binglei Zhao, Houying Qi, Guang Xu, Mian Ma, Xiwei Zhao, Feng Mei, Sulong Xu, and Jinghe Hu. 2025. A hybrid cross-stage coordination pre-ranking model for online recommendation systems. In *Companion Proceedings of the ACM on Web Conference 2025*, pages 621–630.
- Guorui Zhou, Honghui Bao, Jiaming Huang, Jiaxin Deng, Jinghao Zhang, Junda She, Kuo Cai, Lejian Ren, Lu Ren, Qiang Luo, and 1 others. 2025. Openonerec technical report. *arXiv preprint arXiv:2512.24762*.

A Problem Definition

In search advertising, given a user query q and a large candidate ad set $A = \{a_1, a_2 \dots a_n\}$, where n can reach tens of millions, the goal of the retrieval system is to select ads that satisfy user intent and maximize the platform’s commercial value. Traditionally, this involves a four-stage cascaded pipeline: retrieval, relevance, pre-ranking, and ranking, where each stage is optimized independently. This leads to inconsistent objectives and error propagation, degrading overall performance. To overcome these issues, we propose a unified end-to-end framework that integrates retrieval, relevance, and ranking to directly generate the final set of high-quality ads, thereby minimizing inter-stage discrepancies and cumulative errors. For GDR-LLM, training data are constructed using online logs. Specifically, for each query, we collect advertisements that have been clicked by users in the online environment and retrieve their corresponding intent-abstract (int-abs) pairs. By associating each query with the int-abs pairs of the ads it led users to click, we create a high-quality training dataset consisting of query and int-abs pairs. This data construction method leverages real user interactions to provide relevant and diverse supervision signals for model fine-tuning, thereby improving the model’s ability to understand query-advertisement relationships in practical scenarios.

B Fine-tuning Data

This section primarily describes the datasets used for fine-tuning the GDR-LLM on the query side and the base LLM on the advertisement side. As shown in Table 5, the training data for the base LLM are derived from open-source models such as ChatGPT. For each advertisement title, the model is first prompted to generate the advertising intent, followed by a request to produce a concise abstract of the advertisement. The specific prompts used in this process are detailed in Table 5. Subsequently, both the generated intent and abstract undergo rigorous data cleaning, and are then combined to create multiple intent-abstract pairs corresponding to each advertisement title. This structured approach ensures that the training data are high-quality and relevant, thereby enhancing the model’s performance in downstream tasks.

C Evaluation Metrics

In this paper, ACR (Fan et al., 2019), Hit Ratio (HR@K) (Alsini et al., 2020), and Mean Average Precision (MAP) (Cormack and Lynam, 2006) are used to evaluate the performance of ANGLE. *Ad Coverage Rate* (ACR) is a metric used to evaluate the proportion of unique advertisements that a system is able to retrieve from the total pool of available ads. A higher ACR indicates that the retrieval system can expose users to a broader variety of ads, which is important for both user experience and advertiser fairness. *Ad Pave View* (AdPV) is the number of requests with ad recall, and *Pave View* (PV) is the number of requests.

$$ACR = AdPV/PV \quad (7)$$

Hit Ratio (HR@K) is a metric that measures the proportion of relevant items appearing in the top K results of a retrieval system. Specifically, Hits@K represents the number of relevant ads within the top-K retrieved candidates that belong to the ground truth set and Ground Truth (GT) represents set of candidate ads. A higher HR@K reflects better retrieval accuracy and effectiveness for users.

$$HR@K = \frac{Hits@K}{|GT|} \quad (8)$$

Mean Average Precision (MAP) measures overall ranking quality by averaging the Average Precision (AP) across all queries, as defined in Equation 9.

$$MAP = \frac{1}{Q} \sum_{q \in Q} AP_q \quad (9)$$

Average Precision (AP) for a given query is described in Equation 10. In this formula, Ω_q denotes the set of true relevant items, p_{qj} is the rank position of ad_j in the output list, and $p_{qj} < p_{qi}$ signifies that ad_j is ranked higher than ad_i .

$$AP_q = \frac{1}{|\Omega_q|} \sum_{i \in \Omega_q} \frac{\sum_{j \in \Omega_q} h(p_{qj} < p_{qi}) + 1}{p_{qi}} \quad (10)$$

D Online Inference Support

For online inference deployment, we have constructed a dedicated GPU cluster comprising hundreds of L40 GPUs. This cluster has been optimized for workload distribution and peak GPU utilization, achieving a utilization rate of nearly 90%.

The trained models are converted to FP8 format, enabling each L40 GPU to process approximately 30 queries per second. We also optimize each stage of the inference pipeline, including input preprocessing, model inference, output postprocessing, and network transmission, in order to reduce overall end-to-end latency and ensure that each request can be completed within 60 ms.

Model Used	Prompt	Output	Data Size
GDR-LLM query-side	As a search system, when a user enters a query = "12306", please first analyze the user's search intent. Then, recommend suitable advertisements based on that intent, and generate corresponding ad abstract information for display.	From online click logs: Train Ticket Booking_12306 Official; High-Speed Rail Ticket_12306 Official; Online Ticketing_12306 Official; Travel Services_12306 Official.....	5k
Base LLM ads-side	Task Requirements: Generate advertising titles based on user search queries and extract a abstract for each title. The abstract should follow the format: Commercial Intent_Brand/Core Service, where "Commercial Intent" specifies a concrete product or service category, and "Brand/Core Service" clearly identifies the brand or core service. Correct Example: Organic Vegetables_Meituan Grocery (product category + brand), Incorrect Example: Organic Vegetables_Direct from Source ("Direct from Source" is a selling point, not a brand or service). Business Point Guidelines: Must directly reflect the user's search intent (e.g., "Fresh Food Delivery," "Grocery Shopping App") Avoid vague terms (e.g., "Quality Guaranteed," "Fast Delivery"). abstract Guidelines: Must include a brand name (e.g., "Meituan Grocery") or core service name (e.g., "Gaming"). Given the query "12306", please perform the following steps: 1. Identify specific business points from the query, ensuring alignment with the above criteria. 2.Generate ad titles conveying commercial information for each business point. 3.Extract a compliant abstract from each title. 4.Combine each business point with its abstract in the format: Business Point_Brand/Core Service. 5.On the last line, concatenate all results separated by semicolons.	Output from open source LLM models (e.g. ChatGPT): Step-by-step execution: 1. Business point analysis: * Train ticket booking (User need: purchasing train tickets/high-speed rail tickets) * High-speed rail ticket discounts (User need: purchasing tickets at lower prices) * Online ticketing platforms (User need: booking tickets through official channels) * Travel services (User need: train schedule inquiries/ticket availability monitoring) 2. Advertisement headline generation: * Train ticket booking: Official 12306 booking: High-speed rail tickets starting from 50% off, online seat selection with no service fees! * High-speed rail ticket discounts: 12306 limited-time offer! All student/adult tickets at 20% off * Online ticketing platforms: 12306 official booking: Real-time ticket availability + standby ticket grabbing, click to purchase * Travel services: 12306 official services: Train schedules, ticket availability monitoring, worry-free travel! 3. Title summarization and extraction: * Train ticket booking : Train Ticket Booking_12306 Official * High-speed rail ticket discounts : High-Speed Rail Ticket_12306 Official * Online ticketing platform : Online Ticketing_12306 Official * Travel services : Travel Services_12306 Official Final output: Business Point_Brand/Core Service: Train Ticket Booking_12306 Official; High-Speed Rail Ticket_12306 Official; Online Ticketing_12306 Official; Travel Services_12306 Official	430k

Table 5: Details of fine-tuning data.