

# Towards Boosting LLMs-driven Relevance Modeling with Progressive Retrieved Behavior-augmented Prompting

Zeyuan Chen<sup>1</sup>, Haiyan Wu<sup>2</sup>, Kaixin Wu<sup>1</sup>, Wei Chen<sup>1</sup>, Mingjie Zhong<sup>1</sup>, Jia Xu<sup>1</sup>,  
Zhongyi Liu<sup>1</sup>, Wei Zhang<sup>3</sup>

<sup>1</sup>Ant Group, <sup>2</sup>Alibaba Group,

<sup>3</sup>School of Computer Science and Technology, East China Normal University

{chenzeyuan, daniel.wkx, qianmu.cw, mingjie.zmj, steve.xuj, zhongyi.lzy}@antgroup.com,

wuhaiyan.why@taobao.com, zhangwei.thu2011@gmail.com

## Abstract

Relevance modeling is a critical component for enhancing user experience in search engines, with the primary objective of identifying items that align with users' queries. Traditional models only rely on the semantic congruence between queries and items to ascertain relevance. However, this approach represents merely one aspect of the relevance judgement, and is insufficient in isolation. Even powerful Large Language Models (LLMs) still cannot accurately judge the relevance of a query and an item from a semantic perspective. To augment LLMs-driven relevance modeling, this study proposes leveraging user interactions recorded in search logs to yield insights into users' implicit search intentions. The challenge lies in the effective prompting of LLMs to capture dynamic search intentions, which poses several obstacles in real-world relevance scenarios, i.e., the absence of domain-specific knowledge, the inadequacy of an isolated prompt, and the prohibitive costs associated with deploying LLMs. In response, we propose *ProRBP*, a novel *Progressive Retrieved Behavior-augmented Prompting* framework for integrating search scenario-oriented knowledge with LLMs effectively. Specifically, we perform the user-driven behavior neighbors retrieval from the daily search logs to obtain domain-specific knowledge in time, retrieving candidates that users consider to meet their expectations. Then, we guide LLMs for relevance modeling by employing advanced prompting techniques that progressively improve the outputs of the LLMs, followed by a progressive aggregation with comprehensive consideration of diverse aspects. For online serving, we have developed an industrial application framework tailored for the deployment of LLMs in relevance modeling. Experiments on real-world industry data and online A/B testing demonstrate our proposal achieves promising performance.

## 1 Introduction

In today's landscape of excessive information, search engines have become critical for online content platforms, allowing users to swiftly find preferred content that match their search queries. To ensure a user-friendly experience, relevance modeling is crucial to preserve a satisfactory connection between a query and the displayed outcomes, forming a core element of search engine functionality.

In the relevant literature, foundational studies (Robertson et al., 2009; Shah and Pomerantz, 2010; Svore and Burges, 2009) have engaged in feature engineering to accomplish text matching, yet they lacked sufficient generalization and accuracy. Subsequently, deep learning-based approaches have risen as a new paradigm, with two primary categories: *representation-based* approaches (Shen et al., 2014; Palangi et al., 2014; Rao et al., 2019) and *interaction-based* approaches (Parikh et al., 2016; Chen et al., 2016; Hu et al., 2014; Pang et al., 2016). Lately, pre-trained architectures such as BERT (Devlin et al., 2018) have achieved significant progress in Natural Language Understanding (NLU) tasks. As a result, several studies (Yao et al., 2022; Lu et al., 2020; Reimers and Gurevych, 2019; Jin et al., 2023) are introduced that aim to capture the semantic relationships between queries and items. Most recently, Large Language Models (LLMs) have showcased their exceptional capabilities across a wide range of Natural Language Processing (NLP) applications. These models, such as GPT (Radford et al., 2019), LLaMA (Touvron et al., 2023), and GLM (Du et al., 2022), are trained on massive corpora of texts, which enables them to maintain an exhaustive world knowledge. Nonetheless, identifying user search intentions accurately remains challenging when relying solely on semantic understanding, due to the absence of specialized domain knowledge required for complex industrial search scenarios. The texts of queries and items in

---

Correspondence to Jia Xu and Wei Zhang.

Alipay search scenario are quite short and ambiguous, making it hard to convey effective information contained in their identity. For example, given a query “Zhe Yi”, the abbreviation of a hospital, it is hard to comprehend the actual semantics. But its historical clicked items include “the first affiliated hospital of Zhejiang University”, indicating strong correlations between them to help search intention identifying. As such, leveraging behavior data to assist relevance modeling is a natural strategy.

Existing studies (Chen et al., 2023; Zeng et al., 2022; Zhu et al., 2021; Li et al., 2021; Pang et al., 2022) have primarily conducted the use of user behavior data. But they all consider constructing the pre-training dataset or the topology structure based on click behaviors without integrating semantics fully and effectively. Despite the success of LLMs, it’s still uncertain how well they can integrate world knowledge and specialized domain knowledge represented by user behavior data to master relevance modeling. To this end, this paper aims to investigate the potential of LLMs in relevance modeling with user search behavior. By utilizing strategic prompting techniques, specialized domain knowledge could be easily injected into LLMs for relevance modeling, whereas the performance varies due to the following unresolved issues: (i) *The acquisition of domain-specific knowledge*. Though domain-specific knowledge is vital in improving search scenario-oriented relevance modeling capabilities of LLMs, not all knowledge is beneficial. The noisy user behavior data may mislead LLMs to undesired judgements. Moreover, specialized domain knowledge of search scenarios undergoes rapid changes on a daily basis. The limited capacity of LLMs to adapt swiftly to these changes presents a significant obstacle to their ability to render accurate relevance judgments. (ii) *The inadequacy of an isolated prompt*. Despite LLMs could derive the relevance degree exploiting an isolated prompt, LLMs exhibit insensitivity to input, meaning they lack awareness of the aspects from which to infer relevance. In addition, the isolated prompts place greater demands on the quality of the prompts itself. Except for the aforementioned issues, deploying LLMs affordably in industrial scenarios is also a consideration worth addressing.

To address the above problems, we propose a novel *Progressive Retrieved Behavior-augmented Prompting* framework for integrating search scenario-oriented knowledge with LLMs (dubbed as *ProRBP*). To acquire domain-specific knowl-

edge in time, we perform a user-driven behavior neighbor retrieval from the daily updated search logs, retrieving candidates that users consider to meet their expectations currently. Then we anticipate employing advanced prompting techniques that progressively improve the outputs of the LLMs, followed by a progressive aggregation with comprehensive consideration of diverse aspects to form a holistic relevance model. As for the online serving of LLMs, we design an industrial implementation framework enabling LLMs to fully handle search relevance scenarios with the affordable cost.

In summary, we make the following contributions of this paper:

- To the best of our knowledge, we are the first to successfully investigate the potential of LLMs with user behavior data to master relevance modeling.
- We propose *ProRBP* with two novel plug-in modules. Firstly, a user-driven behavior neighbors retrieval is developed to acquire domain-specific knowledge in time. Secondly, the proposal of progressive prompting and aggregation can strengthen the judgement of relevance for LLMs.
- We explore an industrial implementation A.1 enabling LLMs to fully handle search relevance scenarios in Alipay search with the affordable cost and latency.
- We demonstrate *ProRBP* framework achieves superior performance through experiments on real-world industry data and online A/B testing. It has been deployed online and outperforms prior approaches (Chen et al., 2023) on core metrics.

## 2 Related Work

Relevance modeling in search can be viewed as a text matching problem as the sub-domain of information retrieval (IR). The majority of work focuses on distinctions at the semantic level, while a minority of methods judge the relevance from a behavioral perspective.

Current semantics-driven approaches can be classified into two aspects: *feature-based* approaches and *deep learning-based* approaches. The first category is centered on manual-crafted features such as TF-IDF similarity and BM25 (Svore and Burges, 2009). Despite their usefulness, these feature-based

approaches have limited generalization ability due to their domain-specific features and require significant labor resources.

In order to address the limitations of the above approaches, deep learning-based approaches emerge as the new paradigm, which can be broadly classified into *representation-based* approaches and *interaction-based* approaches. The former focuses on learning a low-dimensional representation of data while the latter emphasizes capturing the interaction between inputs. For instance, DSSM (Shen et al., 2014) is a classical two-tower representation-based model that encodes the query and the document separately. In this paradigm, recurrent (Palangi et al., 2014; Tai et al., 2015) and convolutional (Hu et al., 2014; Shen et al., 2014) networks are adopted to extract low-dimensional semantic representations. For these methods, the encoding of each input is carried out independently of the others. Consequently, these models face challenges in modeling complex relationships. To overcome this limitation, interaction-based models are proposed. DecompAtt (Parikh et al., 2016) leverages attention network to align and aggregate representations. In parallel, recurrent (Chen et al., 2016) and convolutional (Hu et al., 2014; Pang et al., 2016) networks are employed for modeling complex interactions.

In recent times, pre-trained models like BERT (Devlin et al., 2018) have made remarkable strides in Natural Language Understanding (NLU). As a result, representation-based (Yao et al., 2022; Lu et al., 2020; Reimers and Gurevych, 2019; Jin et al., 2023) and interaction-based architectures (Wang et al., 2019) are proposed to leverage the capabilities of these models to encode semantic correlations between queries and items. Most recently, Large Languages Models (LLMs) like GPT (Radford et al., 2019), LLaMA (Touvron et al., 2023), BLOOM (Scao et al., 2022) and GLM (Du et al., 2022) trained on massive corpora of texts have shown their superior ability in language understanding, generation, interaction, and reasoning tasks. (Sun et al., 2023) investigates the potential of utilizing LLMs for searching and demonstrates that appropriately instructs ChatGPT and GPT-4 can produce competitive and even superior results to supervised methods widely used information retrieval benchmarks. (Chen et al., 2023) tries to deal with long-tail query-item matching through LLMs efficiently and effectively. In these research work, they tend to exploit world knowledge stored

in parameters of LLMs to judge the relevance between the query and item. However, general LLMs can not adapt to the industrial scenario due to the lack of domain-specific knowledge and insensitivity to the relevance judgement. In this work, we try to explore the LLMs-driven relevance modeling comprehensively in Alipay search engine to meet the above requirements.

In addition to textual information, there are a few related works that aim to integrate user behavior data into their models. The utilization of user behavior data can provide valuable insights into search intention, which can then be used to enhance the relevance of the search engines. MASM (Yao et al., 2021) leverages the historical behavior data to complete model pre-training as a weak-supervision signal with a newly proposed training objective. (Zhu et al., 2021; Li et al., 2021; Pang et al., 2022) try the incorporation of click graphs to enhance the effectiveness of search systems. (Chen et al., 2023) endeavors to exploit behavior neighbors while considering interaction granularity and topology structure. However, no work has fully integrated LLMs with user behavior comprehensively in relevance modeling. We target to do this.

### 3 Problem Formulation

Assume we have the target query  $q$  and target item  $i$  needed to predict the relevance degree exploiting LLMs. In essence, referring to PET (Schick and Schütze, 2020), it could be formulated as follows: with the designed prompt  $\tau(q, i)$ , LLMs can determine which verbalizer  $v$  (i.e., “relevant” or “irrelevant”) is the most likely substitute for the mask based on the likelihood  $\mathcal{P}(v|\tau(q, i))$ . The most naive prompt in the relevance modeling could be formulated as:

$$\tau(q, i) = Is [q] and [i] related? [mask], \quad (1)$$

The relevance label  $y_{qi} \in \{0, 1\}$  can be associated with a verbalizer (i.e., “irrelevant” or “relevant”) from the vocabulary of LLMs to denote the relevance degree between  $q$  and  $i$ . To enable the adaptation of general LLMs to the relevance modeling task, supervised fine-tuning operation is selected using the cross-entropy loss function  $\mathcal{L}_{ce}$ . Then the relevance degree could be given from LLMs for subsequent applications in Alipay search scenario. It worth noting that the above formulation is the basic explanation of relevance modeling using LLMs.




Figure 1: The pipeline of user-driven behavior neighbor retrieval.

We will further explore novel ways below based on this basic.

## 4 Methodology

In this section, we elaborate the proposed *ProRBP* framework with its novel plug-in modules as depicted in Figure 2.

### 4.1 User-driven Behavior Neighbor Retrieval

Generally, LLMs possess an exhaustive world knowledge with the benefit of massive corpora of texts pre-training. However, LLMs still struggle to understand user search intentions for short and ambiguous queries and items from Alipay search due to the lack of specialized and rapidly evolving domain knowledge. A promising approach for addressing the above is retrieval-augmented language modeling (Gao et al., 2023), grounding the LLMs during generation by conditioning on relevant candidates retrieved from an external knowledge source. Drawing inspiration from this, we devise a user-driven behavior neighbor retrieval module. This module could retrieve the daily search logs of users to obtain daily changing behavior neighbors that users consider relevant. The pipeline of user-driven behavior neighbor retrieval is depicted as Figure 1. Due to the limited number of items displayed for a query in Alipay search scenario, almost all results can be seen by users. With this in mind, we could analyze that a higher click-through rate reflects that users believe the corresponding query-item pairs can better meet their search intents and needs when the exposure PV (i.e., page view) reaches a certain quantity. Thus, we filter out the query-item pairs with less than 100 exposure PV and formulate the remaining search logs as high-confidence logs. Then, we utilize the high-confidence logs from the past month to calculate the click-through rate for the exposed query-item pairs. To mitigate the effect of noises, query-item pairs are selected above a click-through threshold (e.g., 0.2) and the neighbors are arranged in

descending order based on click-through rate from the query and item perspective, respectively. And we only choose top-K (e.g., 20) neighbors for corresponding query and item. Through this approach, it is possible to select dual behavior neighbors with high confidence, which in turn can assist LLMs in the process of in-context learning (Ram et al., 2023; Wang et al., 2023; Liao et al., 2024; Liu et al., 2024; Wei et al., 2024). To ensure the timeliness of knowledge, we utilize the daily updated search logs to repeat the above operation every day and construct daily separate indexes for the dual behavioral neighbors of queries and items. The neighbors could be retrieved from indexes for the corresponding query  $q$  and item  $i$ , denoted as  $\mathcal{N}^q$  and  $\mathcal{N}^i$  respectively. Need to clarify is that though we can obtain the behavior neighbors for the majority of queries or items, there is still a small portion of queries or items whose behavior neighbors cannot be obtained. In response to this situation, we will set them as empty in the prompt. Besides, we could obtain the attributes (i.e., brand, keyword, intent) of the query  $q$  and item  $i$  from the industrial knowledge base, which could be denoted as  $\mathcal{A}^q$  and  $\mathcal{A}^i$ . They can be treated as a supplement to the domain-specific information. Then we could construct our daily prompt  $\tau(q, i, \mathcal{A}^q, \mathcal{A}^i, \mathcal{N}^q, \mathcal{N}^i)$  based on the above information. The daily prompt will be send to LLMs for further operation.

### 4.2 Progressive Prompting and Aggregation

Despite the fact that LLMs can determine a relevance score using an isolated prompt, they exhibit insensitivity to the input. It indicates that they lack the awareness required to decide the aspects from which to infer relevance. Furthermore, relying on isolated prompts imposes higher demands on the quality of the prompts itself, as the diversity of domain knowledge and the sensitivity to slight modifications of prompts may lead to unexpected results. As a result, it's expected that using progressive prompting and aggregation process to gradually guide and consider diverse aspects will eventually lead to the creation of a unified model that assesses relevance.

**Progressive Prompting.** It is designed for making LLMs sensitive to the diverse aspects for the relevance judgement and improving the robustness of the model performance. Specifically, we firstly decompose the mentioned prompt above and construct least-to-most prompts step by step (Zhou et al., 2022; Wei et al., 2022) as:

$$\tau(\mathcal{N}^q, \mathcal{N}^i) \hookrightarrow \tau(\mathcal{A}^q, \mathcal{A}^i, \mathcal{N}^q, \mathcal{N}^i) \hookrightarrow \tau(q, i, \mathcal{A}^q, \mathcal{A}^i, \mathcal{N}^q, \mathcal{N}^i), \quad (2)$$

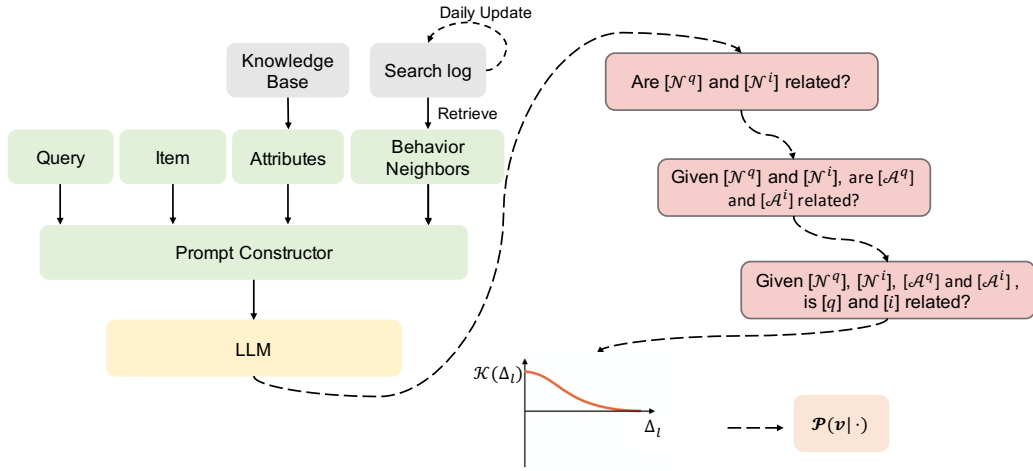


Figure 2: The proposed framework *ProRBP*.

The prompts are documented as shown in Figure 2. Besides, different prompting documents could obtain the consistent improvement in our local experiments exploiting this module.

Then LLMs could reason the likelihood of the verbalizer exploiting the least-to-most prompts step by step for the sensitivity to relevance judgement and stable prediction results. Simply, all least-to-most prompts share the same relevance label  $y_{qi}$  and all least-to-most sub-tasks are supervised by the cross-entropy loss in parallel explicitly. The aforementioned method can be extended to  $L$  progressive prompts and the supervised loss could be denoted as  $\mathcal{L}_{auxi} = \sum_{l=1}^L \mathcal{L}_{ce}^l$ , where  $\mathcal{L}_{ce}^l$  represents the loss of the sub-task corresponding to  $l$ -th prompt.

**Progressive Aggregation.** Once we obtain the  $L$  least-to-most probabilities, we expect to learn the progressive relationship of solutions to sub-tasks. Intuitively, Our least-to-most paradigm is incremental not only in terms of information volume, but also in importance. Hence, we adopt the kernel function  $\mathcal{K}$  to model the incremental tendency and the choice of kernel function can be a Gaussian kernel function, exponential kernel function, logarithmic decay kernel function, etc. After experimental comparison, we have chosen the exponential kernel function  $\mathcal{K}(\Delta_l) = \text{Exp}(\Delta_l|\lambda)$ , where  $\lambda$  is a learnable parameter and  $\Delta_l$  is the degree of attenuation of the  $l$ -th prompt defined by us. The overall relevance score could be acquired through aggregating the probabilities from the least-to-most sub-tasks progressively as:

$$\mathcal{P}(v|\cdot) = \sum_{l=1}^L \mathcal{K}(\Delta_l) \times \mathcal{P}(v|\tau_l), \quad (3)$$

This relevance score could be treated as the production of overall task, deserving to supervise mainly. The loss function is given by  $\mathcal{L}_{main}$ . The hybrid objective function can be formulated as:

$$\mathcal{L} = \mathcal{L}_{main} + \alpha \mathcal{L}_{auxi}, \quad (4)$$

where  $\alpha$  is a hyper-parameter to control the strength of the least-to-most sub-tasks.

## 5 Experiments

### 5.1 Experimental Setup

**Datasets.** In order to evaluate the performance of all the models with reliability, we select the real-world industry data used in mini apps search scenario of Alipay search engine and present its statistics in Table 1. The dataset is labeled by human annotators, where Good and Bad annotations denote label 1 and 0 respectively. User historical behavior data is sampled from the search logs of the search engine. Although datasets such as WANDS\* and MSLR† are publicly available, they do not contain the requisite user historical behavior data. Hence, we select this in-house data to evaluate the proposed framework. Other related work (Yao et al., 2022, 2021; Zhu et al., 2021; Li et al., 2021; Chen et al., 2023) also selects one in-house data to evaluate the proposed approaches. The partial data of this paper was released before‡.

**Baseline Models.** We select a set of popular NLU (Natural Language Understanding)-

\*<https://github.com/wayfair/WANDS/tree/main>

†<https://www.microsoft.com/en-us/research/project/mslr/>

‡<https://github.com/alipay/BehaviorAugmentedRelevanceModel>

Dataset	# Sample	# Query	# Item	# Good	# Bad
Train	773,744	87,499	88,724	460,610	313,134
Valid	97,032	40,754	25,192	57,914	39,118
Test	96,437	40,618	25,004	57,323	39,114

Table 1: Statistics of the human-annotated dataset.

based, behavior-based and NLG (Natural Language Generation)-based relevance models as baselines. For the NLU-based models, we choose three common models including two-tower and single-tower architectures: DSSM (Shen et al., 2014), ReprBert (Yao et al., 2022), Bert (Devlin et al., 2018). The behavior-based models all consider constructing the pre-training dataset (MASM (Yao et al., 2021)) or the topology structure based on click behaviors (TextGNN (Zhu et al., 2021), AdsGNN (Li et al., 2021), BARL-ASe (Chen et al., 2023)). For the NLG-based models, we choose two foundation models including two set of architectures: the causal decoder (BLOOM (Scao et al., 2022)) and the prefix decoder (GLM (Du et al., 2022)).

**Evaluation Metrics.** We use Area Under Curve (AUC), F1-score (F1), and False Negative Rate (FNR) to measure the multidimensional performance of all models. AUC and F1 are commonly used in the studied area, of which higher metric values represent better model performance. Conversely, lower False Negative Rate (FNR) values are preferable, as they indicate a lower false filtering rate of models. Note that AUC often serves as the most significant metric in our task while the others provide auxiliary supports for our analysis.

**Model Implementations.** We tune our model for 5 epochs with batch size 64 and learning rate  $3e-05$  in the supervised fine-tuning stage of LLMs. For the foundation models, we select the 1.1B BLOOM<sup>§</sup> and GLM with different magnitude of parameters (e.g., 0.3B, 2B and 10B) pre-trained by Alipay. The number of retrieved behavior neighbors is set to 20. We tune the parameter  $\alpha$  within the ranges of  $\{0, 0.05, 0.1, 0.2, 0.3, 0.5, 1.0\}$ . We conduct the experiments on NVIDIA Tesla A100 GPUs.

## 5.2 Experimental Results

**Performance Comparison.** Table 2 shows the overall comparison with baselines. Our findings indicate that DSSM performs poorly since it merely encodes the embeddings of the query and item in-

<sup>§</sup><https://huggingface.co/bigscience/bloom-1b1>

Method	AUC	F1	FNR (-)
DSSM	0.8356	0.8210	0.1389
ReprBert	0.8388	0.8376	0.1280
Bert	0.8540	0.8534	0.1150
MASM	0.8547	0.8318	0.1289
TextGNN	0.8847	0.8489	0.1290
AdsGNN	0.8878	0.8458	0.1454
BARL-ASe	0.9078	0.8658	0.1054
GLM-0.3B	0.8608	0.8585	0.1106
+ <i>ProRBP</i>	0.9105	0.8778	0.1035
BLOOM-1.1B	0.8543	0.8511	0.1174
+ <i>ProRBP</i>	0.8991	0.8675	0.1121
GLM-2B	0.8619	0.8598	0.1065
+ <i>ProRBP</i>	0.9120	0.8776	0.1007
GLM-10B	0.8751	0.8656	0.1006
+ <i>ProRBP</i>	0.9143	0.8801	0.0973

Table 2: Main results on real-world industry data. (-) denotes the lower value corresponds to better performance. Improvements over variants are statistically significant with  $p < 0.05$ .

dependently. By further comparing DSSM with ReprBert, we find the performance is improved to a certain degree. This demonstrates the pre-trained models utilizing the large corpus can bring additional gains. Compared to ReprBert, Bert achieves better performance, as the interaction-based models possess more advanced text relevance modeling abilities than representation-based models.

For the models of MASM, TextGNN, and AdsGNN, they exploit historical behavior data to build pre-training dataset or click graph to enhance the effectiveness of search systems. And they achieve significantly better results than the above-mentioned methods in the metric of AUC rather than F1 and FNR. This can be attributed to the introduction of auxiliary signals, which may bring improvement for the ranking ability of models but introduce some noises leading to the loss of F1 and FNR. The performance differences among them depend on the utilization of behavior data and modeling granularity. The newly proposed model BARL-ASe achieves the best performance in behavior-based relevance models.

GLM-0.3B, BLOOM-1.1B, GLM-2B and GLM-10B denote that we use the corresponding language models with the amount of parameters to perform relevance judgements exploiting the naive prompts as stated in Section 3. They exhibit relatively good performance, demonstrating the posi-

tive effect of massive training datasets and large-scale parameters. Performance variances among them with different parameters appear minimal, possibly because the simplicity of the relevance task doesn't fully tap into the advantages of LLMs. Their performance has a certain gap compared to behavior-based models especially on the core metric AUC. This may be attributed to the lack of domain-specific knowledge and the limitations of an isolated prompt. Thus, *ProRBP* framework is proposed to address the mentioned problem and significantly improve the foundation models' performance. The performance differences between BLOOM and GLM may stem from differences in architecture and training corpora.

Overall, exploiting our framework *ProRBP* could yield the best performance and ensures efficiency and economy. In the online serving A.1, we select GLM-10B+*ProRBP* to perform offline inference and GLM-2B+*ProRBP* to proceed online prediction considering affordable cost and efficiency. Besides, they both obtain better gains under all the metrics on the evaluation dataset compared to the second-best performed model BARL-ASe that was deployed in our scenario before.

Method	AUC	F1	FNR (-)
<i>ProRBP</i>	0.9120	0.8776	0.1007
-BNR	0.8943	0.8656	0.1037
-PPA	0.8872	0.8645	0.1045
-Both	0.8619	0.8598	0.1065

Table 3: Ablation study of GLM-2B+*ProRBP*

**Ablation Study.** To investigate the contributions of key components, we provide the following variants of our complete framework: (1) “-BNR” denotes discarding user-driven behavior neighbor retrieval strategy; (2) similarly, “-PPA” denotes discarding progressive prompting and aggregation module. (3) “-Both” represents removing two modules mentioned above.

The results shown in Table 3 lead to the conclusion that both “BNR” and “PPA” strategies yield significantly positive results. And “PPA” demonstrates greater significance than “BNR”.

### 5.3 Parameter Sensitivity

From the variation trends of Figure 3, we could observe that: (1) As the number of behavioral neighbors increases, the model's performance continues to improve, tending to plateau at around 20. (2)




Figure 3: Result variation with different settings.

The exponential kernel function yields the best result among the mean-pooling (Mean), Gaussian (Norm) and logarithmic (Log) kernel functions. (3) The better results are achieved when setting  $\alpha$  to a suitable value. Too large (e.g., 1.0) or small (e.g., 0.01) value will lead to a performance drop.

### 5.4 Online A/B Testing

The proposed method has been deployed in the relevance stage of Alipay search platform providing search service of mini apps and demonstrated its significant performance gains in online A/B testing compared with the previous model BARL-ASe. The each experiment takes about 3% proportion of Alipay search traffic for two weeks. Compared with the previously deployed model, the proposed method improves the valid PV-CTR<sup>¶</sup> by 0.33% on average without causing an increase in latency. And the results of human annotations show the model can reduce the rate of irrelevant results by 1.07% points on average. The results demonstrate that our proposal can improve the experience of users.

## 6 Conclusion

This paper studies the relevance modeling problem by integrating world knowledge stored in the parameters of LLMs with specialized domain knowledge represented by user behavior data for achieving promising performance. The novel framework *ProRBP* is proposed, which innovatively develops user-driven behavior neighbor retrieval module to learn domain-specific knowledge in time and introduces progressive prompting and aggregation module for considering diverse aspects of the relevance and prediction stability. We explore an industrial implementation to deploy LLMs to handle full-scale search traffics of Alipay with acceptable cost and latency. The comprehensive experiments on real-world industry data and online A/B testing validate the superiority of our proposal and the effectiveness of its main modules.

<sup>¶</sup>the number of valid clicks divided by the number of searches

## References

- Qian Chen, Xiaodan Zhu, Zhenhua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2016. Enhanced lstm for natural language inference. *arXiv preprint arXiv:1609.06038*.
- Zeyuan Chen, Wei Chen, Jia Xu, Zhongyi Liu, and Wei Zhang. 2023. Beyond semantics: Learning a behavior augmented relevance model with self-supervised learning. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 4516–4522.
- 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*.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. Glm: General language model pretraining with autoregressive blank infilling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 320–335.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.
- Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional neural network architectures for matching natural language sentences. *Advances in neural information processing systems*, 27.
- Hanqi Jin, Jiwei Tan, Lixin Liu, Lisong Qiu, Shaowei Yao, Xi Chen, and Xiaoyi Zeng. 2023. Msra: A multi-aspect semantic relevance approach for e-commerce via multimodal pre-training. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 3988–3992.
- Chaozhuo Li, Bochen Pang, Yuming Liu, Hao Sun, Zheng Liu, Xing Xie, Tianqi Yang, Yanling Cui, Liangjie Zhang, and Qi Zhang. 2021. Adsgnn: Behavior-graph augmented relevance modeling in sponsored search. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 223–232.
- Zihan Liao, Hang Yu, Jianguo Li, Jun Wang, and Wei Zhang. 2024. D2llm: Decomposed and distilled large language models for semantic search. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14798–14814.
- Hong Liu, Saisai Gong, Yixin Ji, Kaixin Wu, Jia Xu, and Jinjie Gu. 2024. Boosting llm-based relevance modeling with distribution-aware robust learning. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 4718–4725.
- Wenhao Lu, Jian Jiao, and Ruofei Zhang. 2020. Twinbert: Distilling knowledge to twin-structured compressed bert models for large-scale retrieval. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 2645–2652.
- Hamid Palangi, Li Deng, Yelong Shen, Jianfeng Gao, Xiaodong He, Jianshu Chen, Xinying Song, and R Ward. 2014. Semantic modelling with long-short-term memory for information retrieval. *arXiv preprint arXiv:1412.6629*.
- Bochen Pang, Chaozhuo Li, Yuming Liu, Jianxun Lian, Jianan Zhao, Hao Sun, Weiwei Deng, Xing Xie, and Qi Zhang. 2022. Improving relevance modeling via heterogeneous behavior graph learning in bing ads. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3713–3721.
- Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2016. Text matching as image recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.
- Ankur P Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. *arXiv preprint arXiv:1606.01933*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. In-context retrieval-augmented language models. *arXiv preprint arXiv:2302.00083*.
- Jinfeng Rao, Linqing Liu, Yi Tay, Wei Yang, Peng Shi, and Jimmy Lin. 2019. Bridging the gap between relevance matching and semantic matching for short text similarity modeling. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5370–5381.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.



- Timo Schick and Hinrich Schütze. 2020. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676*.
- Chirag Shah and Jefferey Pomerantz. 2010. Evaluating and predicting answer quality in community qa. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 411–418.
- Yelong Shen, Xiaodong He, Jianfeng Gao, Li Deng, and Grégoire Mesnil. 2014. Learning semantic representations using convolutional neural networks for web search. In *Proceedings of the 23rd international conference on world wide web*, pages 373–374.
- Weiwei Sun, Lingyong Yan, Xinyu Ma, Pengjie Ren, Dawei Yin, and Zhaochun Ren. 2023. Is chatgpt good at search? investigating large language models as re-ranking agent. *arXiv preprint arXiv:2304.09542*.
- Krysta M Svore and Christopher JC Burges. 2009. A machine learning approach for improved bm25 retrieval. In *Proceedings of the 18th ACM conference on Information and knowledge management*, pages 1811–1814.
- Kai Sheng Tai, Richard Socher, and Christopher D Manning. 2015. Improved semantic representations from tree-structured long short-term memory networks. *arXiv preprint arXiv:1503.00075*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Jianing Wang, Chengyu Wang, Chuanqi Tan, Jun Huang, and Ming Gao. 2023. Boosting in-context learning with factual knowledge. *arXiv preprint arXiv:2309.14771*.
- Wei Wang, Bin Bi, Ming Yan, Chen Wu, Zuyi Bao, Jiangnan Xia, Liwei Peng, and Luo Si. 2019. Structbert: Incorporating language structures into pre-training for deep language understanding. *arXiv preprint arXiv:1908.04577*.
- Chen Wei, Yixin Ji, Zeyuan Chen, Jia Xu, and Zhongyi Liu. 2024. Llmgr: Large language model-based generative retrieval in alipay search. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2847–2851.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Shaowei Yao, Jiwei Tan, Xi Chen, Keping Yang, Rong Xiao, Hongbo Deng, and Xiaojun Wan. 2021. Learning a product relevance model from click-through data in e-commerce. In *Proceedings of the Web Conference 2021*, pages 2890–2899.
- Shaowei Yao, Jiwei Tan, Xi Chen, Juhao Zhang, Xiaoyi Zeng, and Keping Yang. 2022. Reprbert: Distilling bert to an efficient representation-based relevance model for e-commerce. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4363–4371.
- Zhiyuan Zeng, Yuzhi Huang, Tianshu Wu, Hongbo Deng, Jian Xu, and Bo Zheng. 2022. Graph-based weakly supervised framework for semantic relevance learning in e-commerce. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 3634–3643.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*.
- Jason Zhu, Yanling Cui, Yuming Liu, Hao Sun, Xue Li, Markus Pelger, Tianqi Yang, Liangjie Zhang, Ruofei Zhang, and Huasha Zhao. 2021. Textgnn: Improving text encoder via graph neural network in sponsored search. In *Proceedings of the Web Conference 2021*, pages 2848–2857.

## A Appendix

### A.1 Industrial Implementation for LLMs

In real-world search scenarios (e.g., Alipay search), it is hard to deploy LLMs to handle all the search traffics with acceptable cost and latency. However, the judgment of relevance is objective and non-personalized. This implies that the relevance scores for the same query-item pair should remain consistent for all users, unlike recommendation algorithms that are personalized for each individual. As long as we can obtain the relevance scores for the query-item pairs, we can perform online services for all users, which greatly reduces the volume of online requests. And for the certain query or item, their semantic information is relatively stable and does not vary much. Inspired by this, we come to a efficient and effective solution (i.e., online and offline collaborative service) with the affordable cost and latency. Figure 4 concretely illustrates online serving process of our proposed online and offline collaborative service when a certain query-item pair appears.

Specifically, we utilize LLMs (e.g., 10B parameters) to perform offline inference on daily search




Figure 4: The illustration of online and offline collaborative service.

logs exploiting the daily prompts and update relevance scores stored in the database. The cost of offline inference is minimal, so we can choose LLMs with a larger number of parameters to enhance the accuracy of relevance score calculations. Meanwhile, we would deploy distilled smaller LLMs (e.g., 2B parameters) for online serving. Since the method of distillation is not the focus of this paper, we have omitted the introduction of the method here. With this model, it is possible to make online predictions for query-item pairs that do not have stored scores in the database. When a user enter a query in Alipay search, the service will firstly look up the database from the offline service and obtain relevance scores based query and candidate items. Once the relevance scores of the query-item pair can not be obtained from the database, the online service will be requested for predicting the relevance score timely. Note that approximately 95% of the traffic will be handled by the offline service, while the remaining traffic will be handled by the online service. Through this method, we have achieved affordable efficiency and resource requirements for exploiting LLMs to handle relevance judgements of industrial scenarios.

## A.2 Baseline Models

- **DSSM** is a two-tower representation-based model. It encodes the embedding of a given query and item independently and computes the relevance score accordingly.
- **ReprBert** is a representation-based Bert model that utilizes novel interaction strategies to achieve a balance between representation interactions and model latency.
- **Bert** has achieved significant progress on NLP tasks as an interaction-based model. Here we concatenate the query and item as the input of

the model.

- **MASM** leverages the historical behavior data to complete model pre-training as a weak-supervision signal with a newly proposed training objective.
- **TextGNN** extends the two-tower model with the complementary graph information from user historical behaviors.
- **AdsGNN** further proposes three aggregation methods for the user behavior graph from different perspectives.
- **BARL-ASe** proposes dual behavior-neighbors augmented relevance model with self-supervised learning. And it exploits LLMs to deal with long-tailed query-item pairs.
- **BLOOM** is a language model trained on 46 natural languages and 13 programming languages with the causal decoder architecture.
- **GLM** is a language model with both powerful natural language understanding and generation capabilities based on the prefix decoder.