Breaking the Length Barrier: LLM-Enhanced CTR Prediction in Long Textual User Behaviors

Binzong Geng* Ant Group Hangzhou, China gengbinzong.gbz@antgroup.com

> Yong He Ant Group Hangzhou, China heyong.h@antgroup.com

Zhaoxin Huan* Ant Group Hangzhou, China zhaoxin.hzx@antgroup.com

Liang Zhang Ant Group Hangzhou, China zhuyue.zl@antfin.com

Xiaolu Zhang Ant Group Hangzhou, China yueyin.zxl@antfin.com

Fajie Yuan Westlake University Hangzhou, China yuanfajie@westlake.edu.cn

Linjian Mo[†] Ant Group Hangzhou, China linyi01@antgroup.com

ABSTRACT

With the rise of large language models (LLMs), recent works have leveraged LLMs to improve the performance of click-through rate (CTR) prediction. However, we argue that a critical obstacle remains in deploying LLMs for practical use: the efficiency of LLMs when processing long textual user behaviors. As user sequences grow longer, the current efficiency of LLMs is inadequate for training on billions of users and items. To break through the efficiency barrier of LLMs, we propose Behavior Aggregated Hierarchical Encoding (BAHE) to enhance the efficiency of LLM-based CTR modeling. Specifically, BAHE proposes a novel hierarchical architecture that decouples the encoding of user behaviors from inter-behavior interactions. Firstly, to prevent computational redundancy from repeated encoding of identical user behaviors, BAHE employs the LLM's pretrained shallow layers to extract embeddings of the most granular, atomic user behaviors from extensive user sequences and stores them in the offline database. Subsequently, the deeper, trainable layers of the LLM facilitate intricate inter-behavior interactions, thereby generating comprehensive user embeddings. This separation allows the learning of high-level user representations to be independent of low-level behavior encoding, significantly reducing computational complexity. Finally, these refined user embeddings, in conjunction with correspondingly processed item embeddings, are incorporated into the CTR model to compute the CTR scores. Extensive experimental results show that BAHE reduces training

Jun Zhou[†]

Ant Group

Hangzhou, China

jun.zhoujun@antfin.com

Conference'17, July 2017, Washington, DC, USA

© 2024 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-x/YY/MM...\$15.00

https://doi.org/10.1145/nnnnnnnnnnn

time and memory by five times for CTR models using LLMs, especially with longer user sequences. BAHE has been deployed in a real-world system, allowing for daily updates of 50 million CTR data on 8 A100 GPUs, making LLMs practical for industrial CTR prediction.

KEYWORDS

Large Language Models; Click-through rate prediction; Efficiency

ACM Reference Format:

Binzong Geng, Zhaoxin Huan, Xiaolu Zhang, Yong He, Liang Zhang, Fajie Yuan, Jun Zhou, and Linjian Mo. 2024. Breaking the Length Barrier: LLM-Enhanced CTR Prediction in Long Textual User Behaviors. In Proceedings of ACM Conference (Conference'17). ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/nnnnnnnnnnnn

1 INTRODUCTION

Click-through rate (CTR) prediction plays a significant role in various domains, including online advertising, search engines, and recommendation systems. Recently, large language models (LLMs) have achieved remarkable results across various domains [3, 11, 13, 17, 20, 26]. Numerous studies have investigated the application of LLMs in CTR prediction, leveraging LLMs' powerful semantic understanding and world knowledge to enhance CTR modeling. For example, M6-Rec [5] reconstructs the interactions between users and items in textual prompts and utilizes M6 to conduct several recommendation tasks including CTR prediction. Both CTRL [14] and FLIP [23] integrate semantic information from LLMs into traditional ID-based models [25] using contrastive learning and masked language modeling. KAR [24] develops a three-stage framework based on the reasoning and factual knowledge of LLMs, transferring knowledge from LLMs into the CTR model through embeddings. These works demonstrate the tremendous potential and broad prospects for LLMs in enhancing CTR prediction.

However, we argue that a challenge for LLMs in practical CTR prediction remains unresolved: the efficiency bottleneck of LLMs

^{*}Both authors contributed equally to this research. [†]Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

when dealing with long user behavior sequences. This issue significantly hampers the real-world application of LLMs. In traditional CTR modeling, it is widely known that integrating diverse, long-term user sequences into the model can enhance its performance [4, 18, 19, 27]. However, incorporating longer texts into LLMbased CTR models may improve performance but also introduces a bottleneck, as extended texts significantly slow down training and inference, making such models unsuitable for large-scale deployment. This reason has led the LLM-based CTR methods mentioned before to compromise by using smaller language models and shorter user sequences [14, 15, 23, 24]. Moreover, some LLM-based recommendation methods are not directly tailored to CTR tasks, such as sequential recommendation [6, 8, 9, 12, 13, 16, 22]. These methods are more focused on item matching or retrieval, where the user sequences only consist of items that have been interacted with. As a result, the issue of long user sequences is relatively less severe in these works.

To address these challenges, in this paper, we propose Behavior Aggregated Hierarchical Encoding (BAHE) to tackle the performance bottleneck of LLMs with long user sequences, facilitating the application of LLMs in real-world CTR prediction. Specifically, BAHE proposes a novel hierarchical architecture that decouples the encoding of user behaviors from inter-behavior interactions. Firstly, to understand the semantic meaning of each user behavior and to prevent computational redundancy from repeated encoding of identical user behavior, BAHE employs the LLM's pre-trained low layers to extract embeddings of the most granular, atomic user behaviors from extensive user sequences and stores them in the offline database. In this way, BAHE converts the encoding from token-level to behavior-level, substantially reducing sequence length and enhancing the reusability of behavior. Subsequently, BAHE retrieves all of a user's behaviors from the offline restored atomic behavior database and utilizes the deeper, trainable layers of the LLM to facilitate intricate inter-behavior interactions, thereby learning user preferences and generating comprehensive user embeddings. This separation allows the learning of high-level user representations to be independent of low-level behavior encoding, significantly reducing computational complexity. Finally, these refined user embeddings, in conjunction with correspondingly processed item embeddings, are incorporated into the CTR model to compute the CTR scores.Significantly, atomic behaviors rarely undergo changes, which permits infrequent updates within the lower layers of the LLM. These updates can proceed asynchronously relative to those in the higher layers, thereby boosting computational efficiency. Furthermore, the input sequence for the LLM's higher layers is significantly compacted, reducing from the original token count to a smaller number of behaviors. This concise representation not only shortens the sequence length the LLM needs to process but also enhances the model's overall efficiency.

BAHE has been successfully implemented in real-world industrial CTR prediction, where the training time for LLM-based CTR with over 50 million data has been reduced from the initial 5 days to just 1 day, enabling daily model updates and scheduling. Overall, our main contributions are as follows:

• We investigate a significant and unresolved problem: the efficiency bottleneck in LLM-based CTR modeling with long user sequences. We find that the main reasons for the performance bottleneck due to the repetitive encoding of user behaviors and the strong coupling between behavior representation extraction and behavior interaction modeling.

- We propose the BAHE method, which is a novel hierarchical structure to LLM by decoupling the representation extraction of atomic behaviors from the learning of behavior interactions. BAHE solves redundant representations of the same behavior across different users and significantly reduces the length of input sequences.
- Extensive offline experiments validate that BAHE significantly enhances the efficiency of LLM-based CTR models. Moreover, online experiments fully demonstrate BAHE's ability to reduce computational resources in real-world industrial CTR prediction.

2 PROPOSED METHOD

2.1 **Problem Definition**

LLM-based CTR prediction aims to estimate the probability of a user clicking on an item based on textual features. For each user $u \in U$, the user's textual behavior sequence is s_u , with the total token length of s_u being $l_u \cdot s_u$ is composed of N behavior sequences from different domains (such as click, favorite, add-to-cart actions) denoted as $s_u = [s_{un}|0 < n < N]$, with each sequence s_{un} containing an average of M user behaviors $s_{un} = [a_{unm}|0 < m < M]$, and each behavior a_{unm} having an average of K tokens, which means $l_u = N \times M \times K$. For an item $i \in I$, the textual features of i is t_i , comprised of the item's title, with the overall length l_i , which is significantly shorter than l_u . We define H to be the set of distinct atomic behaviors across all users:

$$H = \{a_{umn} | u \in U, 0 < m < M, 0 < n < N\}$$
(1)

Each user's behavior sequence s_u consists of these atomic behaviors and includes the titles of items they've interacted with, hence $t_i \in H$. The objective of LLM-based CTR modeling is to minimize L:

$$L = \frac{1}{|D|} \sum_{u \in U, i \in I} l(Y, F(s_u, t_i))$$
(2)

where l represents the loss function, |D| denotes the total number of training samples, and Y indicates the click label.

2.2 Behavior Aggregated Hierarchical Encoding

Given users *i* and *j*, with atomic behavior sequences $s_i = [a_1, a_2, a_3]$ and $s_j = [a_3, a_1, a_2]$ respectively, Traditional LLM-based CTR modeling struggle with efficiency due to:

- **Redundant behavior encoding**: The same behaviors are redundantly encoded across different users' sequences. For example, s_i and s_j both contain a_1 , a_2 , and a_3 , causing unnecessary repetition in encoding and a waste of computational resources.
- **Tight coupling**: Behaviors like *a*₁, *a*₂, and *a*₃ have fixed meanings, while their sequence varies per user. Existing approaches couple representation extraction with sequence understanding, causing regular, costly updates when behaviors change.

To address these issues, we propose the Behavior Aggregated Hierarchical Encoding (BAHE) approach, and Figure 1 shows the architecture of BAHE. BAHE upgrades the LLM's encoding from Breaking the Length Barrier: LLM-Enhanced CTR Prediction in Long Textual User Behaviors



Figure 1: Architecture of the proposed BAHE method.

token-level to behavior-level, dramatically reducing the encoding length and increasing the cross-user reusability of behaviors. Additionally, BAHE stratifies the LLM and decouples representation extraction from behavior interaction, boosting computational efficiency and adaptability.

2.2.1 Atomic Behavior Encoding (ABE). BAHE first encodes all behaviors from the set H using the LLM's pre-trained low layers, denoted as $LLM_{L_{low}}$, then stores them offline to serve as a behavior embedding table E. Subsequently, the LLM's high layers will utilize the E as a replacement for the original token embedding table to learn interactions among user behaviors. The encoding of atomic behaviors is illustrated as follows:

$$E_{a_{i}} = F_{p}(LLM_{L_{low}}(a_{i})) \quad E = \{a_{i} : E_{a_{i}} | a_{i} \in H\}$$
(3)

where a_i is an atomic behavior defined in H, composed of K text tokens. $LLM_{L_{low}}(a_i) \in \mathbb{R}^{K \times d}$ corresponds to the L_{low} hidden states, where d is the dimension. $F_p : \mathbb{R}^{K \times d} \Rightarrow \mathbb{R}^d$ is the pooling function. The behavior embedding table is $E \in \mathbb{R}^{|H| \times d}$, where |H| is the total number of atomic behaviors. In this way, BAHE transforms encoding from the token level to the behavior level, substantially reducing the encoding length from the number of tokens to the number of atomic behaviors.

2.2.2 Behavior Aggregation (BA). After obtaining the atomic behavior embedding table *E*, BAHE retrieves the corresponding E_{a_i} for each a_i in user *u*'s n-th sequence s_{un} , and then concatenates them as representation for s_{un} :

$$E(s_{un}) = [E(a_{un1}) \oplus E(a_{un2}) \oplus \ldots \oplus E(a_{unM})]$$
(4)

Conference'17, July 2017, Washington, DC, USA

BAHE next employs the LLM's higher layers $LLM_{L_{high}}$ to model the interaction between behaviors, and obtains the overall representation for user sequence s_{un} :

$$Q_{un} = F_d(F_p(LLM_{L_{high}}(E(s_{un}))))$$
(5)

 $Q_{un} \in \mathbb{R}^d$ is the representation of the n-th sequence. $F_d : \mathbb{R}^d \Rightarrow \mathbb{R}^d$ is a dimensionality reduction function that aims to transform highdimensional LLM hidden states into lower-dimensional, facilitating its use in subsequent models.

2.2.3 *Feature Parallel (FP).* To avoid the exponential growth in LLM's attention computations as the number of user sequences N increases, BAHE utilizes a parallel and independent way to process each user sequence through the $LLM_{L_{high}}$ and concatenate them to get the final user representation:

$$Q_u = [Q_{u1} \oplus Q_{u2} \oplus \ldots \oplus Q_{uN}] \tag{6}$$

Following the same method described above, BAHE can also obtain the representation of the item, denoted as Q_i .

2.2.4 *CTR Modeling.* After obtaining the user representation Q_u and item representation Q_i , BAHE feeds their concatenation into the CTR model F_{θ} to compute the final CTR score y:

$$y = F_{\theta}(Q_u \oplus Q_i) \tag{7}$$

BAHE is model-agnostic, which can apply to any embedding-based CTR model. For illustrative purposes, BAHE chose the simple but efficient DNN (Deep Neural Network) as default. BAHE optimizes the loss function defined in Equation 2 to learn the final LLM-based CTR model. After training, BAHE uses Q_u and Q_i to improve downstream model performance by providing additional semantic knowledge [24].

2.3 Complexity Analysis

We analyze the time complexity of BAHE against traditional LLMbased CTR models. With *N* textual behavior features per user, *M* behaviors per sequence, *H* atomic actions, and *K* tokens per action. The original complexity is $O(L(NMK)^2)$, where *L* is the number of layers in the LLM. BAHE splits this into two stages: atomic behavior encoding $O(L_{low}(HK^2))$, and behavior aggregation with feature parallel $O(L_{high}(NM^2))$, where $L = L_{low} + L_{high}$. The efficiency improvement brought by BAHE is $O(L_{low}((N^2M^2 - H)K^2)) + O(L_{high}((N^2K^2 - N)M^2))$.

BAHE simplifies the LLM's lower-level encoding, lowering complexity by $(N^2M^2 - H)K^2$ as it encodes just H atomic behaviors instead of redundantly processing N^2M^2 identical user behaviors. Furthermore, for LLM's higher-level semantic understanding, BAHE processes N sequence features in parallel and employs behavior aggregation, cutting computations by $(N^2K^2 - N)M^2$.

3 EXPERIMENT

3.1 Experimental Setup

3.1.1 Dataset. We evaluated BAHE using a real-world industrial dataset with around 50 million CTR records collected over a week. The data, split into training, validation, and test sets by log time, includes 6 text features like user bills, searches, and mini-program

Table 1: Performance of BAHE at different text lengths. "AUC" represents the modeling performance of LLM, while "AUC_d" indicates the performance when transferring LLM representations to downstream models.

	Text Length=1024			Text Length=2048			
Model	AUC	GPU-h	Mem(G)	AUC	GPU-h	Mem(G)	AUC_d
DNN	-	-	-	-	-	-	0.7219
LLM-CTR	0.7161	448	43.8	0.7276	928	75.4	0.7323
+ FP	0.7143	420	36.8	0.7326	864	67.8	0.7369
+ ABE	0.7150	256	23.0	0.7332	416	38.0	0.7372
+ BA(BAHE)	0.7132	116	9.8	0.7309	164	12.6	0.7352

visits, along with item titles. Each user sequence has 50 user behaviors, averaging 5 tokens each, summing up to 10 million atomic behaviors. The label indicates whether the user clicked on the item.

3.1.2 Baseline Methods. The baseline methods select mainstream LLM-based CTR modeling (denoted as LLM-CTR) [7, 14, 23, 24], which concatenate multiple user text sequences into single long sequences, and then perform CTR modeling based on the LLM. Since BAHE is model-agnostic, it can be applied to any LLM-based models.Furthermore, to test the enhancement that LLM provides to the downstream CTR model, we also present the performance of the downstream CTR model, which is denoted as DNN.

3.1.3 Evaluation Metrics. We assess performance using the Area Under the Curve (AUC) [2], standard for CTR. Efficiency is evaluated by training GPU hours (GPU-h) [21] and memory usage. Furthermore, as mentioned in Section 2.2.4, we also show AUC_d to demonstrate LLM's benefits for downstream CTR models.

3.1.4 Experimental Details. Our backbone LLM leverages the opensource Qwen-1_8B¹ [1]. As detailed in Section 2.2, we employ an MLP with [2048, 512, 128] units as F_d to condense LLM embeddings, and another MLP as F_{θ} with [256, 32, 1] units for CTR predictions, using mean pooling as default F_p . We fine-tune using Lora Tuning [6, 10] at rank 64, batch size 16, one training epoch, and a 5e-5 learning rate with cosine decay. All tests run on 8 A100 GPUs.

3.2 Performance

3.2.1 Offline Performance. Table 1 compares BAHE with the baseline. The findings are: Firstly, BAHE cuts training time by 5x, from 928 to 164 GPU hours, and slashes GPU memory to a sixth of the baseline. This marks a significant improvement in the efficiency of LLMs when processing long sequences. Secondly, BAHE's reuse of behavior representations improves key behavior capture, substantially raising AUC for both LLMs and downstream models. BAHE thus strikes an effective balance between efficiency and performance. Lastly, the individual components of BAHE-Feature Parallel (FP), Atomic Behavior Encoding (ABE), and Behavior Aggregation (BA)-contributes to these improvements, highlighting their importance. In summary, BAHE enhances both efficiency and effectiveness for LLMs in long sequence applications, paving the way for their practical use.



Figure 2: (a) presents a comparison of the normalized training time and memory usage (left y-axis) and the AUC (right yaxis) of BAHE with the baseline for different user sequence lengths (x-axis). Figure (b) shows the AUC comparison of BAHE utilizing different behavior pooling methods.

3.2.2 Online Deployment. To demonstrate BAHE's practical benefits, we deploy BAHE on a large-scale e-commerce platform's advertisement CTR prediction and conduct a two-week A/B test. BAHE allowed daily LLM training on 50 million CTR records, outperforming the baseline model's weekly training capacity. Consequently, BAHE unlocks a greater potential for LLM-based CTR, resulting in a 9.65% increase in online CTR and 2.41% rise in advertising CPM.

3.3 Empirical Analysis

3.3.1 Comparison at Different Sequence Lengths. To give a more detailed analysis of the performance enhancements attributable to BAHE, Figure 2(a) presents a comparison between BAHE and the baseline at different sequence lengths. Firstly, AUC improves with longer user sequences, showing that extended texts enhance LLM-based CTR. Secondly, BAHE's performance boosts are greater for longer texts, indicating its effectiveness in handling increased lengths.

3.3.2 Comparison of Different Pooling Methods. During the extraction of atomic behavioral representations and subsequent behavior aggregation, the choice of pooling method (denoted as F_p in Equations 3 and 5) plays a crucial role. We evaluate two prevalent techniques in Figure 2(b): mean pooling and the last hidden state of the LLM (denoted as EOS).Our findings show mean pooling is superior to EOS, indicating global representations are more effective than the last hidden embedding for generative LLMs.

4 CONCLUSION

To tackle the challenge of efficiency in LLM-based CTR models when dealing with users' extensive text sequences, this paper proposes a novel BAHE method. BAHE enhances the reusability of behavioral representations across users by encoding atomic behaviors. It employs LLM's hierarchical encoding technique to separate the learning of behavior representations from the inter-behavior modeling, thereby boosting computational efficiency. Extensive online and offline experimental results demonstrate that BAHE not only achieves a more than 5 times increase in efficiency but also enhances the CTR performance, offering fresh insights for the practical deployment of LLM-based CTR models.

¹https://huggingface.co/Qwen/Qwen-1_8B

Breaking the Length Barrier: LLM-Enhanced CTR Prediction in Long Textual User Behaviors

REFERENCES

- [1] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. Qwen Technical Report. arXiv preprint arXiv:2309.16609 (2023).
- [2] Andrew P Bradley. 1997. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern recognition* 30, 7 (1997), 1145–1159.
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. Advances in neural information processing systems 33 (2020), 1877–1901.
- [4] Yue Cao, Xiaojiang Zhou, Jiaqi Feng, Peihao Huang, Yao Xiao, Dayao Chen, and Sheng Chen. 2022. Sampling Is All You Need on Modeling Long-Term User Behaviors for CTR Prediction. In Proceedings of the 31st ACM International Conference on Information & Knowledge Management, Atlanta, GA, USA, October 17-21, 2022. ACM.
- [5] Zeyu Cui, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. M6-rec: Generative pretrained language models are open-ended recommender systems. arXiv preprint arXiv:2205.08084 (2022).
- [6] Junchen Fu, Fajie Yuan, Yu Song, Zheng Yuan, Mingyue Cheng, Shenghui Cheng, Jiaqi Zhang, Jie Wang, and Yunzhu Pan. 2023. Exploring Adapter-based Transfer Learning for Recommender Systems: Empirical Studies and Practical Insights. arXiv preprint arXiv:2305.15036 (2023).
- [7] Zichuan Fu, Xiangyang Li, Chuhan Wu, Yichao Wang, Kuicai Dong, Xiangyu Zhao, Mengchen Zhao, Huifeng Guo, and Ruiming Tang. 2023. A Unified Framework for Multi-Domain CTR Prediction via Large Language Models. arXiv preprint arXiv:2312.10743 (2023).
- [8] Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5). In Proceedings of the 16th ACM Conference on Recommender Systems. 299–315.
- [9] Yupeng Hou, Shanlei Mu, Wayne Xin Zhao, Yaliang Li, Bolin Ding, and Ji-Rong Wen. 2022. Towards universal sequence representation learning for recommender systems. In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 585–593.
- [10] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022.* OpenReview.net.
- [11] Jie Huang and Kevin Chen-Chuan Chang. 2023. Towards Reasoning in Large Language Models: A Survey. In Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023. Association for Computational Linguistics.
- [12] Jiacheng Li, Ming Wang, Jin Li, Jinmiao Fu, Xin Shen, Jingbo Shang, and Julian J. McAuley. 2023. Text Is All You Need: Learning Language Representations for Sequential Recommendation. In Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2023, Long Beach, CA, USA, August 6-10, 2023. ACM.
- [13] Ruyu Li, Wenhao Deng, Yu Cheng, Zheng Yuan, Jiaqi Zhang, and Fajie Yuan. 2023. Exploring the Upper Limits of Text-Based Collaborative Filtering Using Large Language Models: Discoveries and Insights. arXiv preprint arXiv:2305.11700 (2023).
- [14] Xiangyang Li, Bo Chen, Lu Hou, and Ruiming Tang. 2023. CTRL: Connect Tabular and Language Model for CTR Prediction. arXiv preprint arXiv:2306.02841 (2023).
- [15] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019).
- [16] Zhenghao Liu, Sen Mei, Chenyan Xiong, Xiaohua Li, Shi Yu, Zhiyuan Liu, Yu Gu, and Ge Yu. 2023. Text Matching Improves Sequential Recommendation by Reducing Popularity Biases. In Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. 1534–1544.
- [17] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. Advances in Neural Information Processing Systems 35 (2022), 27730–27744.
- [18] Qi Pi, Weijie Bian, Guorui Zhou, Xiaoqiang Zhu, and Kun Gai. 2019. Practice on Long Sequential User Behavior Modeling for Click-Through Rate Prediction. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019. ACM.

- [19] Qi Pi, Guorui Zhou, Yujing Zhang, Zhe Wang, Lejian Ren, Ying Fan, Xiaoqiang Zhu, and Kun Gai. 2020. Search-based User Interest Modeling with Lifelong Sequential Behavior Data for Click-Through Rate Prediction. In CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020. ACM.
- [20] 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 (2019), 9.
- [21] Teven Le Scao, Thomas Wang, Daniel Hesslow, Stas Bekman, M. Saiful Bari, Stella Biderman, Hady Elsahar, Niklas Muennighoff, Jason Phang, Ofir Press, Colin Raffel, Victor Sah, Sheng Shen, Lintang Sutawika, Jaesung Tae, Zheng Xin Yong, Julien Launay, and Iz Beltagy. 2022. What Language Model to Train if You Have One Million GPU Hours?. In Findings of the Association for Computational Linguistics: EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022. Association for Computational Linguistics.
- [22] Kyuyong Shin, Hanock Kwak, Su Young Kim, Max Nihlén Ramström, Jisu Jeong, Jung-Woo Ha, and Kyung-Min Kim. 2023. Scaling law for recommendation models: Towards general-purpose user representations. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 37. 4596–4604.
- [23] Hangyu Wang, Jianghao Lin, Xiangyang Li, Bo Chen, Chenxu Zhu, Ruiming Tang, Weinan Zhang, and Yong Yu. 2023. FLIP: Towards Fine-grained Alignment between ID-based Models and Pretrained Language Models for CTR Prediction. arXiv e-prints (2023), arXiv-2310.
- [24] Yunjia Xi, Weiwen Liu, Jianghao Lin, Jieming Zhu, Bo Chen, Ruiming Tang, Weinan Zhang, Rui Zhang, and Yong Yu. 2023. Towards Open-World Recommendation with Knowledge Augmentation from Large Language Models. arXiv preprint arXiv:2306.10933 (2023).
- [25] Zheng Yuan, Fajie Yuan, Yu Song, Youhua Li, Junchen Fu, Fei Yang, Yunzhu Pan, and Yongxin Ni. 2023. Where to go next for recommender systems? id-vs. modality-based recommender models revisited. arXiv preprint arXiv:2303.13835 (2023).
- [26] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. A Survey of Large Language Models. CoRR (2023).
- [27] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining. 1059–1068.