

Uniboost: Global Coordination with Value Alignment for Fair and Efficient Traffic Allocation

Ge Fan
ge.fan@outlook.com
Taobao & Tmall Group of Alibaba
Hangzhou, China

Cong Luo
hengying.lc@taobao.com
Taobao & Tmall Group of Alibaba
Hangzhou, China

Jialin Liu
miao.ljl@taobao.com
Taobao & Tmall Group of Alibaba
Hangzhou, China

Nan Zhao
yueqi.zn@taobao.com
Taobao & Tmall Group of Alibaba
Hangzhou, China

Yang Fu
yang.fy@taobao.com
Taobao & Tmall Group of Alibaba
Hangzhou, China

Yuning Jiang*
mengzhu.jyn@taobao.com
Taobao & Tmall Group of Alibaba
Beijing, China

Kai Meng
mengkai.meng@taobao.com
Taobao & Tmall Group of Alibaba
Hangzhou, China

Huiping Chu*
yueqi.chp@taobao.com
Taobao & Tmall Group of Alibaba
Hangzhou, China

Bo Zheng
bozheng@alibaba-inc.com
Taobao & Tmall Group of Alibaba
Beijing, China

Abstract

With the rapid evolution of internet services, recommendation systems have become indispensable. In particular, the blending (re-ranking) stage plays a pivotal role in allocating traffic across diverse business objectives. However, existing approaches often suffer from coupled allocation plans, score inflation, and a lack of interpretability. To address these challenges, we propose Uniboost, a unified traffic allocation framework. Uniboost introduces a posterior value alignment mechanism that calibrates abstract model scores to anchor metrics with explicit business semantics, significantly enhancing interpretability. Furthermore, it employs an independent linear boosting paradigm to decouple complex weighting schemes, enabling precise attribution of each plan's contribution. We validate the effectiveness of Uniboost through online A/B tests and in-depth data analysis, demonstrating three key findings: 1) Reducing the overall weight of weighted scores effectively mitigates unintended business interference, yielding a more efficient micro-level traffic allocation strategy; 2) Post-hoc analyses and aggregated dashboards provide intuitive, macro-level insights that guide the design of the overall traffic allocation mechanism; 3) The proposed "Effective Completion Score" serves as an easily obtainable post-metric that offers a reliable anchor for content recommendation pipelines. Collectively, our experiments show that Uniboost not only improves traffic allocation efficiency and recommendation performance at the micro level but also provides macro-level guidance for system iteration. Thus, this work provides an efficient and controllable traffic regulation solution for large-scale industrial recommendation systems.

*Corresponding Author.



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

SIGIR '26, Melbourne, VIC, Australia

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2599-9/2026/07

<https://doi.org/10.1145/3805712.3808411>

CCS Concepts

• **Information systems** → **Information retrieval**; **Data mining**; **Recommender systems**; *Summarization*; • **Human-centered computing** → *Interactive systems and tools*.

Keywords

Traffic Allocation, Recommendation System, Value Alignment

ACM Reference Format:

Ge Fan, Nan Zhao, Kai Meng, Cong Luo, Yang Fu, Huiping Chu, Jialin Liu, Yuning Jiang, and Bo Zheng. 2026. Uniboost: Global Coordination with Value Alignment for Fair and Efficient Traffic Allocation. In *Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '26)*, July 20–24, 2026, Melbourne, VIC, Australia. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3805712.3808411>

1 Introduction

With the rapid advancement of internet technologies, search and recommendation systems have become integral to online applications, driving user engagement and business growth. As the core engine of these systems, recommendation algorithms have undergone significant evolution. Traditional systems relied on collaborative filtering to recommend items based on user behavior similarity [4, 5, 7, 13]. Today, however, industrial-grade recommendation pipelines have evolved into complex multi-stage architectures [8, 15, 16, 21, 25], typically comprising Matching, Ranking, and Blending (Re-ranking) stages, as illustrated in Figure 1. The Matching stage retrieves a broad set of candidate items from a massive pool. The Ranking stage employs sophisticated models to score these candidates precisely, filtering for the most relevant content. Finally, the Blending stage merges candidates from diverse sources and applies traffic allocation mechanisms to select the final content for user exposure.

In practical industrial systems, different content types (e.g., ads vs. organic videos) often operate on independent pipelines due to variations in data distribution and business objectives. For instance, the Ranking system for organic video prioritizes user watch time, whereas the Ranking system for ads focuses on ad revenue

[2, 6, 14, 17–19]. Since ads typically exhibit lower posterior business metrics (e.g., completion rate) compared to organic videos, ranking solely by organic metrics would severely limit ad exposure, reducing commercial revenue [12, 23]. To balance user experience with commercial goals, the Blending system often employs guaranteed delivery mechanisms. Specifically, regulation algorithms (e.g., PID controllers [1, 10]) calculate a regulation score based on historical ad exposure rates and add it to the ad blending score to boost visibility [11, 24]. Furthermore, to maintain a healthy ecosystem and mitigate the Matthew Effect [9, 20, 22], the system must allocate traffic to cold-start content and explore user potential interests. This is typically implemented during the blending stage through boost weighting, where a scaling factor w is applied to the original blending score to increase the exposure likelihood of such content.

However, the proliferation of operational campaigns (e.g., promotions, alliances) introduces a series of short-term traffic allocation plans. In traditional pipelines, these plans accumulated weights cause blending scores to inflate, losing their physical semantics and making them difficult to interpret. Besides, multiple plans are tightly coupled, interfering with each other and hindering the iteration and optimization of the Blending system. While recent works have attempted to address similar issues in ad systems by aligning user value with commercial value [3], these solutions cannot be directly applied to content recommendation systems due to the following challenges:

- (1) **Posterior Target Selection.** Unlike ad systems where the target (revenue) is stable, content recommendation involves complex, multi-faceted goals (e.g., watch time, clicks, comments). The impact of these targets on the Boost framework remains under-explored, complicating the selection of a robust posterior target.
- (2) **Unified Weighting Framework.** Existing weighting schemes are fragmented; ad guaranteed delivery uses PID methods, while cold-start exploration uses Boost weighting. Designing a unified framework that accommodates these diverse mechanisms with minimal disruption to existing business logic is non-trivial.
- (3) **Attribution of Weighting Costs.** Due to the coupling of weighting plans, their effects overlap. Accurately evaluating and tracking the cost and contribution of each individual plan remains a significant system design challenge.

To address these challenges, we propose Uniboost, a novel traffic allocation framework that unifies weighting and regulation for diverse content types within the Blending system.

First, through correlation analysis with various posterior targets, we identify "Effective Completion Rate" as our Anchor Metric. By aligning the distribution of blending scores with this completion rate, we transform abstract scores into values with real-world business semantics without altering the ranking order. Second, based on these aligned scores, we design a Unified Boosting Paradigm with Bias. This paradigm not only covers existing weighting schemes but also ensures all weighted scores retain physical interpretability. Uniboost calculates independent boosting gains for each business plan and aggregates them via linear summation. This enables precise attribution, allowing us to quantify the contribution of each scheme to the final ranking.

In deployment, the online serving stage sorts and outputs content based on these weighted scores to improve micro-level performance. The near-line stage collects weighted and raw scores to update distribution alignment parameters. The offline stage aggregates data to provide macro-level guidance for system iteration.

In summary, our contributions are as follows:

- (1) We propose Uniboost, a unified weighting framework that imbues complex blending scores with business value through posterior value alignment. It decouples complex online weighting schemes, providing reliable quantitative evaluation for each individual plan.
- (2) We design an independent linear boosting paradigm that enables precise attribution of traffic allocation costs. This resolves the score inflation issue and significantly enhances the interpretability of the blending mechanism.
- (3) We deploy Uniboost in a large-scale recommendation system via A/B testing. Results demonstrate that our approach not only improves traffic allocation efficiency at the micro level but also provides actionable guidance for system iteration at the macro level.

2 Proposed Method

This section details the architecture of our proposed Uniboost system. As illustrated in Figure 1, Uniboost is deployed at the Blending stage of the recommendation system. Its primary objective is to support flexible business intervention strategies while maintaining the overall efficiency of the recommendation pipeline. The framework comprises two core components: the Online Serving Pipeline and the Near-line/Offline Statistics System, forming a closed-loop mechanism from strategy execution to performance feedback.

2.1 Online Serving Pipeline

The online serving process constitutes the core of Uniboost, designed to complete the blending of candidate items within millisecond-level latency. The workflow consists of four key stages.

2.1.1 Model Scoring. First, the system aggregates outputs from upstream retrieval and blending stages to form the final candidate set \mathcal{V} . This set is fed into the blending model to generate a raw prediction score y'_v for each candidate item $v \in \mathcal{V}$. This score reflects the model's estimation of the user's overall preference.

2.1.2 Value Alignment. Raw model scores y'_v often lack explicit business semantics, making direct weighted intervention uncontrollable. To address this, we introduce a Value Alignment Module. This module maps the model scores to a specific posterior target space (defined as the "Anchor Target," e.g., "Effective Completion Rate") by utilizing global alignment parameters π . The alignment is formulated as:

$$y_v = F(y'_v | \pi) = \frac{y'_v \mu_{anchor}}{\mu_{score}} \quad (1)$$

where μ_{score} and μ_{anchor} denote the global means of the raw blending scores and the anchor target scores, respectively.

This step calibrates abstract model scores into Expected Values with clear business implications. This ensures that subsequent intervention weights possess interpretable physical significance, bridging the gap between model optimization and business goals.

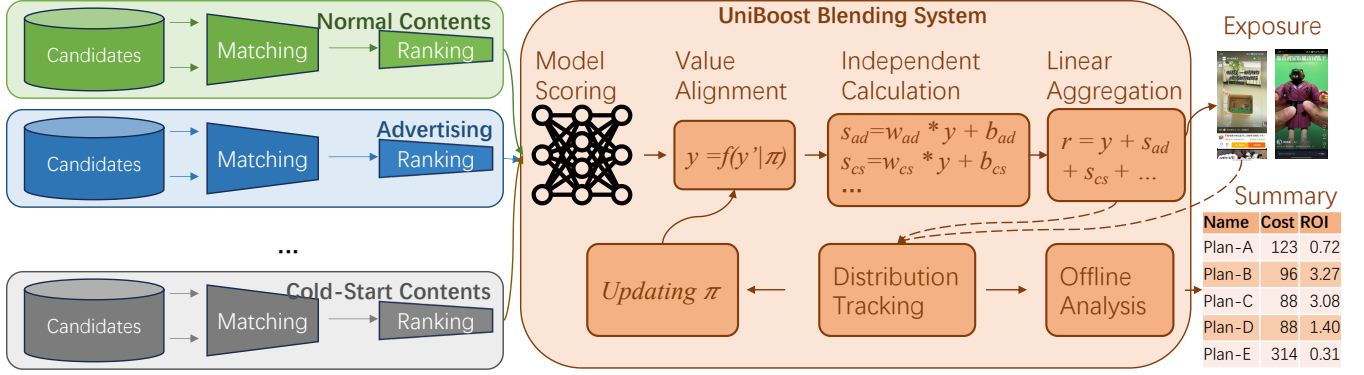


Figure 1: Overview of UniBoost Blending System. The figure illustrates the boosting process for two example content types: advertisements (ad) and cold-start (cs) contents. When $w_{ad} = 0$ and b_{ad} is set to the output score from the PID-based ad delivery controller, the original PID regulation mechanism for ads is preserved. Conversely, when a non-zero weight w_{cs} is assigned and $b_{cs} = 0$ the formulation reduces to the conventional Boosting approach used for cold-start content.

2.1.3 Independent Boosting Calculation. Within the unified value space, the system calculates independent weighted gains for distinct business plans. For any given plan $p \in S_p$, the boost score $s_{p,v}$ for item v is computed as

$$s_{p,v} = \mathbf{1}_p(v) * (w_p * y_v + b_p) \quad (2)$$

where $\mathbf{1}_p(v)$ is the indicator function, taking the value 1 if item v belongs to the plan and 0 otherwise; w_p and b_p are learnable hyperparameters representing the weight coefficient and bias term, respectively.

Since y_v is aligned to the anchor metric, the calculated boost score $s_{p,v}$ shares the same dimensionality as the anchor target. This implies that all business interventions are quantified as specific contributions to the anchor metric, achieving standardized measurement of strategy effects.

2.1.4 Linear Aggregation. Finally, the system fuses the base value and all business plan gains via linear summation to obtain the final ranking score r_v

$$r_v = y_v + \sum_{p \in S_{plan}} s_{p,v} \quad (3)$$

where S_{plan} represents the set of all configured weighting schemes online.

We employ linear summation rather than non-linear fusion to guarantee Additivity. This allows for precise Attribution, enabling us to quantify the contribution of specific schemes to the final ranking result. This feature significantly facilitates downstream strategy tuning and optimization.

Ultimately, the candidate set \mathcal{V} is sorted in descending order based on r_v , truncated, and exposed to the user.

2.2 Near-line/Offline Monitoring System

To ensure system stability and support iterative strategy optimization, UniBoost establishes a comprehensive data statistics and feedback mechanism, consisting of two main phases.

2.2.1 Distribution Tracking & Caching. This module tracks and caches data distributions across key pipeline stages in real-time, including raw blending scores y'_v , aligned value scores y_v and weighted boost scores $s_{p,v}$ for each plan. These fine-grained data points provide critical support for monitoring Distribution Shifts and diagnosing online anomalies.

2.2.2 Aggregation & ROI Analysis. This module aggregates tracking data for downstream strategy evaluation. Specifically, the system statistics the overall resource consumption (e.g., traffic share) and business performance (e.g., improvement in the anchor metric) for each plan. Based on this, we calculate the Return on Investment (ROI) for each weighting plan. For example, suppose we have a plan p aimed at increasing exposure Visit View (VV). Then, the ROI for this plan can be defined as:

$$ROI^{VV} = \frac{\Delta_p^{VV}}{\text{Cost}_p} \quad (4)$$

This metric directly reflects the deployment efficiency of each plan, guiding operators to dynamically adjust parameters. Consequently, this enables globally optimal resource allocation based on data-driven insights.

3 Experiments and Discussion

To comprehensively evaluate the impact of the proposed UniBoost framework on the end-to-end recommendation pipeline, we conducted large-scale online A/B testing to validate its effectiveness from both micro and macro perspectives. Specifically, we established a control group using the original weighted pipeline, in which the scores alignment is performed only after weighting, and a treatment group implemented with the unified UniBoost weighting scheme. We selected the Effective Completion Rate as the anchor metric for content recommendation. Additionally, we performed an in-depth online data analysis to justify the selection of this anchor metric. Overall, this section addresses the following three Research Questions (RQs):

Table 1: Relative Improvement in Online A/B Test.

Metrics	VV	Valued VV	Duration	Valued Score
Overall	+1.69%	+3.07%	+0.65%	+2.54%

- RQ1 Can Uniboost outperform baseline schemes by providing a more efficient traffic allocation strategy at the micro level, thereby enhancing recommendation performance?
- RQ2 Can the proposed system provide macro-level guidance for the overall traffic allocation mechanism through the analysis of individual weighting schemes?
- RQ3 Is "Effective Completion Rate" a robust choice for the anchor metric in this framework?

3.1 Overall Performance

The results of the online A/B experiments are summarized in Table 1. The key metrics are defined as follows: Visit View (VV) represents the total number of content impressions; Valued VV refers to the count of contents played for more than 3 seconds; Duration indicates the total consumption time; and Valued Score is a core composite metric of the scenario, constructed from diverse user behaviors such as clicks and comments, used to evaluate the overall user experience.

The results demonstrate that the proposed method achieved consistent improvements across all core metrics, confirming the overall system-level benefit of Uniboost. Monitoring data reveals that compared to the Control Group, the Treatment Group achieved a 92.21% reduction in ad weighting scores and a 95.81% reduction in organic video weighting scores, yet maintained stable exposure shares for all weighted content categories. This indicates that Uniboost effectively mitigates the weight explosion issue prevalent in the original system. By significantly reducing the interference of regulation strategies on core business objectives while guaranteeing required exposure levels, Uniboost enhances the overall performance of the recommendation system. These results corroborate Uniboost’s capability to decouple inter-dependencies among various boosting plans, addressing the limitations of traditional boosting. Consequently, the proposed scheme enables more efficient traffic allocation at the micro level, significantly enhancing the overall efficiency of the recommendation pipeline.

3.2 Macro-level Guidance via Offline Analysis

To verify whether the proposed scheme can provide macro-level guidance for traffic mechanism regulation, we statistically analyzed the Cost, VV Lift, and ROI of each business plan. Based on this analysis, we identified Plan-A as having the lowest ROI. We subsequently constructed an ablation study by removing Plan-A from the online configuration. The experimental results are presented in Table 2.

Table 2: Relative Improvement of removing Plan-A.

Metrics	VV	Valued VV	Duration	Valued Score
w/o Plan-A	+0.95%	+2.83%	+3.49%	+4.13%

As observed, removing Plan-A led to a significant improvement in core online metrics. Overall, exposure VV shows a slight increase, while Valued VV, Duration, and Valued Score significantly improve, indicating that the content recommended by the new pipeline is indeed well-received by users. This indicates that Plan-A indeed incurred high costs for marginal gains, thereby negatively impacting the overall performance of the recommendation pipeline. These results corroborate that the analytical reports generated by Uniboost not only track the costs and benefits of individual schemes but also provide actionable guidance for macro-level traffic allocation. Consequently, this enhances the overall efficiency of the recommendation chain and accelerates the iteration speed of strategy optimization.

3.3 Selection of the Anchor Metric

We analyzed the correlation between the blending model prediction scores and common post-metrics, including Effective Completion, Play Duration, Click, Buy, Interaction, and Slide. Figure 2 illustrates the distribution differences of various post-metrics under different prediction scores, where the x-axis represents the model prediction score and the y-axis represents the probability of the post-metric.

We observed that Buy, Interaction, and Slide are unsuitable as anchor metrics due to their extreme sparsity and high volatility across prediction scores. While Interaction and Play Duration are relatively denser, their distributions diverge from the score distribution, particularly in the long-tail region (i.e., high score values), where the prediction error increases significantly. In contrast, the Effective Completion metric remains relatively stable. Its values are uniformly distributed across various intervals of the blending scores, and the calibration error remains low and stable across all intervals. This validates that in video recommendation feeds, Effective Completion serves as a robust anchor data point for aligning the value of blending scores.

4 Conclusion

This work focus on the blending stage within the recommendation system pipeline. In this paper, we propose Uniboost, a unified traffic allocation framework grounded in posterior value alignment. Uniboost tackles the challenge of interpretability in traditional re-ranking by identifying stable posterior targets. It aligns abstract model output scores—which often lack physical semantics—with anchor metrics that possess real-world business value. This alignment significantly enhances the transparency and controllability of the traffic allocation mechanism. Building on this foundation, Uniboost replaces the original complex, coupled, and diverse boosting processes with an independent linear boosting scheme. This innovation provides a more efficient and streamlined solution for overall traffic regulation, decoupling inter-dependent strategies for clearer optimization.

Extensive experiments and analyses demonstrate the dual benefits of our approach. At the micro level, Uniboost allocates traffic more efficiently, improving the operational performance of the overall pipeline. At the macro level, it provides strategic guidance for traffic mechanism design through long-term tracking and analysis of weighting costs for each scheme. Consequently, Uniboost not only optimizes immediate recommendation performance but

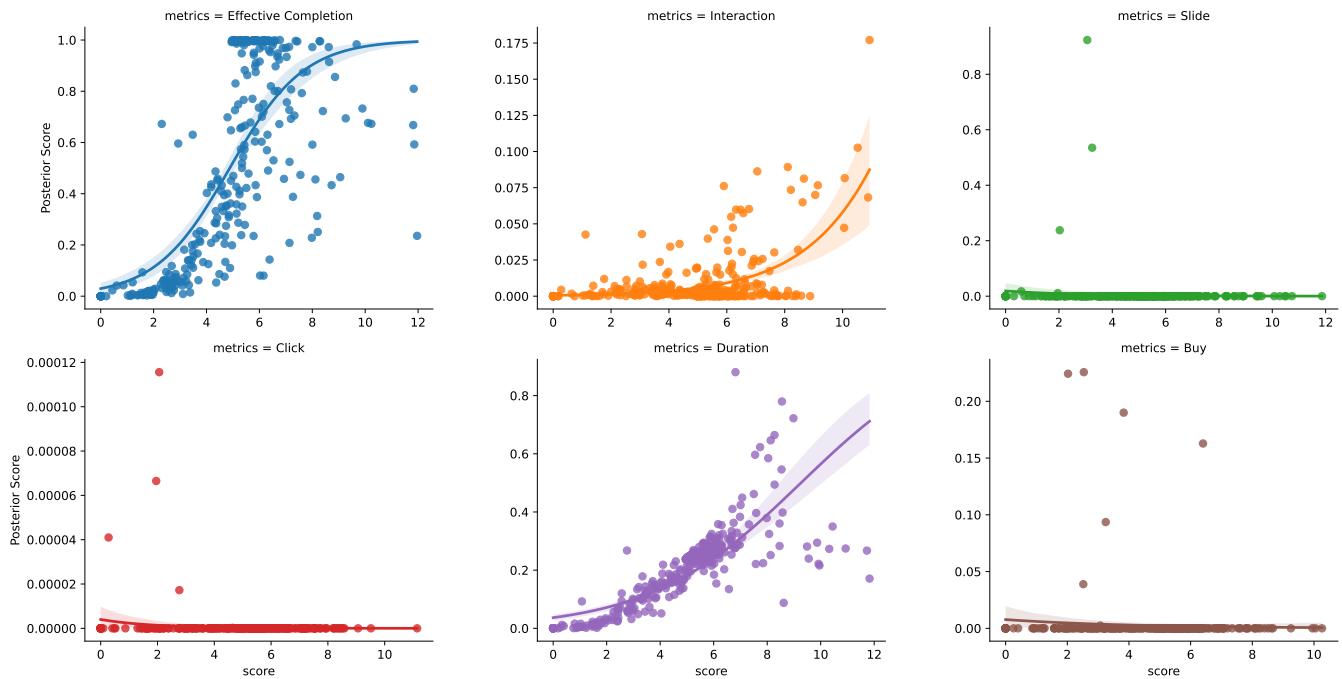


Figure 2: Posterior Business Metrics vs. Original Blending Score y' .

also supports the sustainable iteration and evolution of the entire recommendation system. This solution has been fully deployed in the Taobao Content Feeds recommendation system and is stably operating as the core traffic allocation system.

Acknowledgment

We are deeply grateful to Zhibo Xiao, Dimin Wang and Jialin Zhu for their generous sharing of insights on value alignment in item recommendation systems. Their thought-provoking discussions and practical perspectives, developed through internal collaborative work, have greatly inspired our approach in this study. We sincerely thank all the reviewers and the area chairs for their insightful comments and constructive suggestions, which have greatly improved the quality of this paper. We are also deeply grateful to all collaborators involved in this project, including, but not limited to, the Server Engineering team, Data Science team, Product team, and Operations team, whose support was instrumental in enabling the successful execution of this research.

The authors acknowledge the peoples of the Woi Wurrung and Boon Wurrung language groups of the eastern Kulin Nation on whose unceded lands ACM SIGIR 2026 was hosted. We pay our respects to their Elders past and present, and extend that respect to all Aboriginal and Torres Strait Islander peoples today and their continuing connection to land, sea, sky, and community

References

- [1] Kiam Heong Ang, Gregory Chong, and Yun Li. 2005. PID control system analysis, design, and technology. *IEEE transactions on control systems technology* 13, 4 (2005), 559–576.
- [2] Francesco Busolin, Claudio Lucchese, Franco Maria Nardini, Salvatore Orlando, Raffaele Perego, Salvatore Trani, and Alberto Veneri. 2025. Efficient re-ranking with cross-encoders via early exit. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2534–2544.
- [3] Dagui Chen, Qi Yan, Chunjie Chen, Zhenzhe Zheng, Yangsu Liu, Zhenjia Ma, Chuan Yu, Jian Xu, and Bo Zheng. 2022. Hierarchically constrained adaptive ad exposure in feeds. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 3003–3012.
- [4] Junyang Chen, Ziyi Chen, Mengzhu Wang, Ge Fan, Guo Zhong, Ou Liu, Wenfeng Du, Zhenghua Xu, and Zhiguo Gong. 2023. A neural inference of user social interest for item recommendation. *Data Science and Engineering* 8, 3 (2023), 223–233.
- [5] Junyang Chen, Ge Fan, Zhiguo Gong, Xueliang Li, Victor CM Leung, Mengzhu Wang, and Ming Yang. 2023. A topic-aware graph-based neural network for user interest summarization and item recommendation in social media. In *International Conference on Database Systems for Advanced Applications*. Springer, 537–546.
- [6] Junyang Chen, Zhiguo Gong, Yuanman Li, Huanjian Zhang, Hongyong Yu, Junzhang Zhu, Ge Fan, Xiao-Ming Wu, and Kaishun Wu. 2022. Meta-path based neighbors for behavioral target generalization in sequential recommendation. *IEEE Transactions on Network Science and Engineering* 9, 3 (2022), 1658–1667.
- [7] Ge Fan, Biao Geng, Jianrong Tao, Kai Wang, Changjie Fan, and Wei Zeng. 2022. PPPNE: Personalized proximity preserved network embedding. *Neurocomputing* 472 (2022), 103–112.
- [8] Ge Fan, Chaoyun Zhang, Junyang Chen, Baopu Li, Zenglin Xu, Yingjie Li, Luyu Peng, and Zhiguo Gong. 2022. Field-aware variational autoencoders for billion-scale user representation learning. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 3413–3425.
- [9] Ge Fan, Chaoyun Zhang, Kai Wang, and Junyang Chen. 2022. MV-HAN: A hybrid attentive networks based multi-view learning model for large-scale contents recommendation. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*. 1–5.
- [10] Tamas Jambor, Jun Wang, and Neal Lathia. 2012. Using control theory for stable and efficient recommender systems. In *Proceedings of the 21st international conference on World Wide Web*. 11–20.
- [11] Niklas Karlsson. 2018. Control of periodic systems in online advertising. In *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 5928–5933.
- [12] Xuejian Li, Ze Wang, Bingqi Zhu, Fei He, Yongkang Wang, and Xingxing Wang. 2024. Deep automated mechanism design for integrating ad auction and allocation in feed. In *Proceedings of the 47th international ACM SIGIR conference on research and development in information retrieval*. 1211–1220.
- [13] Greg Linden, Brent Smith, and Jeremy York. 2003. Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet computing* 7, 1 (2003),

76–80.

- [14] Zhenhui Liu, Chunyuan Yuan, Ming Pang, Zheng Fang, Li Yuan, Xue Jiang, Changping Peng, Zhangang Lin, Zheng Luo, and Jingping Shao. 2025. Multi-objective Aligned Bidword Generation Model for E-commerce Search Advertising. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 4299–4303.
- [15] Ning Ma, Mustafa Ispir, Yuan Li, Yongpeng Yang, Zhe Chen, Derek Zhiyuan Cheng, Lan Nie, and Kishor Barman. 2022. An online multi-task learning framework for google feed ads auction models. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3477–3485.
- [16] Maria Movin and Claudia Hauff. 2025. Zero-Shot Reranking with Large Language Models and Precomputed Ranking Features: Opportunities and Limitations. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2276–2286.
- [17] Ajit Puthenpuhussery, Changsung Kang, Alessandro Magnani, Tian Zhang, Hongwei Shang, Nitin Yadav, Prijith Chandran, Bhavin Madhani, Yuan-Tai Fu, He Wang, et al. 2025. Large Scale Deployment of BERT Based Cross Encoder Model for Re-Ranking in Walmart Search Engine. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 4365–4369.
- [18] Srivaths Ranganathan, Chieh Lo, Bernardo Cunha, Nikhil Khani, Li Wei, Aniruddh Nath, Shawn Andrews, Gergo Varady, Yanwei Song, Jochen Klingenhoefer, et al. 2025. Zero-shot Cross-domain Knowledge Distillation: A Case study on YouTube Music. In *Proceedings of the Nineteenth ACM Conference on Recommender Systems*. 1122–1125.
- [19] Viet Anh Tran, Guillaume Salha-Galvan, Bruno Sguerra, and Romain Hennequin. 2023. Attention mixtures for time-aware sequential recommendation. In *Proceedings of the 46th international ACM SIGIR conference on research and development in information retrieval*. 1821–1826.
- [20] Huimu Wang, Mingming Li, Dadong Miao, Songlin Wang, Guoyu Tang, Lin Liu, Sulong Xu, and Jinghe Hu. 2024. A Preference-oriented Diversity Model Based on Mutual-information in Re-ranking for E-commerce Search. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2895–2899.
- [21] Ze Wang, Guogang Liao, Xiaowen Shi, Xiaoxu Wu, Chuheng Zhang, Bingqi Zhu, Yongkang Wang, Xingxing Wang, and Dong Wang. 2022. Hybrid transfer in deep reinforcement learning for ads allocation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 4560–4564.
- [22] Enqiang Xu, Yiming Qiu, Junyang Bai, Ping Zhang, Dadong Miao, Songlin Wang, Guoyu Tang, Lin Liu, and MingMing Li. 2024. Optimizing e-commerce search: Toward a generalizable and rank-consistent pre-ranking model. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2875–2879.
- [23] Jinyun Yan, Zhiyuan Xu, Birjodh Tiwana, and Shaunak Chatterjee. 2020. Ads allocation in feed via constrained optimization. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 3386–3394.
- [24] Weinan Zhang, Yifei Rong, Jun Wang, Tianchi Zhu, and Xiaofan Wang. 2016. Feedback control of real-time display advertising. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*. 407–416.
- [25] Xu Zhao, Ruibo Ma, Jiaqi Chen, Weiqi Zhao, Ping Yang, and Yao Hu. 2025. Multi-Granularity Distribution Modeling for Video Watch Time Prediction via Exponential-Gaussian Mixture Network. In *Proceedings of the Nineteenth ACM Conference on Recommender Systems*. 309–318.