SHORT-PAPER

COUPA: An Industrial Recommender System for Online to Offline Service Platforms

SICONG XIE, Ant group, Hangzhou, Zhejiang, China

BINBIN HU, Ant group, Hangzhou, Zhejiang, China

FENGZE LI, Ant group, Hangzhou, Zhejiang, China

ZIQI LIU, Ant group, Hangzhou, Zhejiang, China

ZHIQIANG ZHANG, Ant group, Hangzhou, Zhejiang, China

WENLIANG ZHONG, Ant group, Hangzhou, Zhejiang, China

View all

Published: 19 July 2023

Citation in BibTeX format

SIGIR '23: The 46th International ACM SIGIR Conference on Research and Development in Information Retrieval
July 23 - 27, 2023
Taipei, Taiwan

Conference Sponsors:
SIGIR

SIGIR '23: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (July 2023)
https://doi.org/10.1145/3539618.3591828
ISBN: 9781450394086

COUPA: An Industrial Recommender System for Online to Offline Service Platforms

Sicong Xie*
Binbin Hu*
sicong.xsc@antfin.com
bin.hbb@antfin.com
Ant Group
Hangzhou, China

Fengze Li
Ant Group
Hangzhou, China
fengze.lfz@antfin.com

Ziqi Liu
Ant Group
Hangzhou, China
ziqiliu@antfin.com

Zhiqiang Zhang
Ant Group
Hangzhou, China
lingyao.zzq@antfin.com

Wenliang Zhong
Ant Group
Hangzhou, China
yice.zwl@antfin.com

Jun Zhou†
Ant Group
Hangzhou, China
jun.zhoujun@antfin.com

ABSTRACT

Aiming at helping users locally discover retail services (*e.g.,* entertainment and dining) on Online to Offline (O2O) service platforms, we propose **COUPA**, an industrial system targeting for characterizing user preference with inspiring considerations of time and position aware preferences. We carefully implement and deploy COUPA in Alipay with a cooperation of edge, streaming and batch computing, as well as a two-stage online serving mode, to support several popular recommendation scenarios. Extensive experiments reveal the superior performance of COUPA for recommendation.

CCS CONCEPTS

• **Information systems** → **Information systems applications**.

KEYWORDS

O2O service platform, Temporal pattern, Position bias

ACM Reference Format:
Sicong Xie, Binbin Hu, Fengze Li, Ziqi Liu, Zhiqiang Zhang, Wenliang Zhong, and Jun Zhou. 2023. COUPA: An Industrial Recommender System for Online to Offline Service Platforms. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23), July 23–27, 2023, Taipei, Taiwan.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3539618.3591828

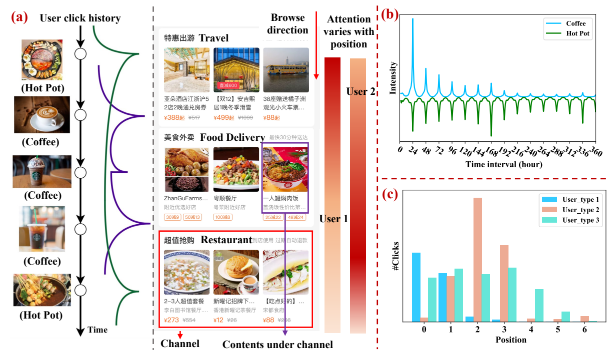

Figure 1: (a) The illustrated example of recommendation scenarios in Alipay. (b): The temporal pattern analysis on "Hot Pot" and "Coffee". (c): The analysis for position bias on channels for "User type 1", "User type 2" and "User type 3". The smaller the number of *x*-axis, the higher the position.

1 INTRODUCTION

Although current user preference modelling based recommendation methods have achieved superior performance to some extent [6], they are still unsuitable to industrial recommendation scenarios for Online to Offline (O2O) service platforms (*e.g.,* Meituan, Grubhub and Uber Eats), where recurrence based temporal patterns are ubiquitous. In the case of Alipay, whose goal is to guide users for locally discovering retail services including entertainment, travelling, delivery and other services, there is a clear and strong need to capture dynamic preference over time for recommendation. Concretely, we illustrate an example in Figure 1 (a). There are numbers of channels (*e.g.,* "Food Delivery" and "Travel" channel) in our scenarios, each of which represents a certain business field. Besides, several items are displayed under each channel, which are called the contents of the channel. Taking the contents of "Food Delivery" channel as an example (as shown in the left part of Figure 1 (a)), users may be interested in daily necessities (*e.g.,* "Coffee") or weekly intents (*e.g.,* "HotPot"), leading to different click demands at different times. On the other hand, as shown in Figure 1 (a), the feeds-like styles cause the fact that users typically scan the screen of mobile phones from top to bottom, revealing distinct attention distributions. That is, users may prefer clicking items with certain positions regardless of the relevance, which brings so-called personalized position biases to hurt the recommendation effectiveness.

*Equal contributions.
†Corresponding author.

SIGIR '23, July 23–27, 2023, Taipei, Taiwan

ACM ISBN 978-1-4503-9408-6/23/07…$15.00
https://doi.org/10.1145/3539618.3591828

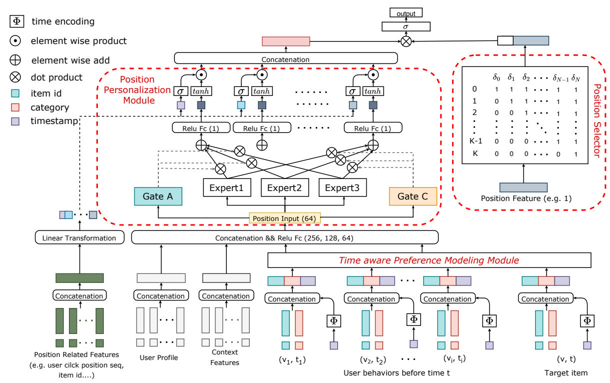

Figure 2: The overall architecture of the proposed COUPA.

In light of the above observations, we aim at designing an industrial recommender system for O2O service platforms, centering on characterizing user preference from time and position bias aware perspectives simultaneously, facing the following challenges: (**C1**) *Time aware preference modelling over continuous time*; (**C2**) *Position aware preference modelling in a personalized way*; and (**C3**) *System design for complicated temporal interaction scenarios.* To this end, we propose **COUPA**, a novel COntinuoUs time and Position bias Aware recommender system targeted for O2O service platforms (*e.g.,* Alipay). Specifically, we propose a novel continuous time aware point process with an attention mechanism to explore the excites of inhabits from previous interaction records in continuous time (**C1**) and a position selector component, cooperated with a position personalization module is elaborately designed to perform personalized position debiasing in our scenarios (**C2**) To satisfy the requirements of online serving, the implementation and deployment are carefully designed to not only collect user behaviors as well as corresponding positions in a real-time manner, but also perform efficient online inference in a two-stage mode (**C3**). We conduct extensive offline and online experiments, demonstrating the superior performance of COUPA.

2 THE PROPOSED COUPA

In this section, we present COUPA, illustrated in Figure 2. Following common strategies [3, 9], given a user u and an item v, we represent the embedding for position related features as $\boldsymbol{p}_{u,v}$, user profile as $\boldsymbol{e}_u$, context features as $\boldsymbol{c}_{u,v}$ and item profile as $\boldsymbol{e}_v$, respectively.

Time aware Preference Modelling. Generally, users' preferences are affected by the items they have already interacted (*e.g.,* click and purchase) with and the effects evolve as time passes. Since modelling such preferences is highly time sensitive and the time encoding should be generalized to any unseen timestamp, we firstly represent timestamps in a continuous, higher dimensional space for preserving the temporal patterns via the functional time encoding function [7, 11, 17], *i.e.,* $\Phi(t) : \mathbb{R}^+ \rightarrow \mathbb{R}^d$. Subsequently, we propose a novel continuous time aware point process module with a well-designed continuous time aware attention mechanism.

i) *Continuous time aware attention mechanism.* We devise it upon the masked self-attention architecture [15] to learn the adaptive weights over continuous time conditioned on the involved user behaviors and target item. Formally, given a triplet $\langle u, v, t\rangle$, we denote the historical interaction records before time t as $S_{u,v} = \{(v_i, t_i)|t_i < t\}$. Following the original self-attention mechanism [15], we obtain the temporal sequence matrix at time t to take account of the relationships of user behaviors $S_{u,v}$ and target item v.

$$\mathbf{Z}_{u,v}(t) = [\boldsymbol{e}_{v_0} \,||\, \Phi(t-t_0);\ ...; \boldsymbol{e}_{v_{N-1}} \,||\, \Phi(t-t_{N-1}); \boldsymbol{e}_v \,||\, \Phi(0)]^T, \quad (1)$$

where "||" denotes concatenation operation, $N = |\mathcal{S}_u|$ is the length of user behaviors and $\boldsymbol{e}_*$ is the original embedding for items. Then, we produce the final representation that summarizes the influence of user behaviors using the scaled dot-product attention:

$$\mathbf{h}_{u,v}(t) = \mathrm{softmax}(\mathbf{Q}_{u,v}(t)\mathbf{K}^T_{u,v}(t)/\sqrt{d})\mathbf{V}_{u,v}(t), \quad (2)$$

where $\mathbf{Q}_{u,v}(t), \mathbf{K}_{u,v}(t)$ and $\mathbf{V}_{u,v}(t)$ respectively denotes the "query", "key" and "value" matrix *w.r.t.* user u and item v in time t, which are linear projections of the temporal sequence matrix $\mathbf{Z}_{u,v}(t)$ [15].

ii) *Continuous time aware point process.* Here, we aim to estimate the likelihood of the target item v based on the historical interaction records $\mathcal{S}_{u,v}$ through a continuous time aware point process, whose major role is to construct conditional intensity function $\lambda_v(t|\mathcal{S}_{u,v})$ [2, 13]. For convenience, we rewrite the above intensity function as $\lambda_v(t|\mathcal{S}_{u,v}) = \psi_v(t - t_i|\mathbf{h}_{u,v})$, where $\psi_v(\cdot)$ is a non-negative function, commonly implemented as the exponential function in previous works. Instead of formulating such a specific functional form, which only models the exponential effects (decrease or increase) of historical behaviors toward the target item, following the idea in [12], we exploit a more complex way to enhance the model capability. That is, we directly model the cumulative intensity function , which can be differentiated for the final intensity function, *i.e.,* $\Psi_{u,v}(\tau|\mathbf{h}_{u,v}) = \int_0^\tau \psi(s|\mathbf{h}_{u,v})ds$, where $\tau = t - t_i$ denotes the interval since the last interaction. Obviously, we adopt an intensity-free formulation to model the user's time aware preference towards target items, which is more suitable to complex scenarios in real-world applications. Due to the ability of modelling non-linear functions, we implement the cumulative intensity function $\Psi_{u,v}(\cdot)$ with a feed-forward neural network as follows:

$$\begin{aligned}&\Psi_{u,v}(\tau|\mathbf{h}_{u,v}) = g(\mathbf{W}_L...g(\mathbf{W}[\boldsymbol{e}_u \,||\, \mathbf{h}_{u,v} \,||\, \tau] + \mathbf{b}_1) + b_L),\\ &s.t. \quad \mathbf{W}_1, ..., \mathbf{W}_L \geq 0, b_1, ..., b_L \geq 0.\end{aligned} \quad (3)$$

Here, $\boldsymbol{e}_u$ is the user embedding derived from his/her original profile features and $g(\cdot)$ is the ReLU activation function. At last, we obtain our final intensity function as follows,

$$\lambda^*_v(t) = \psi_v(\tau|\mathbf{h}_{u,v}) = \partial\Psi_{u,v}(\tau|\mathbf{h}_{u,v})/\partial\tau. \quad (4)$$

Position Bias aware Modelling. Through the above intuitive analysis, user's preference of clicking at certain positions may bring in personalized position bias, while previous works always simply assume that each position is dependent, which ignores the heavy effects of personalized features related to users. To fill this gap, we aim at performing position debiasing in a personalized manner, where a position selector component equipped with a position personalization module is elaborately designed.

i) *Position personalization module.* Since positions differ from each other based on the features of specific users and items, the position uplifts in different positions can be regarded as multiple tasks,

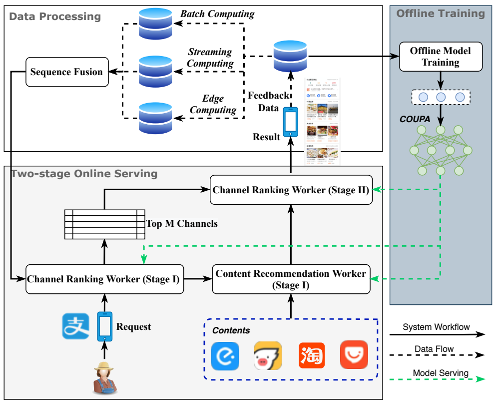

Figure 3: Online deployment of COUPA in Alipay.

whose differences and relations can be naturally captured by the Multi-gate Mixture-of-Experts (MMoE) [10]. Hence, we formulate so-called k-th position uplift for user u and item v as follows:

$$\hat{\delta}_{u,v,k} = ReLU(\boldsymbol{v} \sum_{i=1}^{n} \text{softmax}(\mathbf{W}_k \mathbf{x}_{u,v})_i \cdot f_i(\mathbf{x}_{u,v}) + b), \tag{5}$$

where $\hat{\delta}_{u,v,k}$ refers to the position uplift in position k, $f_i(\cdot)$ refers to the output of i-th experts, n is the number of experts, $\mathbf{W}_k \in \mathbb{R}^{d \times d}$ is a trainable matrix for the k-th position, $\boldsymbol{v}$ and b are also trainable parameters and $\mathbf{x}_{u,v}$ is generated by a 3-layer fully connected neural network with the input of $\mathbf{e}_u$, $\mathbf{c}_{u,v}$ and $\mathbf{h}_{u,v}(t)$.

To further enhance personalization, we incorporate a few position related features (*i.e.,* $\boldsymbol{p}_{u,v}$) and design a Gated Linear Units (GLU) [1] block to control the information passed from features and the position uplift from MMoE block. Specifically, a linear transformation is applied for the position related features $\boldsymbol{p}_{u,v}$ to obtain the gate units $\epsilon_{u,v,k}$ for k-th position. Then, the final position uplift can be calculated with sigmoid function $\sigma(\cdot)$ as:

$$\delta_{u,v,k} = \sigma(\epsilon_{u,v,k}) \cdot tanh(\hat{\delta}_{u,v,k}). \tag{6}$$

ii) *Position selector.* Knowledge transfer (or parameter sharing) has been proven to be potential for facilitating model learning. We take this inspiration to obtain the output $\mu_{u,v,k}$ of k-th position by summarize uplifts from subsequent positions (*i.e.,* $i \geq k$) as $\mu_{u,v,k} = \sum_{i=k}^{K} \delta_{u,v,i}$. Clearly, the uplift for the k-th position $\delta_{u,v,k}$ is involved in the samples with position $i (i \leq k)$, where the learned information can be shared among these samples. To simplify the formulation and speed up the numerical computation, we construct a position matrix, which is denoted as $\mathrm{S}_k = [0_{(0)}, \cdots, 0_{(k-1)}, 1_{(k)}, \cdots, 1_{(K)}]$ for position k. Then, given a user u and an item v with the position k, we estimate the click likelihood as follows:

$$\hat{r}_{u,v,k} = \sigma(\mu_{u,v,k}) = \sigma(\mathrm{S}_k \cdot \Delta_{u,v}^T), \tag{7}$$

where $\Delta_{u,v} = [\delta_{u,v,0}, \cdots, \delta_{u,v,K}]$ is the position uplift vector.

Model Learning. In COUPA, we intend to maximize the following posterior probability of model parameters Θ with observed interaction records $\mathcal{R} = \{u, v, k, t\}$ involving target user u and item v with corresponding position k and click timestamp t: $p(\Theta|\mathcal{R}) \propto p(\Theta) \cdot p(\mathcal{R}|\Theta)$, While the first $p(\Theta)$ can be regarded as a regularizer to avoid overfitting, we mainly focus on the estimation of the second term $p(\mathcal{R}|\Theta)$, which is factorized by minimizing its negative logarithm:

$$\begin{aligned} -\log p(\mathcal{R}|\Theta) &= -\sum_{(u,v,k,t) \in \mathcal{R}} (\log p(v|k, u, \Theta) + \log p(t|u, v, k, \Theta)) \\ &= \sum_{(u,v,k,t) \in \mathcal{R}} C(r_{u,v,k}, \hat{r}_{u,v,k}) \\ &\quad - \log \lambda_v^*(t) + \sum_{v' \sim P_{neg}} \int_0^T \lambda_{v'}^*(t) dt). \end{aligned} \tag{8}$$

Here, $C(\cdot, \cdot)$ is the cross entropy function and P_{neg} is the noise distribution for the generator of negative samples.

Table 1: Offline temporal recommendation performance.

	Food Delivery		Travel		Channel	
#User / #Item	1.38M / 1.93M		1.32M / 17K		8.39M / 17	
#Train / #Test	4.35M / 3.38M		3.84M / 3.18M		31.6M / 1.68M	
Metrics	GAUC	RI	GAUC	RI	GAUC	RI
DeepFM [3]	0.7534	+9.87%	0.7198	+15.92%	0.8837	+0.60%
GRU4Rec [5]	0.7682	+3.80%	0.7373	+7.37%	0.8839	+0.55%
DIN [21]	0.7682	+3.80%	0.7371	+7.46%	0.8846	+0.36%
DIEN [20]	0.7701	+3.07%	0.7421	+5.24%	0.8849	+0.29%
SASRec [8]	0.7688	+3.57%	0.7362	+7.87%	0.8843	+0.44%
Time-LSTM [23]	0.7698	+3.18%	0.7399	+6.21%	0.8842	+0.47%
COUPA$_{\mathcal{T}}$	**0.7784****	-	**0.7548****	-	**0.8860***	-

Table 2: Online performance comparison for position debiasing. We report the improvement ratio over YoutubeRank.

	Food Delivery		Travel		Channel	
Metrics	CTR	IPV	CTR	IPV	CTR	IPV
YoutubeRank [19]	-	-	-	-	-	-
PAL [4]	+1.07%	+1.03%	+0.78%	+0.84%	+1.26%	+1.50%
COUPA$_{\mathcal{P}}$	**+1.44%**	**+1.28%**	**+1.22%**	**+1.17%**	**+2.04%**	**+2.18%**

3 ONLINE DEPLOYMENT OF COUPA

Here, we present the deployment of COUPA in Alipay (Figure 3).

Data Processing. User feedback data are necessary to train the recommender system, where the trade-off between data storage and timely feedback need to be balanced. Hence, our principle is to store 90-day click data in huge volume via cheap storage that only guarantees latency in days, and timely feedback data in tiny volume via edges that guarantees latency in seconds. As such, we design the following batch, streaming and edge computing in a cooperative manner, followed by a sequence fusion procedure.

Batch computing. Based on the offline and low-cost MaxCompute platform, we perform batch computing for generating user click

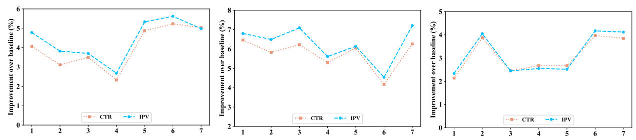

(a) Food Delivery (b) Travel (c) Channel

Figure 4: Online performance on real scenarios in Alipay.

sequences in the last 90 days, involving hundreds of billions of interaction records, with a latency of days.

Streaming computing. Streaming computing is mainly based on the frigate system (similar to Kafka and Storm system) in the Ant Group, which directs at generating user click sequences in the last 2 days. Due to the effects of log back and real-time reporting, this part of the data has a latency of tens of seconds.

Edge computing. Due to the urgent real-time demand of COUPA, we employ edge computing for collecting click sequences on the client in completely real time with a latency of only few seconds. Specifically, user click data will be captured by the client at once after reported, and recorded into the local database. As soon as a user proposes a request, the client timely processes his/her click data in the last 3 hours and sends them to online servers.

Through the *sequence fusion*, a complete sequence in the last 90 days with chronological order is passed to COUPA as input.

Two-stage Online Serving. This module is designed to relax delays of the online system for COUPA with two main stages. *Stage I*: Channel Ranking Worker performs the coarse-grained ranking on channels with dozens of candidates based on COUPA, ignoring features of contents under channels. Then only top M (*i.e.,* $M = 5$) channels are selected to request their corresponding contents (*i.e.,* "Food Deliver" and "Travel" channels) recommender systems with tens of millions of candidates in Content Recommendation Worker, whose ranking model is COUPA. *Stage II*: With contents under channels determined in *Stage I* as well as collected content features, Channel Ranking Worker performs the fine-grained ranking considering content features on top M channels to decide their displayed orders. In this way, system latency is greatly relaxed and ensures the stability of COUPA for online inference.

4 EXPERIMENTS

Performance Comparison over Temporal Recommendation. We report the comparison results *w.r.t.* **GAUC** [16, 22] and relative improvement (**RI** [14, 18]) on three offline datasets in Table 1 (Since the performance of position debiasing cannot be accurately evaluated offline, we only show the performance of the variant only modeling user's temporal preference over continuous time *i.e.,*COUPA$_{\mathcal{T}}$). COUPA$_{\mathcal{T}}$ is consistently better than all the baselines with statistical significance at the level of 0.01 (0.05) (marked by "** (or *)" in Table 1), demonstrating its effectiveness to learn time aware preference for recommendation.

Performance Comparison over Position Debiasing. We examine the performance of position debiasing *w.r.t.* **IPV** and **CTR** on three online scenarios (*i.e.,* Food Delivery, Travel and Channels). For a fair comparison, we only report the performance of COUPA$_{\mathcal{P}}$, aiming at position debiasing task. We perform the evaluation within 7 days and report average results in Table 2. From the results, we observe that COUPA$_{\mathcal{P}}$ achieves the best performance with statistical significance in all scenarios in terms of both IPV and CTR metrics, clearly demonstrating the superiorities of COUPA$_{\mathcal{P}}$ for performing position debiasing in a personalized manner.

Online Performance. We deploy COUPA into three scenarios (*i.e.,* Food delivery, Travel and Channels) in Alipay, comparing it with the existing deployed baseline in our real system from "2022/01/13" to "2022/01/20". From the results in Figure 4, We observe that, compared to the best baseline used in our real system, COUPA consistently and significantly yields performance improvement by a large margin in three scenarios across all online metrics.

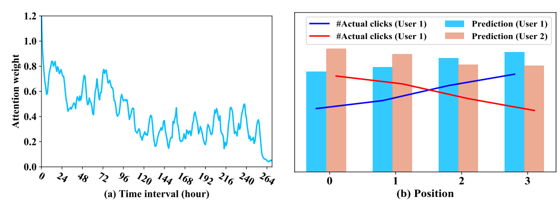

Figure 5: Case studies for COUPA.

Case Study. (**Temporal Perspective**) We select a certain user from the Food Delivery dataset and plot his/her attention weights *w.r.t.* time interval of historical behaviors in Figure 5 (a). With the increase of time interval, the attention weight decreases with a periodical trend. (**Position Bias**) We select two typical user clusters and plot the average prediction score for each position k(*i.e.,* $\sum_v \mu_{u,v,k}$) as well as their actual click trends in Figure 5 (b). As the position changed, we find the distinct trends of average prediction scores for the two selected users (*i.e.,* increase for "User 1" and decrease for "User 2"), consistent with their trends of actual clicks *w.r.t.* each position. Both analyses verify the strong capability of COUPA for summarizing influence from historical interactions and capturing the personalized bias.

5 CONCLUSION

In this paper, we proposed COUPA for O2O service platforms and carefully deployed it in Alipay, which comprehensively takes user preferences towards temporal patterns and position biases into consideration. Extensive offline and online experiments demonstrated the superiority of COUPA.

6 PRESENTER'S BIO

Ant Group aims to create the infrastructure and platform to support the digital transformation of the service industry, enabling all consumers and small businesses to have equal access to financial and other services that are inclusive, green and sustainable.
Binbin Hu is an algorithm expert in the Ant Group, devoting to graph machine learning and recommendation. Binbin Hu receives the master degree from BUPT and has published more than 30 papers in top venues, *e.g.,* KDD, SIGIR and WWW.

REFERENCES

[1] Yann N Dauphin, Angela Fan, Michael Auli, and David Grangier. 2017. Language modeling with gated convolutional networks. In *ICML*. 933–941.
[2] Nan Du, Hanjun Dai, Rakshit Trivedi, Utkarsh Upadhyay, Manuel Gomez-Rodriguez, and Le Song. 2016. Recurrent marked temporal point processes: Embedding event history to vector. In *SIGKDD*. 1555–1564.
[3] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. 2017. DeepFM: a factorization-machine based neural network for CTR prediction. In *IJCAI*. 1725–1731.
[4] Huifeng Guo, Jinkai Yu, Qing Liu, Ruiming Tang, and Yuzhou Zhang. 2019. PAL: a position-bias aware learning framework for CTR prediction in live recommender systems. In *RecSys*. 452–456.
[5] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based recommendations with recurrent neural networks. In *ICLR*.
[6] Yupeng Hou, Binbin Hu, Zhiqiang Zhang, and Wayne Xin Zhao. 2022. Core: simple and effective session-based recommendation within consistent representation space. In *SIGIR*. 1796–1801.
[7] Binbin Hu, Zhengwei Wu, Jun Zhou, Ziqi Liu, Zhigang Huangfu, Zhiqiang Zhang, and Chaochao Chen. 2022. MERIT: Learning Multi-level Representations on Temporal Graphs. In *IJCAI*.
[8] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *ICDM*. 197–206.
[9] Jianxun Lian, Xiaohuan Zhou, Fuzheng Zhang, Zhongxia Chen, Xing Xie, and Guangzhong Sun. 2018. xdeepfm: Combining explicit and implicit feature interactions for recommender systems. In *SIGKDD*. 1754–1763.
[10] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Jilin Chen, Lichan Hong, and Ed H Chi. 2018. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *SIGKDD*. 1930–1939.
[11] Ha Quang Minh, Partha Niyogi, and Yuan Yao. 2006. Mercer's theorem, feature maps, and smoothing. In *COLT*. 154–168.
[12] Takahiro Omi, Naonori Ueda, and Kazuyuki Aihara. 2019. Fully neural network based model for general temporal point processes. In *NIPS*. 2120–2129.
[13] Oleksandr Shchur, Ali Caner Türkmen, Tim Januschowski, and Stephan Günnemann. 2021. Neural Temporal Point Processes: A Review. In *IJCAI*. 4585–4593.
[14] Qijie Shen, Wanjie Tao, Jing Zhang, Hong Wen, Zulong Chen, and Quan Lu. 2021. SAR-Net: A Scenario-Aware Ranking Network for Personalized Fair Recommendation in Hundreds of Travel Scenarios. In *CIKM*. 4094–4103.
[15] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*. 5998–6008.
[16] Hong Wen, Jing Zhang, Yuan Wang, Fuyu Lv, Wentian Bao, Quan Lin, and Keping Yang. 2020. Entire space multi-task modeling via post-click behavior decomposition for conversion rate prediction. In *SIGIR*. 2377–2386.
[17] Da Xu, Chuanwei Ruan, Sushant Kumar, Evren Korpeoglu, and Kannan Achan. 2019. Self-attention with functional time representation learning. In *NIPS*. 15915—-15925.
[18] Ling Yan, Wu-jun Li, Gui-Rong Xue, and Dingyi Han. 2014. Coupled group lasso for web-scale ctr prediction in display advertising. In *ICML*. 802–810.
[19] Zhe Zhao, Lichan Hong, Li Wei, Jilin Chen, Aniruddh Nath, Shawn Andrews, Aditee Kumthekar, Maheswaran Sathiamoorthy, Xinyang Yi, and Ed Chi. 2019. Recommending what video to watch next: a multitask ranking system. In *RecSys*. 43–51.
[20] Guorui Zhou, Na Mou, Ying Fan, Qi Pi, Weijie Bian, Chang Zhou, Xiaoqiang Zhu, and Kun Gai. 2019. Deep interest evolution network for click-through rate prediction. In *AAAI*. 5941–5948.
[21] 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 *SIGKDD*. 1059–1068.
[22] Han Zhu, Junqi Jin, Chang Tan, Fei Pan, Yifan Zeng, Han Li, and Kun Gai. 2017. Optimized cost per click in taobao display advertising. In *SIGKDD*. 2191–2200.
[23] Yu Zhu, Hao Li, Yikang Liao, Beidou Wang, Ziyu Guan, Haifeng Liu, and Deng Cai. 2017. What to Do Next: Modeling User Behaviors by Time-LSTM.. In *IJCAI*. 3602–3608.