Understanding and Modeling Passive-Negative Feedback for Short-video Sequential Recommendation

Yunzhu Pan*
University of Electronic Science and Technology of China
Chengdu, China

Chen Gao†
Beijing National Research Center for Information Science and Technology, Department of Electronic Engineering, Tsinghua University
Beijing, China

Jianxin Chang
Beijing Kuaishou Technology Co., Ltd.
Beijing, China

Yanan Niu
Beijing Kuaishou Technology Co., Ltd.
Beijing, China

Yang Song
Beijing Kuaishou Technology Co., Ltd.
Beijing, China

Kun Gai
Unaffiliated
Beijing, China

Depeng Jin
Beijing National Research Center for Information Science and Technology, Department of Electronic Engineering, Tsinghua University
Beijing, China

Yong Li
Beijing National Research Center for Information Science and Technology, Department of Electronic Engineering, Tsinghua University
Beijing, China

ABSTRACT

Sequential recommendation is one of the most important tasks in recommender systems, which aims to recommend the next interacted item with historical behaviors as input. Traditional sequential recommendation always mainly considers the collected positive feedback such as click, purchase, etc. However, in short-video platforms such as TikTok, video viewing behavior may not always represent positive feedback. Specifically, the videos are played automatically, and users passively receive the recommended videos. In this new scenario, users passively express negative feedback by skipping over videos they do not like, which provides valuable information about their preferences. Different from the negative feedback studied in traditional recommender systems, this passive-negative feedback can reflect users' interests and serve as an important supervision signal in extracting users' preferences. Therefore, it is essential to carefully design and utilize it in this novel recommendation scenario. In this work, we first conduct analyses based on a large-scale real-world short-video behavior dataset and illustrate the significance of leveraging passive feedback. We then propose a novel method that deploys the sub-interest encoder, which incorporates positive feedback and passive-negative feedback as supervision signals to learn the user's current active sub-interest. Moreover, we introduce an adaptive fusion layer to integrate various sub-interests effectively. To enhance the robustness of our model, we then introduce a multi-task learning module to simultaneously optimize two kinds of feedback – passive-negative feedback and traditional randomly-sampled negative feedback. The experiments on two large-scale datasets verify that the proposed method can significantly outperform state-of-the-art approaches. The code is released at *https://github.com/PanYZhu/SINE* to benefit the community.

ACM Reference Format:
Yunzhu Pan, Chen Gao, Jianxin Chang, Yanan Niu, Yang Song, Kun Gai, Depeng Jin, Yong Li. 2023. Understanding and Modeling Passive-Negative Feedback for Short-video Sequential Recommendation. In *Seventeenth ACM Conference on Recommender Systems (RecSys '23), September 18–22, 2023, Singapore, Singapore.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3604915.3608814

1 INTRODUCTION

The sequential recommendation is one of the most fundamental tasks, which enhances the basic collaborative filtering with the sequential behaviors [39]. In real-world information systems, sequential recommenders serve as the core of the recommendation engine. The traditional methods of sequential recommendation follow the same paradigm of learning user interests from the behavioral sequence, with positive feedback only, such as [18, 21, 44]. However, the recent success of short-form videos such as TikTok[1] has re-defined the interaction manner of how the user access online content. In these platforms, users passively receive recommended items in a new single-column layout, and the videos are automatically played unless users actively skip over them. As a result, a

[1] https://www.tiktok.com/en/

*Work done when interning at Tsinghua University.
†Chen Gao is the corresponding author (chgao96@gmail.com).

RecSys '23, September 18–22, 2023, Singapore, Singapore

ACM ISBN 979-8-4007-0241-9/23/09…$15.00
https://doi.org/10.1145/3604915.3608814

new type of passive-negative feedback has emerged, where users passively skip videos to find content that interests them. However, this passive-negative feedback cannot be handled in the traditional manner of negative feedback in the previous works [14, 32, 37, 41]. Specifically, watching a video on these platforms may not necessarily indicate the user likes the video, as videos are played automatically, but users can choose to skip over it, which suggests that the user is not interested in this particular video. Therefore, there is a need for novel approaches to handle this unique feedback scenario.

We first conduct empirical analysis from a statistical perspective and the recommendation performance to fully understand this new negative feedback. Our results are based on the collected behavioral dataset from one of the largest short-video platforms, along with the state-of-the-art sequential recommendation model. The empirical analysis[2] first demonstrates that the negative feedback in real-world applications is always passive, *i.e.,* the users choose to skip over the recommended content, and often do not explicitly report it to the platform. Second, the analysis results show that roughly using these kinds of passive-negative feedback in the recommendation model will even lead to a significant performance drop, illustrating a challenge of feedback learning. Last, based on the hierarchical categorical information, we find that the items receiving passive-negative feedback often share similar categories with those that receive positive feedback. That is, the items are not truly hated/disliked by users. On the contrary, users are not satisfied with certain features of the items. For example, users may have consumed similar content before and no longer find it useful or interesting. Therefore, without a unique design, it is difficult for a model to learn the difference between passive-negative and positive items.

The above results show the significance of negative feedback and illustrate the challenge and its explanation of taking this passive-negative feedback into consideration of the sequential recommenders. To address it, we propose a novel method named SINE (short for **S**ub-**IN**terest learning with Negativ**E** feedback). Specifically, we first design a mix-feedback sequential encoder that takes both positive feedback and passive-negative feedback, and extracts sub-interests for the given context with a sub-interest-based self-attention layer. We then propose an adaptive fusion layer that selects the activated sub-interests and deactivates the remaining ones in the user behaviors. Last, to well exploit both the collected passive-negative feedback and the unobserved ones, we adopt a multi-task learning paradigm for the optimization of model parameters. The main contributions of this paper can be summarized as follows,

- In this paper, we approach the new problem of understanding and modeling of user's passive-negative feedback in the sequential recommendation, particularly in the new paradigm of single-column short-video platforms. We collect users' skipping behaviors from short-video platforms as passive-negative feedback , which provide a more accurate reflection of users' preferences. By skipping a video, users actively indicate their disinterest or dissatisfaction with the content, which provides a clear signal to the recommendation system about their preferences. This approach contrasts with previous studies that have relied on exposure-based negative feedback, such as "expose-but-not-click" behavior, which may not always accurately reflect a user's preferences due to the exposure bias problem. We conduct an empirical analysis on a large-scale real-world dataset, which reveals the importance, challenge, and explanations for negative feedback in today's short-video platforms.
- We proposed a method with sub-interest learning, which can well handle the mix-feedback sequence and extract the sub-interests that belong to different subspaces. The joint optimization elegantly takes positive feedback, negative feedback, and unobserved feedback into consideration at the same time.
- We conduct experiments on two large-scale datasets from two mainstream short-video platforms. Extensive results show that our method can steadily and significantly outperform the state-of-the-art recommendation methods. Our further experiments well demonstrate the rationality of each component of our SINE method. The results of recommendation performance also correspond well to our earlier data analysis.

The remainder of this paper is as follows. We first conduct an empirical analysis on the real-world dataset and provide the motivations of the research problem in Section 2. We then formally define the research problem in Section 3 and present our solution in Section 4. We conduct experiments on Section 5 and review the related works in Section 6. Last, we conclude this paper and discuss the important future works in Section 7.

2 DATA ANALYSIS AND MOTIVATION

In this section, we aim to understand the significance of negative feedback and the critical challenge of modeling it in the real-world short-video recommendation. Specifically, we conduct analysis on the behavioral dataset collected from one of the largest micro-video platforms. We first present the ratio of negative behavior compared with other kinds of behaviors, through which we can find the active-negative behavior is very sparse, and the passive-negative behaviors are far easier to collect. We then conduct experiments on the widely-used sequential recommendation model to illustrate the challenge of leveraging the passive-negative feedback data. Finally, we further study the passive-negative feedback to answer how this behavior occurs, serving as guidance for designing more powerful recommendation models.

2.1 Data characteristics of negative feedback

We first obtain the characteristics of negative feedback, which consists of two major forms in a typical recommendation in today's information system. First, the user can actively convey their opinions by selecting options includes *"not interested"*, *"reduce similar recommendations"*, or even *"hate"*. Second, when the user passively receives sequentially recommended videos, the user can skip over the video fast if he/she feels not interested. For these two kinds of negative feedback, we present the interaction number based on a real-world dataset collected from one of the famous short-video platform[3], along with effective-view behavior, *i.e.*, the user has spent

[2]The experiments details are presented in Section 2.

[3]This dataset will be used for evaluation dataset, and its detail will be introduced in Section 5.

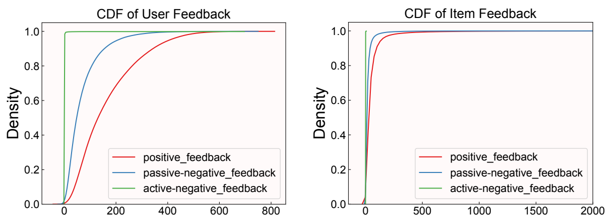

(a) User feedback distribution.

(b) Item feedback distribution.

Figure 1: Distribution of three kinds of feedback on the Kuaishou dataset.

adequate time in watching the video[4] in Fig. 1. From the figure, we can observe that the active-negative feedback is far sparser than the passive-negative feedback, and the density of passive-negative feedback is close to that of positive behavior. The extremely sparse active-negative feedback can be easily understood that users are not always willing to perform additional operations. It means that when designing recommendation models, active-negative feedback can hardly enhance the preference modeling as it is too sparse, while passive-negative feedback may play an important role.

2.2 Study of leveraging passive-negative feedback

However, despite the promising density, the passive-negative feedback is not easy to handle in recommendation system. In this section, we want to design an experiment to demonstrate the challenge of leveraging such kind of feedback. From the perspective of representation learning, the recommendation models need to project users and items into the latent vector in low-dimensional space, and the matching scores between vectors can be used for generating recommendation results. The optimization of the representation learning can treat the passive-negative feedback as negative sample, no matter with point-wise optimization [17, 42] or pair-wise optmization [29]. Therefore, we conduct experiments with the popular SASRec [21] model as the backbone, which adopts self-attention layers to encode users' behavioral sequences. Specifically, in the first case, we use the observed view behavior as positive feedback and randomly sampled items (without observed behavior) as negative feedback. That is, the first case refers to the normal learning procedure. In the second space, we also consider the truly-observed passive-negative behavior as the negative sample, and we carefully control its ratio compared with randomly-sampled samples by a hyper-parameter. We name it SASRec-N (N denotes Negative). We then obtain the performance of two cases, both of which are under the careful and extensive hyper-parameter search. The results are shown in Fig. 2, in which three widely-considered metrics AUC, GAUC and NDCG are used. We can observe from the results that the intuitive design of leveraging passive-negative feedback in the negative sampling and optimization procedure can even harm the recommendation performance. The results demonstrate that it is challenging to well exploit passive-negative feedback even if it seems to make sense.

[4]A common criterion in the industry is the half length of the whole video.

2.3 Analysis of how passive-negative feedback occurs

Given the analysis above, which shows the importance along with the challenges of exploiting the passive-negative feedback, we further explore the reasons behind the occurrence of passive-negative feedback to better address this challenge and inspire model design. We try to answer the question of why the user decides to skip over the video, which is very hard as we can hardly know what the users are thinking about. To address it, we choose to answer a similar but easier question: what is the difference between positive feedback and negative feedback if they occur together? Therefore, we introduce the auxiliary category information of videos and analyze the difference between positive-negative feedback on the category aspect. Specifically, we analyze the categorical difference between positive feedback and two kinds of negative feedback - true negative feedback (*i.e.*, passive-negative feedback) and randomly-sampled negative feedback, based on three levels of categories from coarse-grained to fine-grained. The categorical difference from real-world interaction data is shown in Fig. 2(c). To better compare the differences in categories at different granularities, we divide the results into four cases, explained as follows. Case 1 denotes the negative feedback video having a different level-1 category with the positive feedback video. Case 2 means the positive feedback video and negative feedback video have the same level-1 category but differ in the category level-2. Case 3 indicates that the positive feedback video and negative feedback video have the same level-1 and level-2 category but different level-3 category. Case 4 shows the number of the positive feedback video and negative feedback having the same category at all three levels. We can observe from the results that the categories of a positive-negative pair (occurs very close) are always very similar or even the same. In addition, the number of positive feedback and true negative feedback having same category is larger than that of positive feedback and randomly-sampled negative feedback. This difference is more evident in finer-grained categories. Therefore, we can infer that the recommender system has well estimated the user's interests as the recommended item that received negative feedback is already very similar to the positive item. However, there may be some sub-aspects of the item that do not align with the user's preferences, resulting in the final negative behavior. In other words, the passive-negative item has met a part of the user's preferences but fails to match all of them, causing the user to skip over the item.

In short, we conduct early analysis on a real-world dataset, which first shows the importance of passive-negative feedback in preference learning, then illustrates the challenges of exploitation via experimental results, and finally partly provides the reasons for passive-negative feedback.

3 PROBLEM FORMULATION

In this work, we approach the new problem of modeling the passive-negative feedback and positive feedback in short-video sequential recommendation. Based on the analysis of the real-world data in Section 2, the passive-negative feedback is more promising than the active-negative, considering the data sparsity. Let $\mathcal{U}$ and $\mathcal{I}$ denote the sets of users and items, respectively, of which the

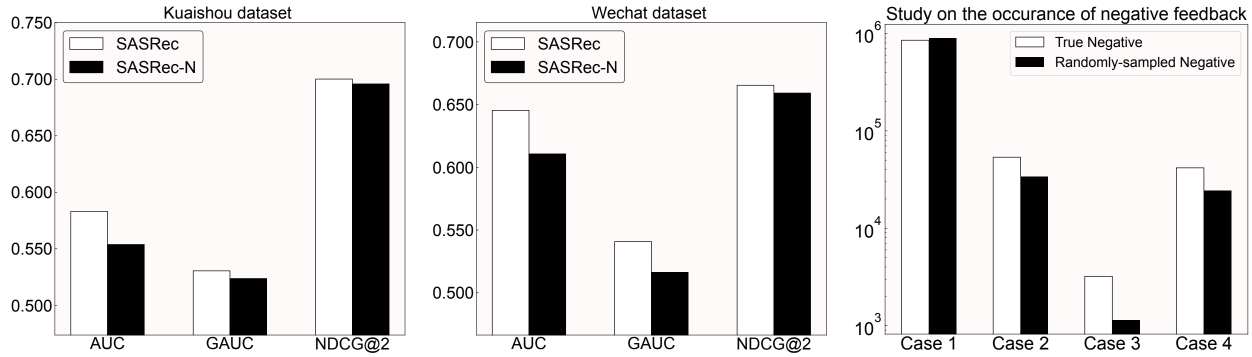

(a) Results of introducing passive-negative feedback on Kuaishou.

(b) Results of introducing passive-negative feedback on Wechat.

(c) Occurrence of passive-negative feedback (case 1: different level-1 categories; case 2: same level-1 but different level-2; case 3: same level-2 but different level-3; case 4: same level-3.)

Figure 2: Data analysis of the passive-negative feedback on real-world datasets via comparing recommendation models (a)(b) and the category-based statistics (c).

Table 1: Frequently used notations in this paper.

Notations	Descriptions
$\lvert \cdot \rvert$	The cardinality of a set
$\langle \cdot \rangle$	Inner product
$\mathcal{U}, \mathcal{I}$	The set of users and items
M, N	The number of users and items
$\mathbf{S}^u, \mathbf{R}^u$	The item sequence and the feedback label of user u
$\mathbf{S}^u_+, \mathbf{S}^u_-$	The item sequence with positive/passive-negative feedback of user u
$E \in \mathbb{R}^{N \times D}$	Item embedding matrix with dimension size of D
$\boldsymbol{e}^u_+, \boldsymbol{e}^u_-$	The item embedding of positive/passive-negative feedback of user u
Z	Sub-interest prototypes
K	Number of sub-interest prototypes
$\hat{z}^{s^u_{+,1}}, \cdots, \hat{z}^{s^u_{+,\lvert S^u_+ \rvert}}$	The corresponded sub-interest of the positive feedback sequence of user u
$\mathbf{Q}, \mathbf{K}, \mathbf{V}, \mathbf{P}$	Query/ Key/Value/position embedding in self-attention
α	Self-attention weight
β	Sub-interests weight
$\widetilde{o}^{u,k}_t$	The k-th sub-interest of user u at timestamp t
r, γ	weight and normalized weight of different sub-interests
L_1, L_2	Loss function on two kind of pairs O_1 and O_2
L_{dis}	The discrepancy loss on sub-interest prototypes

sizes are M and N. Given a user $u \in \mathcal{U}$, its sequence can be denoted as $\mathbf{S}^u = \{s^u_1, s^u_2, \cdots, s^u_{\lvert \mathbf{S}^u \rvert}\}$. Since in the behavioral sequence there are two kinds of feedback, we introduce another sequence $\mathcal{R}^u = \{r^u_1, r^u_2, \cdots, r^u_{\lvert \mathbf{R}^u \rvert}\}$, of which r_u could be 1 for positive feedback and 0 for passive-negative feedback. The frequently-used symbols are listed in Table 1. Therefore, the new-form sequential recommendation can be formally defined as follows.

Input: The sequentially-interacted items of the user u, $\mathbf{S}^u$, along with the feedback type, $\mathcal{R}^u$.

Output: A recommendation model that can estimate the probability, $p^u_{\lvert \mathbf{S}^u \rvert+1,i}$, that the given user u will interact with the target item i at the next time.

4 METHODOLOGY

Inspired and motivated by the analysis of the negative feedback, we propose our method named SINE based on the idea that the passive-negative feedback is caused by the mismatch of specific sub-interests. Specifically, as mentioned above, the videos that received negative feedback shares similar categories with the positive videos, which indicates the recommender system has only partially met the user's preferences but fails to match all of them, causing users to skip over. Our approach, as illustrated in Fig. 3, consists of the following components:

- **Sub-interest based sequential encoder.** Different from the pure-positive feedback in traditional sequential recommenders,

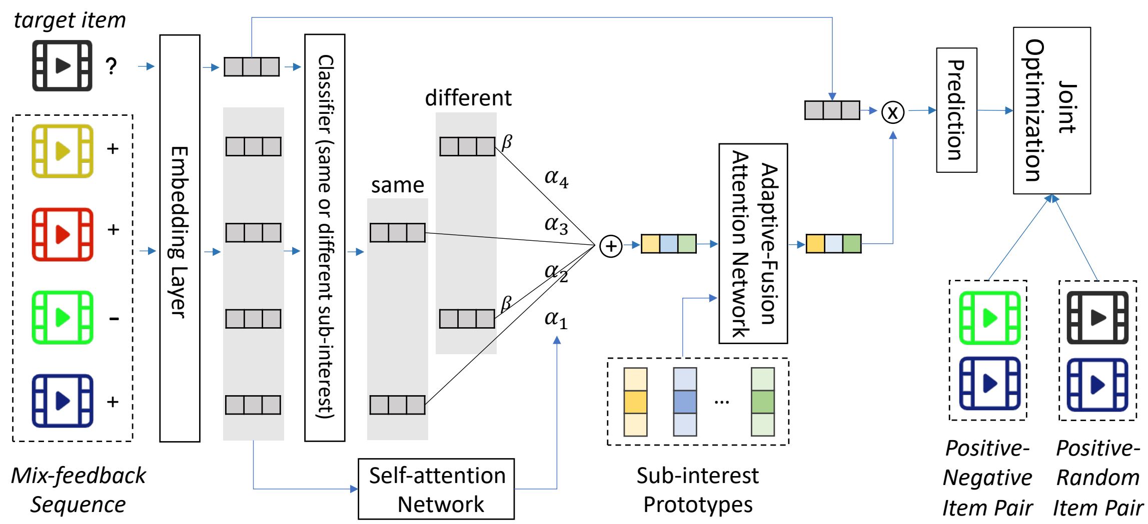

Figure 3: Illustration of our proposed SINE method.

we propose to encode the complex sequence with mixed feedback with a self-attention layer. We propose to project user preferences into multiple sub-spaces, each of which corresponds to a specific aspect that may lead to user behaviors.

- **Adaptive-fusion prediction layer.** With the learned sub-interests, we deploy a prediction layer that can distinguish not only the normal negative feedback but also the passive-negative feedback from positive feedback via learned weights-based interest fusion.
- **Joint optimization.** To well leverage these two kinds of feedback, we propose to jointly optimize the model parameters under a multi-task learning paradigm, in which fitting each kind of data can be treated as a task, and the hyper-parameters can well control the importance of two tasks.

4.1 Subinterest-based Sequential Encoder

4.1.1 Embeddings of users, items, and sub-interest prototypes. First, we build an embedding matrix $\boldsymbol{E} \in \mathbb{R}^{N \times D}$ (D is the dimension size) that assigns low-dimensional vectors to encode each item, and then retrieved embedding of the user sequence S^u at the timestamp t can be represented as follows,

$$\boldsymbol{E}^u = (\boldsymbol{e}_1^u, \boldsymbol{e}_2^u, \cdots, \boldsymbol{e}_t^u), \tag{1}$$

Since the user sequence S^u contains two kinds of feedback, and, for convenience, we use two symbols S_+^u and S_-^u to denote the positive feedback sequence and passive-negative feedback sequence, respectively, defined as follows,

$$\begin{aligned} S_+^u &\in S^u, \text{where } R^u = 1, \\ S_-^u &\in S^u, \text{where } R^u = 0, \end{aligned} \tag{2}$$

and item embedding of positive feedback and passive-negative feedback are denoted as $\boldsymbol{e}_+^u$ and $\boldsymbol{e}_-^u$, respectively.

Generally, a user interacts with a short video due to multiple aspects, such as the style, author, music, etc. Thus, the users can refuse the recommendation by skipping over the video only because one of these aspects does not satisfy the user. To capture the sub-interests towards different aspects, we build the prototype vectors, each of which represents an aspect of users' interests, shown as follows,

$$Z = (z_1, z_2, \cdots, z_K), \tag{3}$$

where K denotes the number of prototypes, which is a controllable hyper-parameter.

To map each feedback in a positive feedback sequence to the corresponding sub-interest (*i.e.,* find the dominate sub-interest that leads to the behavior), we propose to utilize the recent passive-negative feedback to calculate the gap between the matching scores of positive and passive-negative feedback with each sub-interest, calculated as follows,

$$\hat{z}^i = \arg\max_k \boldsymbol{e}_+^u \cdot z_k - \boldsymbol{e}_-^u \cdot z_k, \tag{4}$$

where $\hat{z}^i$ is selected from $1, 2, \cdots, K$ sub-interests. Then the corresponding sub-interests of positive feedback in a user sequence are as follows,

$$\hat{z}^{s^u_{+,1}}, \hat{z}^{s^u_{+,2}}, \cdots, \hat{z}^{s^u_{+,|S^u_+|}}, \tag{5}$$

For positive feedback without recent negative feedback, we omit the second term and directly calculate the similarity between item and prototype embeddings.

4.1.2 Sub-interest enhanced sequence encoder. With the proposed component above, we obtain the specific aspect of user preferences that lead to user feedback, *i.e.,* sub-interest, which can be further exploited to encode the behavioral sequence. Specifically, we propose a multi-head self-attention-based encoder that can well distinguish the complex relations between items, especially for the impact of two kinds of feedback and the sub-interests. The encoder can be formulated as follows,

$$\begin{aligned} \mathbf{X} &= \mathbf{E}^u + \mathbf{P}, \\ \mathbf{Q} = \mathbf{X}\mathbf{W}^{\mathbf{q}}; \mathbf{K} &= \mathbf{X}\mathbf{W}^{\mathbf{k}}; \mathbf{V} = \mathbf{X}\mathbf{W}^{\mathbf{v}}, \end{aligned} \tag{6}$$

where $\mathbf{p}$ denotes the position embedding, $\mathbf{W}^{\mathbf{q}}, \mathbf{W}^{\mathbf{k}}, \mathbf{W}^{\mathbf{v}} \in \mathbb{R}^{D \times D}$ are three projection matrices and $\mathbf{Q}, \mathbf{K}, \mathbf{V} \in \mathbb{R}^{L \times D}$ (L is the length of sequence).

For the sub-interest space, there exist implicit relations between different sub-interests, and thus we design a controllable weight β that models the correlation of sub-interests. First, the self-attention weights for each item is formulated as follows:

$$\alpha = \textbf{Softmax}(\frac{\mathbf{QK}^T}{\sqrt{d}}), \tag{7}$$

Then, we introduce β to learn the correlation of sub-interests, defined as follows:

$$\beta = \begin{cases} \beta_1, \hat{z}^i = \hat{z}^{s^u_{+,|S^u_+|}}, \\ \beta_2, \hat{z}^i \neq \hat{z}^{s^u_{+,|S^u_+|}}, \end{cases} \tag{8}$$

where $\beta_1 + \beta_2 = 1$. Last, we obtain the final attention weights for each item by multiplying β to α, shown as below:

$$\hat{\alpha} = \beta \cdot \alpha, \tag{9}$$

With the obtained $\hat{\alpha}$, we not only consider the items' role in the whole interaction sequence, but also take into account the relationships between sub-interests. Then, we can calculate the encoded vector as follows,

$$\mathbf{o}^u_t = f_{\mathbf{FFN}}(\mathbf{Norm}(\hat{\alpha}\mathbf{V} + \mathbf{q})) + \mathbf{Norm}(\hat{\alpha}\mathbf{V} + \mathbf{q}), \tag{10}$$

where $f_{\mathbf{FFN}}$ denotes a fully-connected layer, and **Norm** means the layer normalization.

After we obtain the encoded user embedding $\mathbf{o}^u_t$, we further project it into the pre-defined sub-interest prototypes space to generate a user-specific sub-interests embeddings, which is based on the context of the user's interaction history. The formulation is as follows,

$$\widetilde{\mathbf{o}}^{u,k}_t = \mathbf{o}^u_t + \sigma(\mathbf{o}^u_t * z_k) \cdot \mathbf{o}^u_t, \tag{11}$$

where k means the k-th sub-interest. After the projection, each user will have K user-specific sub-interests, which is the same number as the sub-interest prototypes.

4.2 Fusion Prediction Layer

As for the real-world information system, user behavior is triggered by multiple aspects, which we have learned the corresponding representations in sub-interest spaces. To accurately predict the next item that users will interact, it is essential to determine which aspect will play a vital role in users' current state. Therefore, we propose to first estimate the importance of sub-interests before fusing them. Specifically, we propose an attention network-based approach as follows,

$$r_k = \sigma(\mathbf{W}[\mathbf{e}_i; \widetilde{\mathbf{o}}^{u,k}_t] + \mathbf{b}), \tag{12}$$

where $\mathbf{W}$ and $\mathbf{b}$ are learnable parameters. The weights are further normalized as follows:

$$\gamma_k = \frac{r_k}{\sum_{k=1}^K r_k}. \tag{13}$$

Then the prediction score based on the learned weights and sub-interests can be formulated as follows,

$$\textbf{Score}(u, i, t) = \sum_{k=1}^{K} (\gamma_k \cdot \widetilde{\mathbf{o}}^{u,k}_t \mathbf{e}_i), \tag{14}$$

where $\mathbf{e}_i$ denotes the embedding of the target item.

Discussion. The learnable weights above are actually the generalized version of an intuitive design that only considers one sub-interest for a given behavior. For example, if the weights are 1 for one sub-interest and 0 for others, the specific sub-interest dominates the user behavior.

4.3 Joint Optimization

From the perspective of the user-algorithm feedback loop, the passive-negative feedback is collected based on the exposure of the already-deployed recommendation algorithms. Therefore, it is essential to leverage all kinds of behaviors to optimize the parameters. In our problem, there is additional negative feedback, *i.e.*, the truly-observed passive-negative feedback, besides the randomly sampled ones. Specifically, the traditional recommenders, no matter collaborative filtering or sequential recommendation, tend to sample unobserved items as negative ones for optimization.

Thus, we propose pairwise learning on both two pairs: {*positive feedback, randomly-sampled negative feedback*} and {*positive feedback, truly-observed negative feedback*}.

The pairwise loss function defined on the two kinds of pairs are as follows:

$$\begin{aligned} L_1 &= \sum_{(u,i,j) \in O_1} -ln\sigma(y(u,i) - y(u,j)), \\ L_2 &= \sum_{(u,i,j) \in O_2} -ln\sigma(y(u,i) - y(u,j)), \end{aligned} \tag{15}$$

where $O_1 = \{(u,i,j) | (u,i) \in R^+, (u,j) \in R^-\}$ denotes the set of training data, where R^+ represents observed behavior and R^- represents unobserved behavior; $O_2 = \{(u,i,j) | (u,i) \in R^+, (u,j) \in R^-\}$ denotes the set of training data, where R^+ represents observed behavior and R^- represents truly-observed passive-negative behavior. Here $\sigma(\cdot)$ denotes the sigmoid function. We also should ensure the disentanglement across the sub-interest prototypes. Therefore, we design another distance correlation loss [34], $L_{dis} = dCor(\mathbf{Z})$ which tries to maximize the distance between prototypes.

The final joint loss function is performed under a paradigm of multi-task learning [45]:

$$L = \lambda_1 L_1 + \lambda_2 L_2 + \lambda_3 L_{dis}, \tag{16}$$

where λ_1, λ_2, and λ_3 are three hyper-parameters that can control the importance of each loss, and we have $\lambda_1 + \lambda_2 + \lambda_3 = 1$.

5 EXPERIMENTS

In this section, we conduct extensive experiments on two real-world datasets in order to answer the following three research questions (RQs).

- RQ1: Can our proposed SINE method outperform the state-of-the-art solutions for sequential recommendation on two real-world datasets?
- RQ2: How about the rationality of each component in our proposed SINE method?
- RQ3: How do the introduced hyper-parameters affect the recommendation performance of our proposed SINE model?

Table 2: Statistics of the two datasets used in experiments.

Dataset	Users	Items	Instances	Positive	Negative	Average Length
Wechat	19,901	94,910	6,200,308	4,050,193	2,150,114	311
Kuaishou	37,497	126,293	9,049,176	6,399,423	2,649,753	241

5.1 Experimental Settings

5.1.1 Datasets. We conduct experiments on two large-scale datasets from two of the largest short-video platforms, Kuaishou and Wechat. Data are collected from the feed under real-world scenario. Basic statistics of the two datasets are summarized in Table 2, where we present the interaction number of both positive feedback and negative feedback. Please note that there is a minor difference in statistics compared with the data analysis in Section 2 since, for model evaluation, we have conducted a data processing of the widely-used N-core filtering [9, 43]. N is set as 10 in our experiments.

- **WeChat**[5]. This dataset was released by the Big Data Challenge in the year 2021, hosted by the recommendation team of WeChat, and it contains the behavioral logs on Wechat Channels (short-video services in the WeChat app) with a time period of two weeks. We select the last and second-last interactions as validation and test items, respectively, and use the remaining interactions as the training set. We define interaction longer than fifty percent of the video duration as positive feedback and interaction shorter than three seconds as passive-negative feedback, following the commonly accepted industrial experience. We have also tried other definitions of positive and passive-negative feedback, and the improvement of our proposed method still holds.
- **Kuaishou**[6]**.** This large-scale dataset is collected from Kuaishou, one of the most famous short-video platforms. There are billions of active users on Kuaishou, with various types of short videos, ranging from movie clips to news, as well as videos uploaded by users themselves. We utilize the behavioral logs collected from October 22 to October 28, 2020, with a one-week period. For the Kuaishou dataset, we have the same data pre-processing as the Wechat dataset.

5.1.2 Metrics. To evaluate our model and baseline models, we adopt three widely-used metrics, including AUC, GAUC and NDCG, defined as follows.

- GAUC [48] is an improved version of AUC, which alleviates the negative impact of unbalanced distributions across users. It evaluates whether the model can well rank positive items higher than negative items.
- NDCG is an improved version of Recall which evaluates whether the model can well rank positive items at top-K positions, and it also considers the specific position by assigning weights to the scores. In our experiments, we set $K = 2$ following existing works [6].

5.1.3 Baselines. We compare the proposed SINE with the following competitive recommenders to evaluate the performance.

- **SASRec** [21]: This model is the state-of-the-art sequential recommendation model with self-attention layers to capture sequential preferences as context vectors.
- **SLi-Rec** [44]: This method deploys two encoders for capturing the long-term and short-term preferences.
- **DIN** [48]: This method proposes an attention network to obtain the similarity between historical items and the target item, to calculate the interaction probability.
- **DIEN** [47]: This method extends DIN by combining a recurrent neural network.
- **CASER** [35]: This method adopts convolutional filters to extract the sequential patterns in user behaviors.
- **GRU4REC** [18]: This method utilizes a GRU network for modeling the users' sequential interactions.
- **CLSR** [46]: This method extends SLi-Rec based on disentangled representation learning, showing the state-of-the-art performance.
- **FeedRec** [38]: This model leverages various types of feedback in the sequential recommendation, and the negative feedback can be roughly treated as one kind of feedback to adapt this model to our problem.
- **SASRec-N** [21]: Although SASRec is not defined for exploiting truly negative feedback, we can still adapt it with negative sampling from both passive-negative feedback and unobserved items. We name it SASRec-N (N denotes Negative).

It is worth mentioning FeedRec [38] can be regarded as a kind of multi-feedback learning, which we will discuss in detail in Section 6. We do not include other existing methods of multi-feedback learning since they mainly focus on collaborative filtering and cannot process the sequential information.

5.1.4 Hyper-parameter Settings. We implement our SINE model and the baselines on the Microsoft Recommenders framework [2]. We use the Adam optimizer [22], carefully searching the choice of learning rate among {0.0005, 0.0007, 0.0009, 0.001, 0.003, 0.005}. The batch size is tuned among {32, 64, 128, 256, 512, 1024}. The embedding size D of all the models is set as 50 following existing work [21], to ensure fair performance comparison. Besides, we find setting the number of sub-interests to two or seven for our SINE model can both achieve good recommendation performance. We use grid-search to carefully find the best hyper-parameters, and we have released code and the best settings of hyper-parameters in *https://github.com/PanYZhu/SINE*, to benefit the community.

5.2 Overall Performance Comparison (RQ1)

We present the overall performance comparison on our adopted datasets in Table 3. Based on the results, we have the following conclusions:

[5]https://algo.weixin.qq.com/
[6]https://www.kuaishou.com

- **Our SINE achieves steady and significant improvement compared with other methods.** On the Kuaishou dataset, SINE outperforms the best baseline, SASRec, by 5.32%, 3.95%, and 1.88%, on AUC, GAUC, and NDCG, respectively. On the WeChat dataset, SINE outperforms the best baseline, SASRec, by 1.56%, average on AUC, GAUC, and NDCG. It is acknowledged 1%-level improvements in AUC, GAUC, and NDCG@2 can be claimed as significant [38, 46, 47].
- **The modeling of negative feedback should be carefully designed.** FeedRec and SASRec-N are two models that can utilize the negative feedback in modeling, but however, FeedRec performances worse than SASRec-N. This can be explained that the negative feedback may worsen the recommendation performance without proper utilization manner. Specifically, the observed passive-negative behaviors reflect weaker "dislike" signals compared with active-negative ones.
- **Only modeling positive feedback may achieve very poor performance.** Some competitive methods, such as CLSR and SLi-Rec, have shown good performance in datasets of traditional sequential recommendation. But, however, these methods perform poorly in Kuaishou and Wechat datasets, since that positive feedback is far sparser in the short-video recommendation. This observation further supports this work's motivation of modeling the passive-negative feedback.

5.3 Ablation Study (RQ2)

In this section, we conduct experiments to study the impact of some key components, including negative feedback learning and adaptive fusion. We show the performance results in Table 4, where "AF" denotes the adaptive fusion-based prediction and "NF" denotes learning from negative feedback.

5.3.1 Adaptive fusion. We remove the adaptive fusion, and then there is always one specific sub-interest activated in determining user behavior. The result shows a performance drop of 0.49-5.56%, and thus it is essential to well handle the role of different sub-interests since each behavior may be affected by multiple sub-interests.

5.3.2 Negative feedback learning. We remove all the negative feedback learning in our SINE model, and then the sub-interest design will not exist. It means that the SINE model will degenerate into a basic model that only leverages positive feedback. The result shows a performance drop of 1.21-6.42%, and it illustrates the significance of our proposed whole framework of sub-interest-based negative feedback learning.

5.4 Hyper-parameter Study (RQ3)

In this section, we conduct experiments on two datasets to study the impact of three important hyper-parameters in our SINE model, including the number of user interests, learning rate, and batch size.

5.4.1 The number of sub-interests. The number of sub-interests is an important hyper-parameter in the our SINE model. To explore its impact on our model's performance, we conduct experiments on the number of sub-interests in the range of {1, 2, 3, 4, 5, 6, 7, 8, 9, 10}. The results on two datasets are displayed in Fig. 4, from which we have the following observations.

- **Our proposed sub-interest based encoder is essential and effective in modeling user's preference.** From the results in Fig. 4, there is a significant performance increase in models that capture two sub-interests compared to models that only capture one sub-interest. This phenomenon can be observed in both datasets across all metrics. As previously discussed, we believe that users have multiple sub-interests, but only one sub-interest is active during a certain period, while the rest of the sub-interests are relatively inactive. Our proposed method utilizes users' passive-negative feedback to distinguish which sub-interest is the current active one. When there is only one sub-interest, our model is regressed to original SASRec model. However, with only two sub-interests being used to capture user's preference, our model achieves distinct improvements compared with one sub-interest, which verifies the effectiveness of our proposed approach.
- **The number of user interests is customized to specific dataset.** The results indicate that the optimal number of sub-interests for modeling the kuaishou dataset is two and for the wechat dataset is seven. As shown in Fig. 4, the overall performance of both datasets exhibit a similar trend. However, upon analyzing each dataset individually, we can observe that the best results are obtained at different numbers of sub-interests. These results demonstrate that different platforms possess unique characteristics.

5.4.2 Learning rate & Batch size. We carefully tune learning rate in the range of {0.0005, 0.0007, 0.0009, 0.001, 0.003, 0.005} and batch size in the range of {32, 64, 128, 256, 512, 1024}, following the existing works. Our experiments revealed that the best performance was achieved with a learning rate of 0.003 and a batch size of 32.

To summarize, we conduct experiments on two large-scale real-world datasets, and the results show our SINE's better performance compared with SOTA models. Further experiments of ablation study verify the rationality of our model design.

6 RELATED WORK

In this section, we would like to discuss the related works based on the following three perspectives, including sequential recommendation, multi-feedback learning in recommendation, and sub-interest learning for recommendation. We emphasize why these methods cannot well address the studied new problem in this paper.

6.1 Sequential Recommendation

Sequential recommendation is defined as to recommend the next interacted item based on the historically-interacted item sequence. The early non-deep-learning approaches [30] used the Markov chain to model the transition between items in one sequence. Recently, deep learning methods [13, 18, 21, 24–26, 33, 35, 47, 49] have become the mainstream solution for sequential recommendation. Kang *et al. [21]* propose to build self-attention layers to encode the users' sequence.

Some recent works [1, 19, 46] argued that existing sequential recommenders tend to focus on the recent interaction and ignore the

Table 3: Overall performance comparison of our SINE model and the baseline models on two datasets (the best performance is marked with bold font, and the best baseline is marked with underline).

Method	Kuaishou			Wechat		
	AUC	GAUC	NDCG@2	AUC	GAUC	NDCG@2
SASRec	0.5830	0.5305	<u>0.7002</u>	<u>0.6454</u>	<u>0.5408</u>	<u>0.6654</u>
CLSR	0.5676	<u>0.5617</u>	0.6428	0.6423	0.5325	0.5573
SLi-Rec	0.5350	0.5252	0.6177	0.5901	0.5361	0.5429
DIN	<u>0.6111</u>	0.5216	0.4766	0.6379	0.5173	0.4557
GRU4REC	0.5483	0.5278	0.5847	0.6193	0.5251	0.551
DIEN	0.5840	0.5346	0.5917	0.6497	0.5361	0.5413
CASER	0.5617	0.5260	0.6136	0.5969	0.5306	0.5496
NCF	0.5049	0.5221	0.5961	0.6203	0.5301	0.5434
SASRec-N	0.5509	0.5200	0.6958	0.5679	0.5245	0.6622
FeedRec	0.5367	0.5102	0.6922	0.5470	0.5249	0.6612
Our SINE	**0.6362**	**0.5700**	**0.7190**	**0.6609**	**0.5623**	**0.6752**

Table 4: Ablation study of our proposed components. "AF" denotes adaptive-fusion-based prediction, and "NF" denotes learning from negative feedback.

Method	Kuaishou			Wechat		
	AUC	GAUC	NDCG	AUC	GAUC	NDCG
w/o AF	0.6204	0.5411	0.7055	0.6053	0.5483	0.6703
w/o NF	0.5720	0.5295	0.6988	0.6076	0.5293	0.6631
Our SINE	**0.6362**	**0.5700**	**0.7190**	**0.6609**	**0.5623**	**0.6752**

long-term user behaviors, and propose various solutions for modeling both long-term and short-term user preferences. In industry, sequential recommenders in the large-scale system are also widely concerned [3, 28]. Pi *et al.* [28] propose to use search behavior to retrieve the similar item, addressing the challenge of modeling too-long sequences. Cao *et al.* [3] further propose new sampling strategies which can shorten the orginal long user behavioral sequences.

However, these works of sequential recommendation always mainly model the sequence of positive feedback, ignoring the widely-existed negative feedback. This could be explained by the fact that most of the existing works make use of the e-commerce datasets; however, in the recent applications of short-video recommendation, the negative behavior (especially for the passive skipping-over behavior) widely exists and is very important for modeling user preferences, which is the focus of this paper.

6.2 Multi-feedback learning in recommendation

It is not new for recommender systems to utilize multiple types of behaviors. Multi-behavior recommendation [12, 20, 40] is one of the relevant research problems, which is defined as to leverage the multiple types of user feedback, such as click, adding-to-cart, purchase, etc., in e-commerce websites or click, like, share, etc. in online social media. For example, Jin *et al.* [20] build a multi-relational user-item graph to represent the multi-form feedback between users and items, and then developed graph neural network models to predict the missing links on the graph. That is, these works tend to study the multi-type positive feedback, leaving negative feedback less explored.

Another close topic is exposure bias-aware recommendation [7, 15, 27], in which the expose-but-not-click is also considered as a kind of feedback, and it is related to the passive-negative feedback in this work. Other works approach the implicit feedback modeling by designing negative sampling strategy via exposure data [10, 11]. The recent negative feedback modeling works [14, 32, 37, 41] also explore this expose-but-not-click negative feedback. Wang *et al.*[37] focuses on unbiased recommendation that introduces negative feedback modeling and Seo *et al.*[32] is a graph-based recommendation. However, there are two folds of critical differences compared with our work. First, the user may have not noticed the exposed items as the utilized datasets in these works still require the user to actively click. Second, the sequential behaviors are not well considered in these works.

6.3 Sub-interest learning for recommendation

Users in real-world always have multiple criteria to judge whether the recommended item satisfies their needs or not. For example, a user in the e-commerce website will make decisions according to the price, brand, functionality, etc. Therefore, different from the traditional recommendation models with only one general interest representation [16], researchers have begun to model multiple distinct interests of users, which are also known as sub-interests [5, 8, 23, 31, 36].

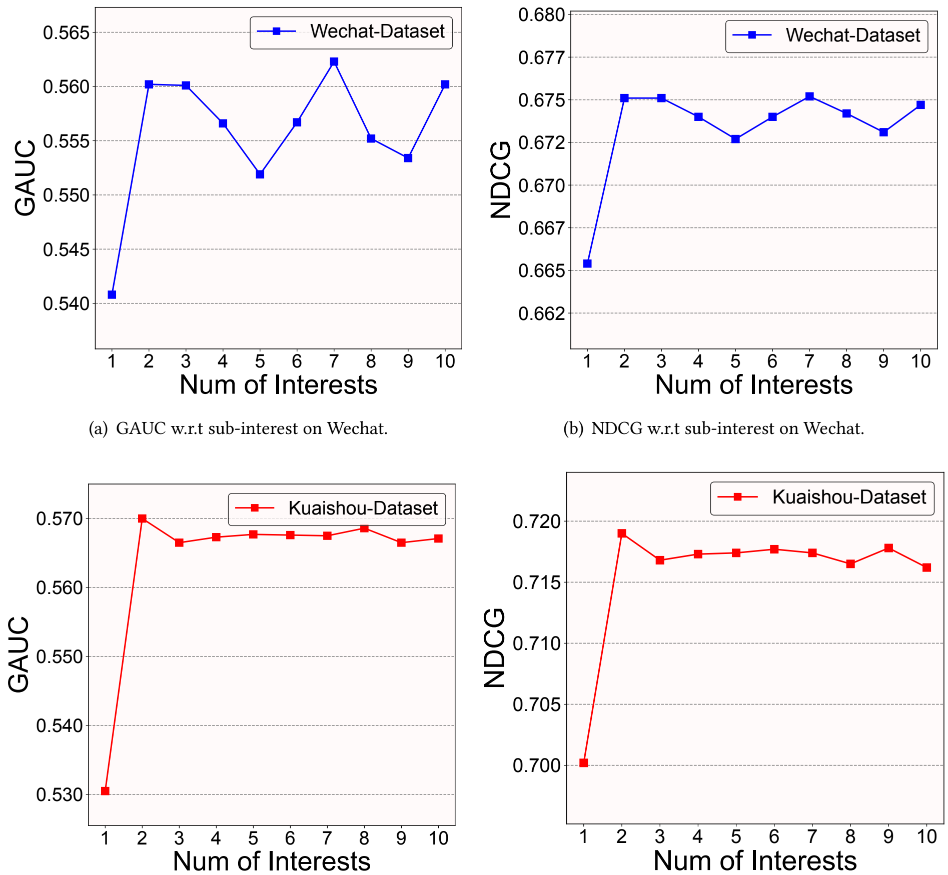

(a) GAUC w.r.t sub-interest on Wechat.

(b) NDCG w.r.t sub-interest on Wechat.

(c) GAUC w.r.t sub-interest on Kuaishou.

(d) NDCG w.r.t sub-interest on Kuaishou.

Figure 4: The impact of the number of user interests on recommendation performance of our SINE model.

Li *et al.* [23] proposed to use capsule network [31] to extract the users' sub-interests and combined it with a label-aware attention layer. Chai *et al.* [5] combines capsule network with user profile information to refine user's sub-interest for enhancing recommendation. Cen *et al.* [4] propose two kinds of architectures, capsule network, and self-attention network, with a controllable procedure to balance the accuracy and diversity of recommendation. Similarly, Chen *et al.* [8]also extracts sub-interests with capsule network and self-attention network. Recently, Tian *et al.* [36] propose to combine multi-interest learning with multi-grained interest learning modules.

However, the sub-interests of these works lack explicit supervision signals, and the learned sub-interests are always not explainable. Different from them, in our work, we use the sub-interests to explain the occurrence of passive-negative feedback, which can serve as good supervision for the representation learning of sub-interests.

7 CONCLUSION AND FUTURE WORK

In this work, we approach a new problem in sequential recommendation systems: understanding and modeling of users' passive-negative feedback in short-video recommendation. We first use the data analysis on large-scale real-world data to demonstrate 1) it is important, 2) it is challenging, and 3) why it is challenging to utilize passive-negative feedback in sequential recommenders. The results show that the passive-negative feedback is similar to the observed positive feedback in video category and thus the intuitive manner of treating it as negative sample does not work well. Motivated by the data analysis, we propose a method with a sub-interest extractor, in which the passive-negative feedback can be modeled as the mismatch of specific sub-interests. The performance comparison shows the best performance of our proposed SINE method. The further results extensively demonstrate the effectiveness of SINE's different components, and the results correspond well to the earlier data analysis.

As for the future work, we first plan to deploy the proposed SINE model to the real-world online recommendation engine, and evaluate the recommendation performance compared with other methods through the A/B tests. We also plan to collect other kinds of explicit positive feedback, such as like behavior, sharing behavior, etc., to enhance the study of negative feedback in recommender systems, which can motivate us to further improve the model design.

ACKNOWLEDGMENTS

This work is partially supported by the National Key Research and Development Program of China under 2022YFB3104702, and the National Natural Science Foundation of China under 62272262, 61972223, U1936217, and U20B2060.

REFERENCES

[1] Mingxiao An, Fangzhao Wu, Chuhan Wu, Kun Zhang, Zheng Liu, and Xing Xie. 2019. Neural news recommendation with long-and short-term user representations. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 336–345.

[2] Andreas Argyriou, Miguel González-Fierro, and Le Zhang. 2020. Microsoft Recommenders: Best Practices for Production-Ready Recommendation Systems. In *Companion Proceedings of the Web Conference 2020*. 50–51.

[3] 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 *CIKM*.

[4] Yukuo Cen, Jianwei Zhang, Xu Zou, Chang Zhou, Hongxia Yang, and Jie Tang. 2020. Controllable multi-interest framework for recommendation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2942–2951.

[5] Zheng Chai, Zhihong Chen, Chenliang Li, Rong Xiao, Houyi Li, Jiawei Wu, Jingxu Chen, and Haihong Tang. 2022. User-Aware Multi-Interest Learning for Candidate Matching in Recommenders. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1326–1335.

[6] Jianxin Chang, Chen Gao, Yu Zheng, Yiqun Hui, Yanan Niu, Yang Song, Depeng Jin, and Yong Li. 2021. Sequential Recommendation with Graph Neural Networks. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 378–387.

[7] Jiawei Chen, Hande Dong, Yang Qiu, Xiangnan He, Xin Xin, Liang Chen, Guli Lin, and Keping Yang. 2021. AutoDebias: Learning to debias for recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 21–30.

[8] Wanyu Chen, Pengjie Ren, Fei Cai, Fei Sun, and Maarten De Rijke. 2021. Multi-interest diversification for end-to-end sequential recommendation. *ACM Transactions on Information Systems (TOIS)* 40, 1 (2021), 1–30.

[9] Junsu Cho, Dongmin Hyun, Seongku Kang, and Hwanjo Yu. 2021. Learning heterogeneous temporal patterns of user preference for timely recommendation. In *Proceedings of the Web Conference 2021*. 1274–1283.

[10] Jingtao Ding, Yuhan Quan, Xiangnan He, Yong Li, and Depeng Jin. 2019. Reinforced Negative Sampling for Recommendation with Exposure Data.. In *IJCAI*. Macao, 2230–2236.

[11] Jingtao Ding, Yuhan Quan, Quanming Yao, Yong Li, and Depeng Jin. 2020. Simplify and robustify negative sampling for implicit collaborative filtering. *Advances in Neural Information Processing Systems* 33 (2020), 1094–1105.

[12] Chen Gao, Yu Zheng, Nian Li, Yinfeng Li, Yingrong Qin, Jinghua Piao, Yuhan Quan, Jianxin Chang, Depeng Jin, Xiangnan He, et al. 2023. A survey of graph neural networks for recommender systems: challenges, methods, and directions. *ACM Transactions on Recommender Systems* 1, 1 (2023), 1–51.

[13] Chen Gao, Yu Zheng, Wenjie Wang, Fuli Feng, Xiangnan He, and Yong Li. 2022. Causal Inference in Recommender Systems: A Survey and Future Directions. *arXiv preprint arXiv:2208.12397* (2022).

[14] Shansan Gong and Kenny Q Zhu. 2022. Positive, Negative and Neutral: Modeling Implicit Feedback in Session-based News Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1185–1195.

[15] Shantanu Gupta, Hao Wang, Zachary Lipton, and Yuyang Wang. 2021. Correcting exposure bias for link recommendation. In *International Conference on Machine Learning*. PMLR, 3953–3963.

[16] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.

[17] Xiangnan He, Hanwang Zhang, Min-Yen Kan, and Tat-Seng Chua. 2016. Fast matrix factorization for online recommendation with implicit feedback. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. 549–558.

[18] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939* (2015).

[19] Linmei Hu, Chen Li, Chuan Shi, Cheng Yang, and Chao Shao. 2020. Graph neural news recommendation with long-term and short-term interest modeling. *Information Processing & Management* 57, 2 (2020), 102142.

[20] Bowen Jin, Chen Gao, Xiangnan He, Depeng Jin, and Yong Li. 2020. Multi-behavior recommendation with graph convolutional networks. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 659–668.

[21] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 197–206.

[22] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[23] Chao Li, Zhiyuan Liu, Mengmeng Wu, Yuchi Xu, Huan Zhao, Pipei Huang, Guoliang Kang, Qiwei Chen, Wei Li, and Dik Lun Lee. 2019. Multi-interest network with dynamic routing for recommendation at Tmall. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 2615–2623.

[24] Jing Li, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Tao Lian, and Jun Ma. 2017. Neural attentive session-based recommendation. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 1419–1428.

[25] Jiacheng Li, Yujie Wang, and Julian McAuley. 2020. Time interval aware self-attention for sequential recommendation. In *Proceedings of the 13th international conference on web search and data mining*. 322–330.

[26] Chen Ma, Liheng Ma, Yingxue Zhang, Jianing Sun, Xue Liu, and Mark Coates. 2020. Memory augmented graph neural networks for sequential recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 5045–5052.

[27] Masoud Mansoury, Himan Abdollahpouri, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. 2021. A graph-based approach for mitigating multi-sided exposure bias in recommender systems. *ACM Transactions on Information Systems (TOIS)* 40, 2 (2021), 1–31.

[28] 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 *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 2685–2692.

[29] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).

[30] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing personalized markov chains for next-basket recommendation. In *Proceedings of the 19th international conference on World wide web*. 811–820.

[31] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. 2017. Dynamic routing between capsules. *Advances in neural information processing systems* 30 (2017).

[32] Changwon Seo, Kyeong-Joong Jeong, Sungsu Lim, and Won-Yong Shin. 2022. SiReN: Sign-Aware Recommendation Using Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems* (2022).

[33] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 1441–1450.

[34] Gábor J Székely, Maria L Rizzo, and Nail K Bakirov. 2007. Measuring and testing dependence by correlation of distances. *The annals of statistics* 35, 6 (2007), 2769–2794.

[35] Jiaxi Tang and Ke Wang. 2018. Personalized top-n sequential recommendation via convolutional sequence embedding. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. 565–573.

[36] Yu Tian, Jianxin Chang, Yanan Niu, Yang Song, and Chenliang Li. 2022. When Multi-Level Meets Multi-Interest: A Multi-Grained Neural Model for Sequential Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1632–1641.

[37] Zhenlei Wang, Shiqi Shen, Zhipeng Wang, Bo Chen, Xu Chen, and Ji-Rong Wen. 2022. Unbiased sequential recommendation with latent confounders. In *Proceedings of the ACM Web Conference 2022*. 2195–2204.

[38] Chuhan Wu, Fangzhao Wu, Tao Qi, Qi Liu, Xuan Tian, Jie Li, Wei He, Yongfeng Huang, and Xing Xie. 2022. Feedrec: News feed recommendation with various user feedbacks. In *Proceedings of the ACM Web Conference 2022*. 2088–2097.

[39] Le Wu, Xiangnan He, Xiang Wang, Kun Zhang, and Meng Wang. 2022. A survey on accuracy-oriented neural recommendation: From collaborative filtering to information-rich recommendation. *IEEE Transactions on Knowledge and Data Engineering* (2022).

[40] Lianghao Xia, Yong Xu, Chao Huang, Peng Dai, and Liefeng Bo. 2021. Graph meta network for multi-behavior recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 757–766.

[41] Ruobing Xie, Cheng Ling, Yalong Wang, Rui Wang, Feng Xia, and Leyu Lin. 2020. Deep Feedback Network for Recommendation. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, Christian Bessiere (Ed.). ijcai.org, 2519–2525. https://doi.org/10.24963/ijcai.2020/349

[42] Hong-Jian Xue, Xinyu Dai, Jianbing Zhang, Shujian Huang, and Jiajun Chen. 2017. Deep matrix factorization models for recommender systems.. In *IJCAI*, Vol. 17. Melbourne, Australia, 3203–3209.

[43] Wenhui Yu, Xiao Lin, Jinfei Liu, Junfeng Ge, Wenwu Ou, and Zheng Qin. 2021. Self-propagation graph neural network for recommendation. *IEEE Transactions on Knowledge and Data Engineering* (2021).

[44] Zeping Yu, Jianxun Lian, Ahmad Mahmoody, Gongshen Liu, and Xing Xie. 2019. Adaptive User Modeling with Long and Short-Term Preferences for Personalized Recommendation.. In *IJCAI*. 4213–4219.

[45] Yu Zhang and Qiang Yang. 2018. An overview of multi-task learning. *National Science Review* 5, 1 (2018), 30–43.

[46] Yu Zheng, Chen Gao, Jianxin Chang, Yanan Niu, Yang Song, Depeng Jin, and Yong Li. 2022. Disentangling Long and Short-Term Interests for Recommendation. In *Proceedings of the ACM Web Conference 2022*. 2256–2267.

[47] 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 *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 5941–5948.

[48] 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.

[49] 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*, Vol. 17. 3602–3608.