



Attacking Pre-trained Recommendation

Yiqing Wu^{*†}

Institute of Computing Technology,
Chinese Academy of Sciences
Beijing, China
wuyiqing20s@ict.ac.cn

Ruobing Xie[†]

WeChat, Tencent
Beijing, China
ruobingxie@tencent.com

Zhao Zhang[‡]

Yongchun Zhu^{*}
Institute of Computing Technology,
Chinese Academy of Sciences
Beijing, China
{zhangzhao2021,zhuyongchun18s}@ict.ac.cn

Fuzhen Zhuang[§]

Institute of Artificial Intelligence,
Beihang University
Beijing, China
zhuangfuzhen@buaa.edu.cn

Jie Zhou

WeChat, Tencent
Beijing, China
withtomzhou@tencent.com

Yongjun Xu

Qing He^{*}
Institute of Computing Technology,
Chinese Academy of Sciences
Beijing, China
{xyj,heqing}@ict.ac.cn

ABSTRACT

Recently, a series of pioneer studies have shown the potency of pre-trained models in sequential recommendation, illuminating the path of building an omniscient unified pre-trained recommendation model for different downstream recommendation tasks. Despite these advancements, the vulnerabilities of classical recommender systems also exist in pre-trained recommendation in a new form, while the security of pre-trained recommendation model is still unexplored, which may threaten its widely practical applications. In this study, we propose a novel framework for backdoor attacking in pre-trained recommendation. We demonstrate the provider of the pre-trained model can easily insert a backdoor in pre-training, thereby increasing the exposure rates of target items to target user groups. Specifically, we design two novel and effective backdoor attacks: basic replacement and prompt-enhanced, under various recommendation pre-training usage scenarios. Experimental results on real-world datasets show that our proposed attack strategies significantly improve the exposure rates of target items to target users by hundreds of times in comparison to the clean model. The source codes are released in <https://github.com/wyqing20/APRec>.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

pre-trained recommendation, attack, security.

^{*}Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS). Also at University of Chinese Academy of Sciences, China.

[†]Yiqing Wu and Ruobing Xie contributed equally to this research.

^{*}Corresponding author.

[§]Fuzhen Zhuang is also at Zhongguancun Laboratory, Beijing, China



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

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

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9408-6/23/07.

<https://doi.org/10.1145/3539618.3591949>

ACM Reference Format:

Yiqing Wu, Ruobing Xie, Zhao Zhang, Yongchun Zhu, Fuzhen Zhuang, Jie Zhou, Yongjun Xu, and Qing He. 2023. Attacking Pre-trained Recommendation. 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/10.1145/3539618.3591949>

1 INTRODUCTION

Recommender systems aim to personally provide appropriate items for users according to their preferences, mainly hidden in their historical behavior sequences. Recently, sequential recommendation (SR) has achieved great success and is widely applied in practice [4, 28], which takes users' historical behavior sequences as inputs and outputs the predicted items. Lots of effective sequential modeling methods have been verified in SR [7, 16]. Recently, inspired by the overwhelming power of large-scale pre-training models [1], some pioneer efforts bring pre-training into sequential recommendation and achieve great successes [17, 19, 21, 25]. Some ambitious works even explore building a unified big pre-training recommendation model for various downstream tasks [2, 5, 20]. The broad usage of recommendation pre-training is promising.

Due to the significant social impact and commercial value of recommender systems, the security of recommender systems is a crucial matter. For example, in E-commerce, an attacker could mislead the model to recommend target items via fake users and cheating behaviors. Existing works have shown the vulnerabilities and security threats of conventional recommendation [8, 9, 12, 23]. In these studies, attacks typically involve two roles: the *platform* and the *users*. Malicious users could poison log data (i.e., user behavior history) by natural interactions to attack the platform and manipulate the system to deliver their desired results [18, 24, 27]. Unfortunately, despite the progress of pre-training recommendation, the security of pre-training recommendation model is still without consideration, and recommendation model encounters new threats under pre-training. Different from conventional recommendation, there is an additional role besides *platform* and the *users* in the novel pre-training recommendation paradigm: *pre-trained model provider*. The pre-trained model may be actively (e.g., plant backdoors) or passively (e.g., fake users) poisoned by certain attacks,

which consequently harms the downstream platforms. In this case, attacks in public pre-training recommendation models are much more covert and hard to be explicitly detected and controlled by the platform since the pre-trained model is fixed after black-box training.

In this work, we conduct a pioneer exploration of the new attack paradigm in pre-training recommendation. We design two backdoor attacks to improve the exposure rates of certain target items to target users (i.e., global or a group of users). Specifically, we first design a straightforward and effective *basic replacement strategy* to generate fake user behavior sequences involving extra target items. Through this basic attack, the exposure rates of target items can be increased hundreds of times. Moreover, to enhance the attack effectiveness for user group attacks, we propose a *prompt-enhanced attacking strategy* against the mainstream tuning paradigm "prompt-tuning". Specifically, we design a three-step training framework that simulates the process of prompt-tuning, enabling facilitate more precise and harder-to-detect attacks on target user groups with minor influence on other user groups in different settings. Additionally, we also propose a preliminary pre-trained model anomaly detection method against these new types of pre-training backdoor attacks.

Overall, we are the first to systematically demonstrate the threat of backdoor attacks on pre-trained recommendation by the proposed backdoor attacking method. We explore this new attacking paradigm in pre-trained recommendation to remind developers and researchers of the security concerns that have been ignored but are extremely important, ringing the alarm bell.

2 PRELIMINARIES & RELATED WORKS

Recommendation Pre-training. Recommendation pre-training mainly has three parts, including *model provider*, *platform*, and *user*. Let S and P be the model provider and platform. We denote the unaligned user sets of model provider and platform as U^S and U^P . In this work, we assume the item set is shared (similar scenarios widely exist in movie/book/music recommendation. Future large-scale pre-trained recommendation models may also involve nearly all items[5]), noted as V . Each user $u \in (U^P \cup U^S)$ has a historical behavior sequence $s_u = \{v_1^u, v_2^u, \dots, v_{|s_u|}^u\}$ (of length $|s_u|$) ordered by time. Besides, each user has m attributes $A_u = \{a_1^u, a_2^u, \dots, a_m^u\}$ (i.e., user profiles such as age and gender). The model provider pre-trains a large sequential model $f_{seq}(s_u^s | \Theta)$ on all pre-training data of U^S , which is then adopted by the downstream platform to serve its users. In general, platforms can use pre-trained models in three ways: direct using, fine-tuning, and prompt-tuning. For direct using, the platform deploys the pre-trained model for downstream tasks without any change. For fine-tuning, the platform tunes the pre-trained model on its private data. For prompt-tuning, the platform conducts a parameter-efficient tuning paradigm with prompts [10, 20, 26] (see Sec.3.2 for details of prompt-tuning).

Threat Model. The attacker's goal can be classified into promotion/demotion attacks and availability attacks[15]. Availability attacks aim to make the recommender system unserviceable, which is easy to detect and meaningless for our pre-training attack setting. Promotion/demotion attacks are designed to promote/demote the recommended frequency of the target items [18, 24, 27]. In this

work, we focus on *promotion attacks*. An attacker's goal is to promote the exposure rate of target items to target users in a top-K recommender system. Here, target users can be all users or specific user groups in the platform.

Private data: The model provider's and platform's behavior data is private and cannot be observed by each other since user data cannot be public due to privacy protection.

3 METHODOLOGY

3.1 Basic Replacement Attack

We first introduce the pre-training part of our study. We adopt next-item prediction as our pre-training task, which is a natural and widely adopted pre-training task in recommendation [2, 5, 22, 29]. Given a user behavior sequence $s_u = \{v_1^u, v_2^u, \dots, v_{|s_u|}^u\}$, the goal of next-item prediction is predicting items that the user may interact with. We formulate it as:

$$u_s^t = f_{seq}(\{v_1^u, \dots, v_t^u\} | \Theta), \quad L = \min_{u \in U^S} \sum_{t=1}^{|s_u|} l(u_s^t, v_{t+1}^u), \quad (1)$$

where f_{seq} is the pre-training model and Θ is the parameters. l is the loss function, we adopt the classical Bayesian Personalized Ranking (BPR) loss as our loss [14].

Previous studies show generating fake user behavior sequences is an effective way to promote target items[6, 24, 27]. As the recommendation model is trained on user behavior data. In recommendation pre-training, intuitively, we expect the generated fake user sequences to have the following two key features: (1) they should be as similar as possible to those of natural users so that the attacked pre-train model is available on the downstream platform, in which users are natural real user. (2) they should have as little impact as possible on the normal recommendation performance. Based on this, we design a simple but effective and efficient random replacement strategy. Specifically, given a behavior sequence $s_u = \{v_1^u, v_2^u, \dots, v_{|s_u|}^u\}$ and corresponding ground truth item $v_{|s_u|+1}^u$, we first find target items that similar with ground truth item. Then we replace this ground truth with a randomly selected item from similar items with probability r . In this way, we can generate natural fake user behavior sequences. Actually, those sequences are created by real users except for the replaced item. Moreover, by adjusting the replacement rate r and the measure of similarity, we can balance the attack effect and recommendation effect. In this paper, we regard the items under the same category as similar items.

3.2 Prompt-enhanced Attack

Prompt tuning is becoming a new popular paradigm for utilizing pre-trained models on downstream tasks [10, 11, 20, 26], which is effective and efficient. Prompt also shows its power to manipulate sequential models' output[1, 13]. Considering the strengths of prompt, we design a prompt-enhanced attack for user group attacks. A prompt is often a small piece of hard text or soft embeddings inserted into the original sequence, which helps to efficiently extract knowledge from the pre-training models for the downstream tasks. During tuning, only the prompts (having much fewer parameters) will be updated, with the whole pre-trained model unchanged. As recommendation emphasize personalized, prompts generally

are generated by personalized information (e.g., user profile). We formulate the prompt-tuning as follows:

$$\begin{aligned} p_1^u, p_2^u, \dots, p_n^u &= \text{Generator}(A_u) \\ s_u^t &= \{p_1^u, p_2^u, \dots, p_n^u, v_1^u, v_2^u, \dots, v_t^u\}. \end{aligned} \quad (2)$$

where the $\{p_1^u, p_2^u, \dots, p_n^u\}$ is generated prompt, and ϑ is the prompt generator's parameters. There we adopt MLP as our generator. We freeze Θ when tuning the prompt.

The prompt can be provided by the model provider. In this case, the prompt can be trained together with the pre-trained model. Our experiments show the effectiveness of prompts, especially in user group attacks. However, typically, the platform trains prompt on its own private data. In this case, it is challenging to implement prompt-enhanced attacks. As the model provider (attacker) does not know the parameters of prompts and private data of the platform. Experiments also show the ineffectiveness of joint training. To solve this challenge, we propose a novel *three-step training framework*: (1) *Step 1*: Pre-train a sequential model on the provider's data. This step aims to build a clean model for downstream tuning. (2) *Step 2*: Freeze the pre-trained model's parameters and conduct prompt tuning on the model provider's private data. The goal is to simulate the prompt-tuning of the platform and obtain fake prompts. (3) *Step 3*: Freeze the prompt's parameters and tune the sequential model with our basic replacement attack strategy. In this way, the attacked pre-trained model will react to the fake prompts and achieve the goal of manipulating the recommendation system.

After the three-step training, we implant the backdoor into the pre-trained model. The backdoor will be triggered after the platform conducts prompt-tuning on its private data.

4 EXPERIMENTS

4.1 Dataset

We evaluate our attack method on two real-world open datasets, namely CIKM and AliEC. We assume that the platform has fewer data and lower user activity, and therefore needs to use pre-trained models. To simulate the roles of the *model provider* and downstream *platform*, we partition the datasets into two subsets. We treat users with fewer than ten interactions as platform users and others as users in the model provider. **CIKM**. The CIKM dataset is an E-commerce recommendation dataset released by Alibaba¹. There are 130 thousand items in this dataset. The model provider has 60 thousand users with 2 million click instances. The platform has 31 thousand users with 200 thousand clicks. In user group attack, users having a certain common attribute (i.e., gender) are viewed as the target user group. **AliEC**. AliEC is an E-commerce recommendation dataset. It contains 109 thousand items. There are 98 thousand users with 8.9 million click instances in the model provider, and over 6 thousand users with 30 thousand click instances in the platform. Users in the same age range are viewed as the target user group.

4.2 Experiments Setting

In this work, we adopt the Transformer-based sequential model SASRec[7] as our base model. As discussed before, **as the scenarios and settings differ from previous attack methods**, it is not

appropriate to compare them with other methods, and therefore we conduct an internal comparison of our proposed method. (1) **Local model (Local)**: Local model is the model that is trained on the platform's local private data. (2) **Pre-trained clean model (Clean)**: *Clean* is the clean model that is trained on the model provider's private data. (3) **Basic Replacement Attack (BRA)**: *BRA* is the pre-trained model attacked by our basic replacement attack, and it is directly deployed on the platform. (4) **Basic Replacement Attack + Fine-tuning (BRA^{FT})**: BRA^{FT} is the model that is fine-tuned on the platform's private data based on *BRA*. (5) **Prompt-enhanced Attack + Directly Using (PEA^D)**: PEA^D is the prompt-enhanced method, where both pre-trained model and prompts given by the model provider are directly used in the platform. (6) **Prompt-enhanced Attack + Prompt-tuning (PEA^{PT})**: PEA^{PT} adopts a three-step training framework specially for prompt-enhanced attack. The platform conducts prompt-tuning with the platform's private data for the prompt-enhanced pre-trained model.

Parameter Settings For parameters, we set the embedding size to 64 for all methods. The replace rate r is 0.5 for all methods. We conduct a grid search for hyper-parameters. The L2 normalization coefficient is set to $1e - 6$. The learning rate is set to $1e - 4$ for recommendation pre-training (i.e., model provider) and $1e - 5$ for recommendation tuning (i.e., platform). We randomly select 300 items as target items in attack. For CIKM, the target user is male. For AliEC, the target users are aged between 20 and 30 years.

Evaluation, we adopt the classical leave-one-out strategy to evaluate the performance [3, 7, 22]. We use widely accepted HIT@N, and NDCG@N to measure the recommendation accuracy.

Table 1: Results of global attack evaluation. All improvements are significant over baselines (t-test with $p < 0.05$).

Dataset	Model	H@5	N@5	H@10	N@10
CIKM	<i>Local</i>	0.0002	0.0001	0.0004	0.0002
	<i>Clean</i>	0.002	0.001	0.004	0.002
	<i>BRA</i>	0.649	0.338	1.323	0.554
	BRA^{FT}	0.456	0.239	0.901	0.381
AliEC	<i>Local</i>	0.002	0.001	0.004	0.002
	<i>Clean</i>	0.010	0.006	0.016	0.008
	<i>BRA</i>	0.789	0.444	1.656	0.721
	BRA^{FT}	0.639	0.357	1.353	0.585

4.3 Global Attack Evaluation

In global attack evaluation, we attempt to promote the exposure rate of target items on all users in the platform. Table 1 shows the performance of our global attacks. We can find that: (1) *BRA* successfully implants a backdoor into pre-trained model. Compared to *Local*, *BRA*'s HIT@5 of target items (indicating attack success rate) increases by more than several hundred times. (2) Compared with *BRA* and BRA^{FT} , the attack effect decreases somewhat through fine-tuning the pre-trained model, but it still improves by several hundred times. This not only demonstrates the effectiveness of our attack but also indicates that fine-tuning is a relatively good defense mechanism against our backdoor attack. (3) For recommendation accuracy, Our *BRA* models' results significantly outperform the *Local* model, even based on the attacked pre-trained models. *BRA*'s

¹<https://tianchi.aliyun.com/competition/entrance/231719/introduction>

Table 2: User Group Attack Evaluation. λ denotes the average attack success rate gaps of target users and non-target users in HIT@N and NDCG@N. All improvements are significant over baselines (t-test with $p < 0.05$).

Dataset	Model	Target user				Non-target user				λ
		H@5	N@5	H@10	N@10	H@5	N@5	H@10	N@10	
CIKM	<i>Clean</i>	0.002	0.001	0.004	0.002	0.002	0.002	0.004	0.002	0.9x
	<i>BRA</i>	0.755	0.388	1.798	0.72	0.234	0.116	0.610	0.236	3.1x
	<i>PEA^D</i>	1.437	0.786	3.056	1.302	0.072	0.037	0.017	0.0	44.6x
	<i>BRA^{FT}</i>	0.033	0.016	0.106	0.040	0.016	0.008	0.047	0.018	2.1x
	<i>PEA^{PT}</i>	0.725	0.393	1.586	0.667	0.064	0.034	0.151	0.061	11.1x
AliEC	<i>Clean</i>	0.009	0.006	0.016	0.008	0.007	0.004	0.014	0.006	1.3x
	<i>BRA</i>	0.272	0.148	0.671	0.269	0.114	0.057	0.288	0.113	2.4x
	<i>PEA^D</i>	1.649	0.972	3.303	1.499	0.033	0.018	0.067	0.029	51.2x
	<i>BRA^{FT}</i>	0.272	0.144	0.660	0.267	0.112	0.051	0.284	0.111	2.5x
	<i>PEA^{PT}</i>	0.597	0.335	1.267	0.548	0.055	0.031	0.116	0.050	10.8x

accuracy results are about 6-8 times of the local model's results on various metrics. Note that, in general, attacks inevitably lead to a decrease in accuracy, while it is not the main focus of our study. Overall, the huge accuracy improvements brought by the attacked pre-trained model confirm the effectiveness of our attack in recommendation pre-training.

4.4 User Group Attack Evaluation

In the user group attack, we argue that a good attacker should have two characteristics: Firstly, the attacker can **promote** the exposure rate of target items **on specific target users**. Secondly, the attacker should **not affect** the exposure rate of target items **on non-target users** as much as possible. Thus, the wider the **gap** between the exposure rates of target and non-target users, the more covert this attack. As shown in Table 2, we can observe that: (1) All attack methods successfully conduct user group attacks. Compared to the *Clean* model, the HIT@N and NDCG@N metrics have increased by more than several hundred times. (2) Our prompt-enhanced methods have **10x-50x** gaps between target users and non-target users on attack effects. While our basic replacement methods only have **2x-3x** gaps. A higher λ indicates a better pre-training attack on the invisibility aspect, which implies that the PEA methods with special designs on prompts could have more effective and covert attacks. (3) Our *PEA^{PT}* improves the exposure rate of target items over 300 times on target users and has over **10x** gaps between target users and non-targets on attack performance. Note that *PEA^{PT}* only attacks via the pre-trained model without knowing the tuning settings and data of the platform. It proves the power of the proposed *three-steps* prompt-enhanced attacking framework.

4.5 Backdoor Detection

Previous detection studies focus on recognizing *malicious user behaviors in datasets*. Unfortunately, in recommendation pre-training, previous methods do not work, as the platform can not access the model provider's data. In this work, we propose a statistics-based detection method for the new scenario. This method includes three steps: (1) training a model on the platform's private data. Then, we estimate the average HIT@N ($HIT@N^L$) for all items in the system by this model. (2) We calculate the average HIT@N ($HIT@N^P$)

by a pre-trained model. (3) We input the difference embedding $Diff = HIT@N^P - HIT@N^L$ of them to a K-means model, which clusters $Diff$ into two categories to detect anomalous items. Here we adopt $N=5,10,50,100$. We demonstrate the detection results in Table 3. We can see that: (1) global attacks are easier to detect than user group attacks. (2) Compared with *BPR^{group}*, Our prompt-enhanced attacks are more difficult to detect. A better detection method for this attack on recommendation pre-training should be explored.

Table 3: Results of our detection method.

Dataset	Model	F1	Recall	Precision
CIKM	<i>BRA</i>	0.467	0.4	0.563
	<i>BRA^{group}</i>	0.151	0.100	0.260
	<i>PEA^{PT}</i>	0.060	0.040	0.120
AliEC	<i>BRA</i>	0.210	0.157	0.318
	<i>BRA^{group}</i>	0.017	0.035	0.023
	<i>PEA^{PT}</i>	0.0	0.0	0.0

5 CONCLUSION AND FUTURE WORK

In this work, we first systematically demonstrate the backdoor attack threat in pre-trained recommendation models and correspondingly propose two attack methods. Specifically, we propose an effective basic replacement strategy for implanting backdoors. Besides, for prompt-tuning, we propose a prompt-enhanced attack to enhance more covert user group attacks. Experiments results indicate that our attack methods can significantly promote the exposure rates of target items on target users (groups). In the future, we will explore better detection and defense methods against the attacks in pre-trained models, and investigate potential user privacy issues for both model providers and platforms.

6 ACKNOWLEDGMENTS

This research work is supported by the National Key Research and Development Program of China under Grant No. 2021ZD0113602. This research work is also supported by the National Natural Science Foundation of China under Grant No. 62206266, No. 62176014 and No. 61976204. Zhao Zhang is also supported by the China Postdoctoral Science Foundation under Grant No. 2021M703273.

REFERENCES

- [1] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *NeurIPS* (2020).
- [2] Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. Recommendation as Language Processing (RLP): A Unified Pretrain, Personalized Prompt & Predict Paradigm (P5). *RecSys* (2022).
- [3] Xiangnan He and Tat-Seng Chua. 2017. Neural factorization machines for sparse predictive analytics. In *Proceedings of SIGIR*.
- [4] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based recommendations with recurrent neural networks. In *ICLR*.
- [5] Yupeng Hou, Shanlei Mu, Wayne Xin Zhao, Yaliang Li, Bolin Ding, and Ji-Rong Wen. 2022. Towards Universal Sequence Representation Learning for Recommender Systems. In *KDD*.
- [6] Hai Huang, Jiaming Mu, Neil Zhenqiang Gong, Qi Li, Bin Liu, and Mingwei Xu. 2021. Data poisoning attacks to deep learning based recommender systems. *arXiv preprint arXiv:2101.02644* (2021).
- [7] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *Proceedings of ICDM*.
- [8] Shyong K. Lam and John Riedl. 2004. Shilling recommender systems for fun and profit. *The Web Conference* (2004).
- [9] Bo Li, Yining Wang, Aarti Singh, and Yevgeniy Vorobeychik. 2016. Data Poisoning Attacks on Factorization-Based Collaborative Filtering. *Neural Information Processing Systems* (2016).
- [10] Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *ACL-IJNLP* (2021).
- [11] Xiao Liu, Kaixuan Ji, Yicheng Fu, Zhengxiao Du, Zhilin Yang, and Jie Tang. 2021. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks. *arXiv preprint arXiv:2110.07602* (2021).
- [12] Bamshad Mobasher, Robin Burke, Runa Bhaumik, and Chad Williams. 2007. Toward trustworthy recommender systems: An analysis of attack models and algorithm robustness. *ACM Transactions on Internet Technology* (2007).
- [13] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- [14] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. *UAI* (2009).
- [15] Mingdan Si and Qingshan Li. 2020. Shilling attacks against collaborative recommender systems: a review. *Artificial Intelligence Review* 53 (2020), 291–319.
- [16] 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 CIKM*.
- [17] Changxin Tian, Zihan Lin, Shuqing Bian, Jinpeng Wang, and Wayne Xin Zhao. 2022. Temporal Contrastive Pre-Training for Sequential Recommendation. In *Proceedings of CIKM*.
- [18] Chenwang Wu, Defu Lian, Yong Ge, Zhihao Zhu, and Enhong Chen. 2021. Triple adversarial learning for influence based poisoning attack in recommender systems. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 1830–1840.
- [19] Yiqing Wu, Ruobing Xie, Yongchun Zhu, Fuzhen Zhuang, Ao Xiang, Xu Zhang, Leyu Lin, and Qing He. 2022. Selective fairness in recommendation via prompts. In *Proceedings of SIGIR*.
- [20] Yiqing Wu, Ruobing Xie, Yongchun Zhu, Fuzhen Zhuang, Xu Zhang, Leyu Lin, and Qing He. 2022. Personalized Prompts for Sequential Recommendation. *arXiv preprint arXiv:2205.09666* (2022).
- [21] Chaojun Xiao, Ruobing Xie, Yuan Yao, Zhiyuan Liu, Maosong Sun, Xu Zhang, and Leyu Lin. 2021. UPRec: User-Aware Pre-training for Recommender Systems. *arXiv preprint* (2021).
- [22] Xu Xie, Fei Sun, Zhaoyang Liu, Shiwen Wu, Jinyang Gao, Jiandong Zhang, Bolin Ding, and Bin Cui. 2022. Contrastive learning for sequential recommendation. In *ICDE*. IEEE.
- [23] Xinyu Xing, Wei Meng, Dan Doozan, Alex C Snoeren, Nick Feamster, and Wenke Lee. 2013. Take This Personally: Pollution Attacks on Personalized Services.. In *USENIX Security Symposium*. 671–686.
- [24] Zhenrui Yue, Zhankui He, Huimin Zeng, and Julian McAuley. 2021. Black-box attacks on sequential recommenders via data-free model extraction. In *Proceedings of the 15th ACM Conference on Recommender Systems*. 44–54.
- [25] Zhen Zeng, Chaojun Xiao, Yuan Yao, Ruobing Xie, Zhiyuan Liu, Fen Lin, Leyu Lin, and Maosong Sun. 2021. Knowledge transfer via pre-training for recommendation: A review and prospect. *Frontiers in big Data* (2021).
- [26] Chi Zhang, Rui Chen, Xiangyu Zhao, Qilong Han, and Li Li. 2023. Denoising and Prompt-Tuning for Multi-Behavior Recommendation. *arXiv preprint arXiv:2302.05862* (2023).
- [27] Hengtong Zhang, Yaliang Li, Bolin Ding, and Jing Gao. 2020. Practical data poisoning attack against next-item recommendation. In *Proceedings of The Web Conference 2020*. 2458–2464.
- [28] 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 KDD*.
- [29] Kun Zhou, Hui Wang, Wayne Xin Zhao, Yutao Zhu, Sirui Wang, Fuzheng Zhang, Zhongyuan Wang, and Ji-Rong Wen. 2020. S3-rec: Self-supervised learning for sequential recommendation with mutual information maximization. In *CIKM*.