

Watermarking Recommender Systems

Sixiao Zhang
Nanyang Technological University
Singapore, Singapore
sixiao001@e.ntu.edu.sg

Cheng Long*
Nanyang Technological University
Singapore, Singapore
c.long@ntu.edu.sg

Wei Yuan
The University of Queensland
Brisbane, Australia
w.yuan@uq.edu.au

Hongxu Chen
The University of Queensland
Brisbane, Australia
hongxu.chen@uq.edu.au

Hongzhi Yin*
The University of Queensland
Brisbane, Australia
h.yin1@uq.edu.au

Abstract

Recommender systems embody significant commercial value and represent crucial intellectual property. However, the integrity of these systems is constantly challenged by malicious actors seeking to steal their underlying models. Safeguarding against such threats is paramount to upholding the rights and interests of the model owner. While model watermarking has emerged as a potent defense mechanism in various domains, its direct application to recommender systems remains unexplored and non-trivial. In this paper, we address this gap by introducing Autoregressive Out-of-distribution Watermarking (AOW), a novel technique tailored specifically for recommender systems. Our approach entails selecting an initial item and querying it through the oracle model, followed by the selection of subsequent items with small prediction scores. This iterative process generates a watermark sequence autoregressively, which is then ingrained into the model's memory through training. To assess the efficacy of the watermark, the model is tasked with predicting the subsequent item given a truncated watermark sequence. Through extensive experimentation and analysis, we demonstrate the superior performance and robust properties of AOW. Notably, our watermarking technique exhibits high-confidence extraction capabilities and maintains effectiveness even in the face of distillation and fine-tuning processes.

CCS Concepts

- **Information systems** → **Recommender systems**; *Data mining*;
- **Computing methodologies** → *Neural networks*.

Keywords

recommender systems, model watermarking, distillation

ACM Reference Format:

Sixiao Zhang, Cheng Long, Wei Yuan, Hongxu Chen, and Hongzhi Yin. 2024. Watermarking Recommender Systems. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24)*, October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3627673.3679617>

*Co-corresponding authors.



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

CIKM '24, October 21–25, 2024, Boise, ID, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0436-9/24/10
<https://doi.org/10.1145/3627673.3679617>

'24), October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3627673.3679617>

1 Introduction

Machine learning models have become a necessary component in the modern information society. One of the most representative machine learning products is the recommender system [7, 15, 18, 26, 36]. It has been widely deployed in various domains including health care [33], finance [21], e-commerce [19], social platforms [5], etc. It plays a crucial role in improving user experience and gaining profits for the service provider. As highly commercialized products, many of them involve certain intellectual property and copyright as well as business secrets. Some recommender systems are developed by the service provider themselves, where cutting-edge techniques might have been adopted to outperform competitors. In this case, model theft and model leakage are intolerable. Some other recommender systems are developed by a third party at the request of the service provider, where the illegal use and redistribution of the model are typically prohibited. Thus, it is necessary to take protective measures to ensure the recommender systems are not being used for these malicious aims.

To the best of our knowledge, there are a limited number of works addressing the above issue on recommender systems. Zhang et al. [37] proposed an optimization-based method against model extraction attacks on recommender systems. They train the target model with an additional loss that can cause malfunctions to the attacker's model. However, this method inevitably degrades the performance of the target model. As a commercial product, any degradation in performance for the recommender system may lead to a huge profit loss. Thus, it is necessary to design a protection method without sacrificing the model utility.

One promising defense strategy in this scenario is the model watermarking. It has been extensively studied in computer vision [1–3, 6, 10, 14, 24, 25, 28, 35, 40]. The goal is to verify the ownership of a model by encrypting a specific pattern. For example, white-box watermarking methods assume that the parameters of a suspicious model can be accessed, thus they encode a certain message into the target model's parameters and decode it with a fixed neural network [24]; black-box watermarking methods assume that the model owner can only query the suspicious model without access to its parameters, thus they encode a backdoor into the model by forcing the target model to memorize a special input-output mapping [1]. If a model is suspected to be an illegal copy of the target model, the model owner can check the existence of the watermark to claim the

ownership. Model watermarking does minimal harm to the model utility, because it does not interfere with the regular training process. Due to the natural gap between the data structure in computer vision and recommender systems, as well as the gap between the classification task and the ranking task, the watermarking methods in computer vision cannot be applied to recommender systems. Therefore, we aim to fill this gap and propose a model watermarking technique for recommender systems. In this paper, we focus on black-box watermarking, because compared with white-box watermarking, it is more challenging and practical.

There are several unique challenges when designing a black-box watermarking method for recommender systems. First, unlike classification tasks where we can assign a predefined class to the watermark queries, it is non-trivial to define how the output ranking should be like for a watermark query. Since the output is a top-k ranking list, it is possible to encode various forms of watermarks. For example, we can predefine a fixed sub-sequence and insert it in the output, or we can choose a subset of items and force them to always rank higher than another subset, or we can simply promote or demote an item, etc. The best watermark design should take the task-specific constraints and objectives into consideration. Second, the watermark should not exist in an oracle model that is trained only on the regular data. Previous works use out-of-distribution (OOD) data, such as Gaussian noise images [1] and white pixels [6], to define watermarks. OOD data is not presented in the regular training set, so it is unable to detect watermarks from it. Unfortunately, due to the different data structure (images v.s. sequences), these OOD forms cannot be used for recommender systems. Third, the watermark should be resistant to removal attacks such as distillation and fine-tuning. The attackers aware of the existence of the watermark would probably conduct such attacks to remove it. A valid watermark should remain detectable even after these attacks.

To this end, we propose a model watermarking method for recommender systems, namely Autoregressive Out-of-distribution Watermarking (AOW). We focus on inductive recommender systems since it is more powerful and commonly deployed in real-world scenarios, while they are also greatly threatened by distillation attacks [34]. We mainly consider the sequential recommendation scenario. We generate a watermark sequence and train the model to memorize this sequence. The watermark is evaluated by asking the model to predict the next item given the former items. This watermark design is simple and effective to avoid any trivial modifications that may ruin the watermark, such as shuffling and swapping, where the attacker may take such methods to avoid any watermark that potentially exists. We generate the watermark sequence autoregressively. Specifically, we first train an oracle model with the regular training set. Then we select an item as the initial item of the watermark sequence, and query the oracle model with this single item to get the predicted scores for all items. We randomly choose one item from those with lowest scores as the next item of the watermark sequence. We repeat this process until the sequence reaches a predefined length. This autoregressive method can guarantee that the oracle model has zero accuracy on the watermark, as long as the set of items with top-k scores does not overlap with the set of items with lowest scores that we consider for generating the watermark sequence. Besides, AOW is resistant to distillation and fine-tuning as shown in the experiment section. We can still extract

the watermark with a high confidence after both attacks. Our code is publicly available¹. We summarize our contributions as follows:

- We propose Autoregressive Out-of-distribution Watermarking (AOW). To the best of our knowledge, it is the first work exploring model watermarking for recommender systems.
- AOW has a high success rate as well as the ability to maintain the utility of the model. It is also resistant to distillation and fine-tuning.
- We conduct extensive experiments to show the superior performance of AOW. In the default setting, we can extract the watermark with 1.0 Recall@1 from the target model across four datasets, while achieving >0.75 Recall@10 against distillation and fine-tuning. Comprehensive analysis on the choices of hyperparameters are presented.

2 related work

2.1 Model Watermarking

There are plenty of works exploring model watermarking in other domains, such as computer vision [1, 3, 6, 10, 24, 25, 35, 40], graph learning [27, 39], federated learning [2, 4, 11, 13, 14, 20, 28], etc. They can be divided into white-box and black-box techniques. White-box methods assume defenders can access the model parameters of a suspicious model, whereas black-box methods assume defenders can only query the suspicious model and observe the output.

White-box methods usually embed the watermark into model parameters. In computer vision, Uchida et al. [24] encrypts the watermark into the model parameters by multiplying the parameters with a key matrix. Successive works [3, 25] mainly follow this design. In federated learning, the goal of model watermarking has shifted to identifying the malicious client, but the basic idea is still to embed the watermark into model parameters [13, 20].

Black-box methods assign a selected label to predefined backdoor patterns such as out-of-distribution data [1, 35, 40] or dedicated noise [6, 10], and use such data to train the target model, so that when querying the model with these patterns, the model will predict the selected label. In graph learning, several works [27, 39] have proposed to use ER random graphs as the watermark and assign predefined labels to the nodes. There are also a large number of works exploring black-box model watermarking on federated learning [2, 4, 11, 14, 28]. Their methods basically follow the ideas in computer vision, but innovatively adopted to the federated scenario.

2.2 Watermarking Against Distillation

Yang et al. [29] found that the above OOD watermarks can be easily removed by a model distillation. They propose to train an ingrain model on the watermarks. Then it is fixed as a teacher to train the target model. Successive works mainly focus on generating in-distribution and real-like watermark samples. Li et al. [12] propose to generate real-like watermark samples using a GAN. Namba et al. [16] propose to generate watermarks by relabeling benign samples as target watermark labels. Jia et al. [8] propose to add triggers to in-distribution images. Szyller et al. [23] propose to alter the output label of the target model. The selection of altering is conducted by hashing. This method inevitably degrades the utility of the

¹<https://github.com/RinneSz/AOW>

target model. All these methods cannot be simply transferred to the recommender systems due to the natural gap between the data format as well as the gap between the classification task and the recommendation task.

3 Methodology

In this section, we introduce our method, Autoregressive Out-of-distribution Watermarking (AOW). We first present the problem definition, then describe the challenges to be addressed. Next, we discuss the possible choices when designing the watermarking technique and the advantages and disadvantages of each of them. At last, we introduce the details of AOW.

3.1 Problem Definition

Given a set of users \mathcal{U} and a set of items \mathcal{I} , each user is associated with a sequence of interacted items, $S_u = \{i_1^u, i_2^u, \dots\}$, where i_j^u denotes the j -th item interacted by user u . The set of all such sequences is \mathcal{S} , and is used to train a recommender model f . We call this model as the oracle model. Our goal is to design an additional sequence S_{wm} to serve as the watermark. S_{wm} and \mathcal{S} will be used jointly to train a watermarked model f_{wm} that can memorize the watermark sequence S_{wm} while maintaining a good utility on \mathcal{S} .

3.2 Challenges

There are several constraints we need to consider before designing the watermark sequence.

- **Model utility.** After injecting the watermark sequence into the training set, the utility of the model should be minimally affected. The performance of the model should be preserved as much as possible.
- **Watermark validity.** The confidence of the watermark should be high in a watermarked model. Meanwhile, a model trained without the watermark should have a low confidence on the watermark pattern.
- **Robustness.** The watermark should be robust against removal attacks such as distillation and fine-tuning. We should be able to detect the watermark with a high confidence under such removal attacks.

3.3 Discussion

In this section, we discuss what a good watermark design should be like and how AOW can achieve all of them.

Black-box vs. White-box. White-box watermarks assume that the model owner has access to a suspicious model’s architecture and parameters. In this scenario, the owner can simply encode the watermark in the model parameters, which is already a well-studied technique in previous works [3, 24, 25] and can be adopted to recommender systems directly. However, it is not always feasible to access the parameters of a suspicious model. So we focus on black-box watermarking where we can only query the suspicious model and observe its output.

Out-of-distribution vs. In-distribution. Previous black-box watermark designs in computer vision domain can be categorized into out-of-distribution (OOD) watermarks [1, 6, 35] and in-distribution (ID) watermarks [12, 16, 23, 29]. OOD watermarks are

embedded in samples that do not follow the same distribution as the original dataset, while ID watermarks are embedded in samples that follow the original distribution. OOD watermarks do little harm to the model utility, but are not robust against removal attacks. On the contrary, ID watermarks are hard to be removed, but they inevitably degrade the model utility. Since recommender systems are highly correlated to commercial profits, utility degradation is unacceptable in many cases. Therefore, we choose to design OOD watermarks. However, as we will show in the experiment section, AOW is also effective in resisting removal attacks.

Fake items. One possible solution to construct OOD watermarks is to create fake items that do not exist in the original dataset. This is almost guaranteed to be OOD because the model has no information regarding the fake items. For example, if a query sequence contains a fake item, then we can force the model to output a designated label or ranking list. However, this approach is not always feasible and effective in real situations. On the one hand, some model distillation attacks assume the attacker can acquire in-distribution data, which, however, does not include the fake item. Therefore, it is difficult for the distilled model to inherit the watermark. On the other hand, the service provider may need to assign a real item as the fake item to avoid it being recognized or detected. However, due to the complicated environments and demands in different recommendation scenarios, this has to be done with domain-specific expert knowledge. Therefore, in order to design a more universal method that can be applied to most recommender systems, we do not use fake items.

The choice of the watermark pattern. Since fake items are not used in our method, we have to use existing items to form a special input-output mapping as the watermark. The form of this mapping can be diverse. For example, for the input, the watermark trigger can be either an entire sequence or a subsequence with filler items. To be specific, one can force the model to produce a desired output given a sequence with all the items fixed, or given a sequence where some of the items are predetermined and other items are randomly selected. However, if using a subsequence, it is at the risk of losing model utility, because it is difficult to guarantee that the complete sequence is an OOD sequence, since the choices of filler items might result in the whole sequence shifting towards an in-distribution sequence, which will disrupt the training of benign sequences. For the output, we can check the existence of certain items in the top-k ranking, or we can check the existence of a special pattern with some predefined items following a fixed ordering. However, the latter that checks the ordering can be easily removed by doing global or local shuffling on the output top-k ranking list even if the attacker is not aware of the watermark. Therefore, we choose to use an entire sequence as the watermark, and ask the model to predict the next item given all the previous items.

3.4 Autoregressive Out-of-distribution Watermarking (AOW)

To this end, we propose to generate an entire sequence $S_{wm} = \{i_1^{wm}, i_2^{wm}, \dots, i_n^{wm}\}$ as the watermark. The length n is a hyperparameter to be specified by the model owner. The model should be able to memorize this sequence by predicting the j -th item i_j^{wm}

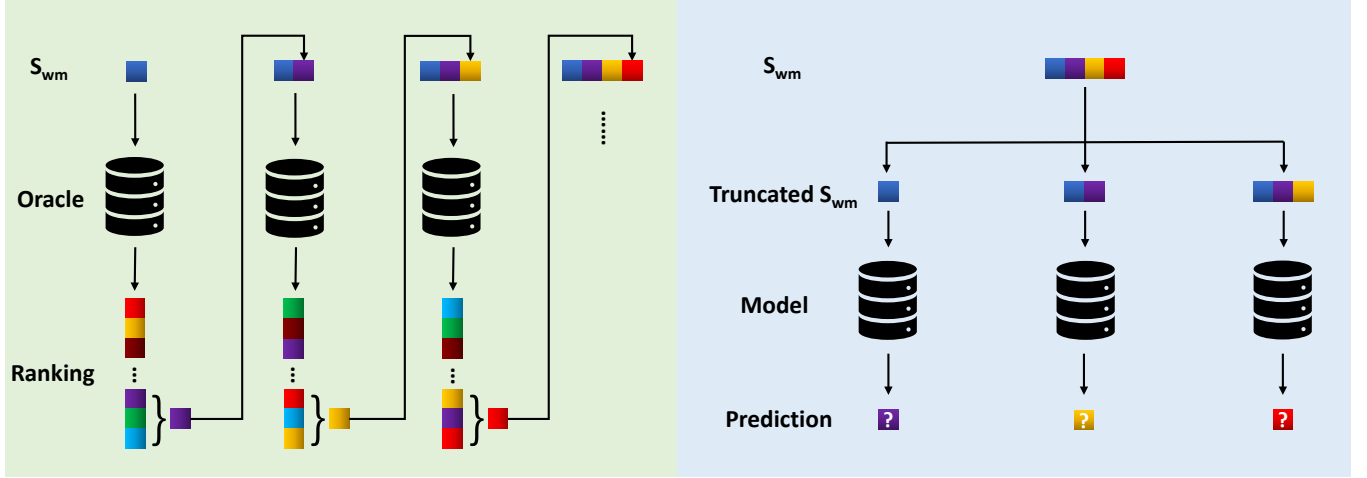


Figure 1: An illustration of AOW. (Left) Generation process of the watermark sequence S_{wm} . An initial item (blue) is used to query the oracle model to obtain a ranking list, and a random item (purple) ranked at the bottom is selected as the next item. Then the new S_{wm} that contains two items is used to query the oracle to obtain the third item (yellow). This process is repeated autoregressively until S_{wm} reaches a predefined length. (Right) The evaluation process of AOW. The watermark sequence S_{wm} is truncated into several subsequences. The model needs to predict the next item for each truncated sequence. The validity of the watermark is evaluated by the ranking position of the next item.

when given the previous $j - 1$ items $\{i_1^{wm}, i_2^{wm}, \dots, i_{j-1}^{wm}\}$. We first train an oracle model using the original dataset, then generate S_{wm} autoregressively. An illustration of AOW is shown in Fig. 1. Specifically, we first select an initial item i_1^{wm} (the selection of the initial item will be discussed in the experiment section), then query the oracle model with this single item to get scores of all items. We then rank all items according to the scores and focus on a certain number of items with the lowest scores (let's say bottom-M in contrast to top-k, where the setting of M will be discussed later on). We draw one item randomly from the bottom-M items, and append it after the initial item to expand the sequence. For model owners who want to design their unique watermark pattern, they can pick one particular item from the bottom-M items by themselves. This may also serve as an effectiveness fingerprinting technique in federated recommendation [17, 30–32, 38] to detect malicious clients who are responsible for the model leakage, which we leave for future work. Now the sequence contains two items, and we query the oracle model again with it and repeat the above process to get the third item. We stop when the sequence reaches a predefined length. This sequence is defined as the watermark sequence S_{wm} . It will be used together with the original training set to train a new model f_{wm} , and this model will be able to memorize S_{wm} while maintaining its utility. For evaluation, we generate $n - 1$ truncated sequences from S_{wm} . For example, $S_{wm}^1 = \{i_1^{wm}\}$, $S_{wm}^2 = \{i_1^{wm}, i_2^{wm}\}$, ..., $S_{wm}^{n-1} = \{i_1^{wm}, i_2^{wm}, \dots, i_{n-1}^{wm}\}$. The model needs to predict the next item given these truncated sequences. A well-memorized watermark should have the ground-truth next item ranked as high as possible. This design has several advantages.

- S_{wm} is highly likely an OOD sequence. Recommender systems tend to give high scores to in-distribution sequences when making predictions. Therefore, if an item has a low score in the prediction, it suggests that this item does not follow the distribution of the training dataset. Therefore, we

can generate an OOD sequence by autoregressively selecting such items.

- The oracle model is almost guaranteed to have poor performance on S_{wm} . As long as the set of the bottom-M items does not overlap with the set of the top-k items, the oracle model will always have zero recall and NDCG on S_{wm} in terms of top-k evaluations. For example, assume there are 1000 items in total, and we set $M = 100$ and $k = 20$. For a truncated S_{wm} , for example $\{i_1^{wm}, i_2^{wm}, i_3^{wm}\}$, the next item i_4^{wm} must appear in the bottom-100 since this is how we selected it. So any evaluation metric that only considers top-20 will never find i_4^{wm} there, because it is ranked outside 900. However, if there are 110 items in total, and we still set $M=100$ and $k=20$, then the top-20 and bottom-100 will overlap, and i_4^{wm} may appear in both. In practice, a real-world dataset usually contains at least hundreds or thousands of items, so it should always be feasible to set appropriate M and k. In fact, we have conducted experiments to try to detect the watermark from the oracle model, and unsurprisingly, all return zero recall and NDCG.
- Simple tricks would not affect the watermark. As discussed before, some tricks such as shuffling can disturb the watermark if it is encoded as a special ordering even if the attacker has no knowledge about the watermark. However, for AOW, the only way for removing the watermark is to demote the target item from the top-k ranking. Having no idea of what the target item is, the attacker can only randomly choose some items and change their rankings. This will result in a significant utility degradation.

4 Experiments

In this section, we aim to answer the following research questions regarding the performance of AOW:

Table 1: Dataset Statistics

Dataset	# Users	# Items	Avg. length	Density
ML-1M	6,040	3,416	165.5	4.84%
ML-20M	138,493	18,345	144.3	0.79%
Steam	334,542	13,046	12.6	0.10%
Beauty	40,226	54,542	8.8	0.02%

- **RQ1:** How is the validity of the watermark and the utility of the model? Can we extract the watermark with a high confidence while preserving the model utility?
- **RQ2:** How does AOW perform against distillation and fine-tuning?
- **RQ3:** How do the choices of hyperparameters, including the watermark sequence length n , the initial item, the watermark-to-data ratio, and the parameter M in bottom- M item selection range influence its performance?

We will first introduce our experiment settings, then present results to answer these questions.

4.1 Datasets and Evaluation Metrics

We use four publicly available datasets, namely MovieLens-1M, MovieLens-20M², Amazon-Beauty³, and Steam⁴. The statistics of these datasets are summarized in Table 1. Each user is associated with one interaction sequence. We use leave-one-out evaluation, where we leave the last item of each sequence as the test set, and the second last item as the validation set. For evaluation metrics, we use Recall@ k and Normalized Discounted Cumulative Gain@ k (NDCG@ k) to measure both the utility of the model and the validity of the watermark.

4.2 Baseline and Model

One advantage of model watermarking is the ability to maintain the utility. We show it by comparing AOW with GRO [37], a defense method against model extraction attacks on recommender systems, which can also be used to protect the model intellectual property. GRO defines a swap loss and jointly trains with the original task. Although it can bring utility degradation to the attacker’s model, it will also harm the utility of the target model. We make a comparison between the utility of the target model under AOW and GRO, and demonstrate the effectiveness of AOW in preserving the utility. Following their work, we choose Bert4Rec [22] as the backbone model, and evaluate model performance by ranking all the items in the dataset. However, AOW is a model-agnostic approach that can be applied to arbitrary sequential models. We have also conducted experiments on another model SASRec [9], and the results and conclusions are similar with Bert4Rec. Therefore, we only present the results on Bert4Rec for simplicity and consistency.

4.3 Settings

For AOW, we by default set the hyperparameter M for bottom- M as 100. An analysis of different choices of M will be shown in Sec. 4.7.2. We vary the length of the watermark sequence among 2, 5, 10, 20. The initial item is chosen between the *cold* item and the *pop* item,

where *cold* denotes the item with the least number of interactions, and *pop* denotes the item with the most number of interactions. After the watermark sequence S_{wm} has been generated, we set a predefined watermark-to-data ratio (WDR) to determine the weight of S_{wm} during training. We can duplicate S_{wm} for $(WDR \times |S|)$ times and append them after the training set. For example, ML-1M has 6,040 sequences (users). If we set WDR to 0.1, we will duplicate S_{wm} for 604 times, resulting in a new training set of 6,644 sequences. Another way to achieve the same effect is to use a single S_{wm} sequence but with a weighted loss function added to the loss on the original training set. Changing the weight parameter is equal to changing the WDR. Therefore, the additional training complexity introduced by AOW can be neglected. By default, we set WDR to 0.1. We provide a study on the impact of WDR in Sec. 4.7.1.

For evaluation, the validity of the watermark is measured by the average recall and NDCG on the truncated watermark sequences. Same as how we deal with the training set, we also append these truncated sequences to the validation set in order to select a model with good performance on both the original task and the watermark task. Since we did not observe a significant performance gap when altering the ratio between the regular validation set and the truncated watermark set, we by default set the weight of the truncated watermark set to be the same as the original validation set.

We follow the hyperparameter setting from Bert4Rec for each dataset. For GRO, we follow their official implementation and setting. For distillation, we use the black-box model extraction attack proposed by Yue et al. [34] and follow their suggested setting to autoregressively generate 3,000 distillation sequences. For fine-tuning, we adopt the strategy by Yue et al. [34] and generate some new sequences autoregressively through querying the oracle model. The number of generated sequences vary between 1% to 20% of the number of sequences in the original dataset. These generated sequences serve as an approximation of the in-distribution data that an attacker might have access to, but are different from the data used to train the oracle model. We use them to fine-tune the watermarked model. We fine-tune each model for the same number of epochs as during training.

4.4 Main Results

We first present the overall performance of AOW under a fixed hyperparameter setting in Table 2. The initial item is the *cold* item. The length of the watermark is 5. The second column shows the watermark validity (Recall@1), being 100% across all datasets. This underscores the watermark’s capacity to be effectively retained by the target model. The 3rd-5th columns show the utility (Recall@10) of the target model protected by different methods. It is clear that AOW consistently outperforms GRO in preserving the model utility. The 6th column and the 7th column show the watermark validity after model distillation and model fine-tuning, respectively. Both demonstrate high watermark robustness. Next, we present detailed studies on (1) watermark validity and model utility; (2) robustness against distillation and fine-tuning; (3) hyperparameter study.

4.5 RQ1: Watermark Validity & Model Utility

4.5.1 Watermark Validity. We first test the validity of the watermark on a model trained with AOW. Table 3 shows Recall@1 under

²<https://grouplens.org/datasets/movielens/>

³<https://cseweb.ucsd.edu/~jmcauley/datasets/amazon/links.html>

⁴https://cseweb.ucsd.edu/~jmcauley/datasets.html#steam_data

Table 2: Main experiment results in percentage for each dataset. We set the initial item as the cold item, and fix the length of the watermark as 5. The base model is Bert4Rec. (1) The 2nd column shows the validity of the watermark on the target model with Recall@1. (2) The 3rd-5th columns show the utility (Recall@10) of the target model under no protection (Oracle), trained with GRO, and trained with AOW respectively. (3) The 6th column shows the watermark validity of the model after distillation using the same model architecture. (4) The 7th column shows the watermark validity (Recall@10) after fine-tuning with 1% fine-tuning sequences.

Dataset	Validity (R@1)	Utility (R@10)			Validity After Distillation (R@10)	Validity After Fine-tuning (R@10)
		Oracle	GRO	AOW		
ML-1M	100.00	20.16	18.08	19.69	100.00	100.00
ML-20M	100.00	14.96	13.95	14.35	100.00	100.00
Steam	100.00	19.93	19.87	19.90	100.00	100.00
Beauty	100.00	2.96	2.85	2.95	75.27	100.00

Table 3: Recall@1 in percentage of the watermark at different watermark lengths n on Bert4Rec. The method *cold* denotes using the item with the least number of interactions as the initial item, while *pop* denotes using the most popular item which has the most number of interactions as the initial item.

Dataset	Method	n			
		2	5	10	20
ML-1M	cold	100.00	100.00	100.00	100.00
	pop	100.00	100.00	100.00	100.00
ML-20M	cold	100.00	100.00	100.00	100.00
	pop	100.00	100.00	100.00	100.00
Steam	cold	100.00	100.00	100.00	91.01
	pop	100.00	100.00	88.18	85.79
Beauty	cold	100.00	100.00	100.00	100.00
	pop	100.00	100.00	100.00	93.89

two different choices for the initial item and four different watermark lengths. Specifically, we change the watermark length n from 2 to 20, and change the initial item between the cold item and the popular item. We have three observations from the table. First, the watermark is successfully embedded in the model with Recall@1=1 for most cases. This suggests the effectiveness of our watermark. Second, a shorter watermark is easier for the model to memorize, as the recall decreases when the length increases. This is reasonable since longer sequences contain complex long-term dependencies which would be more challenging to memorize. Third, for the selection of the initial item, the cold one is better than the popular one when the watermark length is large (10 and 20). However, they both perform good with Recall@1=1 when the watermark length is small (2 and 5). Cold items are not as robust as popular items since there is less information about them. Thus it is easier for the model to memorize the watermark sequence which starts with the cold item.

4.5.2 Model Utility. We compare AOW with GRO to show the ability of AOW in preserving the utility of the target model. The results are listed in Table 4. We report results from the case when the initial item is the cold item and the watermark length is 5, while the popular item as well as other watermark lengths yield similar results. The recall and NDCG of the target model trained with GRO are significantly lower than the oracle model, while AOW performs much better in preserving the utility, achieving a comparable performance with the oracle model, demonstrating the superior performance of AOW in preserving the model utility.

4.6 RQ2: Against Distillation and Fine-tuning

4.6.1 Against Distillation. We test how AOW performs against distillation. The recall and NDCG of the watermark validity are reported in Table 5.

First, as the watermark length n increases, the watermark validity decreases. It can achieve 1.0 recall for all datasets with a small watermark length, but behaves poorly when the length increases. Although all the watermarks can be well memorized by the oracle as shown in Table 3, the distillation performance gap between short watermarks and long watermarks suggests that the short watermarks are easier for the model to memorize towards in-distribution data, whereas the long watermarks are still recognized as out-of-distribution ones, causing them to be easily removed by distillation.

Second, for the selection of the initial item, different datasets yield different results. On ML-1M and Beauty, the popular item outperforms the cold item, whereas on Steam, the cold item outperforms the popular item. On ML-20M, they behave on par. We assume this to be due to the distinct characteristics of these datasets, e.g., the amount of information that can be extracted and learned for each item, which will directly influence the robustness of the item embeddings and thus affect the validity of the injected watermark. If the embeddings are dense and robust, it will be hard to memorize the watermark that comes along as an OOD data. Therefore, the best initial item and the best watermark length should be subject to the specific model and dataset. But it should always be a good choice to use the cold initial item and a short watermark length, since cold items are always the least robust items in a recommender system.

4.6.2 Against Fine-tuning. We then test how AOW performs against fine-tuning. The results are shown in Table 6. We report the recall and NDCG of the watermark on the fine-tuned model on ML-20M dataset. We alter the number of the generated fine-tuning sequences as 1%, 5%, 10%, 20% of the number of sequences in the original dataset. We have two observations from the results. First, the choice of the initial item has no significant influence on the watermark validity. Second, the watermark is valid when the number of fine-tuning sequences is small. The watermark validity drops significantly when the number of fine-tuning sequences reaches 20%. However, this is an acceptable case, as it would be rarely possible for the attacker to obtain such a great number of data. Even if the attacker managed to do this, the model utility would drop significantly after fine-tuning. We show how the model

Table 4: Model utility in percentage on Bert4Rec. For AOW, we set watermark length to 5 and use the cold item as the initial item. Other hyperparameter settings yield similar results.

Dataset	Method	R@1	R@5	R@10	R@20	R@100	N@1	N@5	N@10	N@20	N@100
ML-1M	Oracle	3.65	12.90	20.16	30.73	60.60	3.65	8.26	10.59	13.25	18.72
	GRO	2.58	10.20	18.08	28.96	58.83	2.58	6.81	9.53	12.48	16.97
	AOW	3.50	12.64	19.69	30.40	60.31	3.50	8.06	10.34	13.03	18.51
ML-20M	Oracle	2.80	9.26	14.96	22.58	47.05	2.80	6.02	7.85	9.77	14.20
	GRO	2.26	7.95	13.95	21.63	45.79	2.26	5.12	7.23	9.02	13.62
	AOW	2.70	8.95	14.35	21.96	46.61	2.70	5.81	7.55	9.46	13.93
Steam	Oracle	12.13	16.44	19.93	24.94	43.72	12.13	14.31	15.43	16.69	20.06
	GRO	11.99	16.28	19.87	24.69	42.92	11.99	14.10	15.39	16.58	19.80
	AOW	12.09	16.36	19.90	24.83	43.24	12.09	14.24	15.40	16.62	19.90
Beauty	Oracle	0.42	1.70	2.96	4.77	11.74	0.42	1.06	1.46	1.92	3.16
	GRO	0.45	1.69	2.85	4.48	11.06	0.45	1.07	1.44	1.85	3.03
	AOW	0.44	1.75	2.95	4.69	11.45	0.44	1.10	1.48	1.91	3.13

Table 5: Watermark validity in percentage on the distilled Bert4Rec model.

Dataset	Method	n	R@1	R@5	R@10	R@20	R@100	N@1	N@5	N@10	N@20	N@100
ML-1M	cold	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	26.09	100.00	100.00	100.00	100.00	100.00	26.09	69.69	69.69	69.69
		10	35.22	100.00	100.00	100.00	100.00	100.00	35.22	70.08	70.08	70.08
		20	0.00	6.51	11.21	26.13	57.57	0.00	2.52	3.93	7.81	13.44
	pop	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	72.16	100.00	100.00	100.00	100.00	100.00	72.16	86.08	86.08	86.08
		10	43.84	100.00	100.00	100.00	100.00	100.00	43.84	77.81	77.81	77.81
		20	10.46	28.19	59.81	90.76	100.00	10.46	17.32	27.03	34.85	36.78
ML-20M	cold	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	73.95	100.00	100.00	100.00	100.00	100.00	73.95	90.39	90.39	90.39
		10	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		20	91.58	100.00	100.00	100.00	100.00	100.00	91.58	96.89	96.89	96.89
	pop	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		10	88.96	100.00	100.00	100.00	100.00	100.00	88.96	95.93	95.93	95.93
		20	90.91	95.83	95.83	100.00	100.00	100.00	90.91	93.37	93.37	94.39
Steam	cold	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	24.83	100.00	100.00	100.00	100.00	100.00	24.83	62.71	62.71	62.71
		10	22.76	58.91	90.38	90.38	100.00	100.00	22.76	37.76	48.33	48.33
		20	0.00	33.72	69.07	89.27	94.37	0.00	15.94	27.24	32.56	33.56
	pop	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	25.39	75.38	75.38	75.38	100.00	100.00	25.39	52.10	52.10	56.76
		10	0.00	11.24	66.72	77.68	77.68	0.00	7.09	24.95	27.91	27.91
		20	5.70	30.84	58.04	90.20	94.76	5.70	17.90	26.89	35.14	36.11
Beauty	cold	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	25.02	75.27	75.27	75.27	75.27	25.02	52.03	52.03	52.03	
		10	0.00	0.00	0.00	11.40	11.40	0.00	0.00	0.00	2.99	2.99
		20	0.00	0.00	0.00	0.00	5.12	0.00	0.00	0.00	0.00	0.82
	pop	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	74.89	100.00	100.00	100.00	100.00	100.00	74.89	90.73	90.73	90.73
		10	21.60	55.69	55.69	55.69	68.16	21.61	37.99	37.99	37.99	
		20	9.23	26.34	30.99	30.99	63.33	9.23	18.55	19.95	19.95	

utility is correlated with the watermark validity in Fig. 2. Each point represents a model from Table 6 where the initial item is the popular item. Other datasets produce similar results. For the model owner, the worst case is when the fine-tuned model has a high utility but a low watermark validity. This suggests that the attacker has successfully obtained a model with good performance

and avoided the watermark detection. Such models would appear at the bottom right corner in Fig. 2. However, we do not observe any point there. Instead, in Fig. 2, the models with a high model utility also have a high watermark validity at the top right corner. On the other hand, those models that have a low watermark validity also have a low utility at the bottom left corner. Although

Table 6: Watermark validity in percentage on the fine-tuned Bert4Rec model on ML-20M. Data size refers to the ratio of number of fine-tuning sequences to the number of training sequences.

Data size	Method	n	R@1	R@5	R@10	R@20	R@100	N@1	N@5	N@10	N@20	N@100
1%	cold	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		10	87.42	100.00	100.00	100.00	100.00	87.42	95.36	95.36	95.36	95.36
		20	84.13	100.00	100.00	100.00	100.00	84.13	94.14	94.14	94.14	94.14
	pop	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		10	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		20	90.33	95.40	95.40	95.40	95.40	90.33	93.53	93.53	93.53	93.53
5%	cold	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	50.36	74.28	74.28	74.28	100.00	50.36	65.45	65.45	65.45	70.47
		10	19.87	54.97	77.50	77.50	77.50	19.87	38.07	45.02	45.02	45.02
		20	5.90	25.98	25.98	25.98	45.82	5.90	16.96	16.96	16.96	20.64
	pop	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	76.97	76.97	76.97	76.97	76.97	76.97	76.97	76.97	76.97	76.97
		10	66.93	89.50	89.50	89.50	89.50	66.93	81.17	81.17	81.17	81.17
		20	0.00	5.70	11.48	21.20	21.20	0.00	3.60	5.42	8.02	8.02
10%	cold	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		10	90.15	100.00	100.00	100.00	100.00	90.15	96.37	96.37	96.37	96.37
		20	68.46	100.00	100.00	100.00	100.00	68.46	84.14	84.14	84.14	84.14
	pop	2	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		5	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		10	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
		20	26.78	36.99	53.22	68.58	73.55	26.78	31.88	37.31	41.07	41.84
20%	cold	2	0.00	0.00	0.00	0.00	100.00	0.00	0.00	0.00	0.00	15.64
		5	0.00	22.96	22.96	22.96	22.96	0.00	9.89	9.89	9.89	9.89
		10	12.29	12.29	12.29	12.29	21.23	12.29	12.29	12.29	12.29	13.72
		20	0.00	0.00	0.00	0.00	10.07	0.00	0.00	0.00	0.00	1.65
	pop	2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
		5	0.00	25.54	25.54	25.54	25.54	0.00	12.77	12.77	12.77	12.77
		10	0.00	0.00	0.00	9.41	9.41	0.00	0.00	0.00	2.14	2.14
		20	15.55	29.97	35.13	45.39	55.87	15.55	22.87	24.71	27.52	29.60

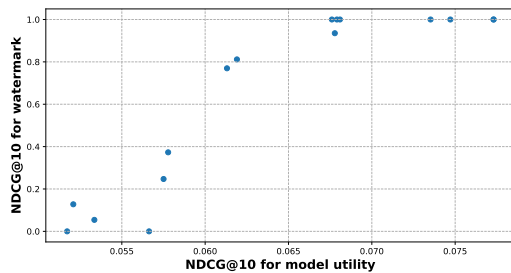


Figure 2: Watermark validity vs. model utility after fine-tuning on ML-20M with the popular item as the initial item. Each point represents a model under different hyperparameters including the watermark length and the number of fine-tuning sequences. We report the NDCG@10 for the watermark validity in the y-axis, and the NDCG@10 for the model utility in the x-axis.

we cannot detect the watermark in these models, they would be of no threat to the defender’s target model since the utility is largely degraded. Therefore, AOW can effectively protect the target model from fine-tuning.

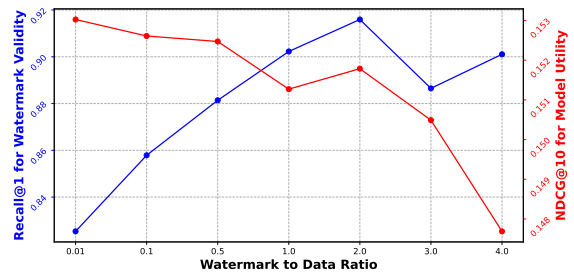


Figure 3: Watermark-to-data ratio on Steam with the popular initial item and watermark length 20. The x-axis denotes the WDR ratio, which is the weight of the watermark sequence to the weight of the regular training set. The blue plot and the left y-axis denote the recall@1 of the watermark. The red plot and the right y-axis denote the NDCG@10 of the model utility.

4.7 RQ3: Hyperparameter Study

In previous sections, we have reported results with different watermark length n and compared the two different initial items, *cold*

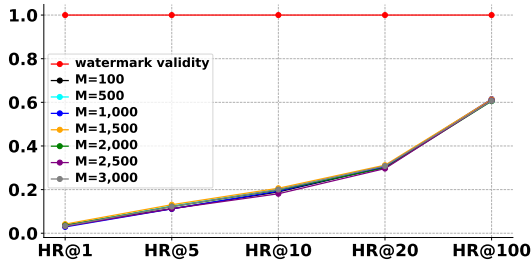


Figure 4: Model utility and watermark validity with different choices of M in selecting the next item from the bottom- M items. The dataset is ML-1M. The watermark length is 20 and the initial item is the cold item. For different M , the watermark validity is consistently 100%, which is represented by the red plot. Other plots represent the model utility of different M .

Table 7: Recall@100 in percentage for detecting the watermark in the oracle model of two different selection ranges for the next item in S_{wm} .

	Bottom-100	Top-100
Recall@100	0	34.02

and *pop*. Therefore, we skip the study on the two hyperparameters and focus on the others in this section.

4.7.1 Watermark-to-data Ratio. We show how the watermark-to-data ratio (WDR) influences the performance of AOW in Fig. 3. We show the results on Steam with the popular initial item and watermark length 20. We can observe that, as the WDR increases (the weight of S_{wm} increases), the Recall@1 of the watermark also increases. However, the recall does not further increase significantly when WDR is larger than 1.0. Meanwhile, the NDCG@10 of the model decreases. Therefore, there is a trade-off between the watermark validity and the model utility as the WDR varies. However, the watermark validity is already at a very high level (recall@1>0.8) even if the WDR is low, so in order to preserve the utility of the model, we would suggest to select a low WDR. We by default set the WDR to 0.1 across all experiments.

4.7.2 Bottom- M Item Selection Range. We show how the choice of M influences the model utility in Fig. 4. We use ML-1M and choose the cold item as the initial item, while setting the watermark length as 20. Other datasets and hyperparameter choices yield similar performance. It is clear that the choice of M does not introduce significant performance gap in terms of model utility. Besides, the watermark validity is consistently 100% for all choices of M .

However, it is not true if we claim that M does not matter. A larger M , in fact, may potentially harm the uniqueness of the watermark, causing it to be unintentionally detected in a benign oracle model. Consider two extreme cases, where we set the selection range of the next item in S_{wm} as bottom-100 and top-100 respectively. We would like to see if we can detect the two watermarks from an oracle model. We set the initial item as the cold item, and the watermark length as 20. We train five oracle models with different random seeds on ML-1M, then generate 10 different watermark sequences from each oracle model and test them on other oracles. We report the average Recall@100 of the watermark in Table 7. The

Recall@100 of bottom-100 is zero, while it is 34.02% for top-100. This suggests that incorporating top items into the selection range is risky as it would probably lead to a non-zero detection rate for the watermark in oracle models.

4.7.3 Summary. Here we provide a summary and recommended settings to all the hyperparameters. The watermark length is recommended to be as small as possible to ensure the memorization of it. For the initial item, cold items are generally a better choice since the model can memorize them easily because they are not as robust as the popular items. But popular items may become a better choice when the situations and demands change. WDR introduces a trade-off between the watermark validity and model utility. Setting it between 0.1 to 1.0 is generally a good choice to preserve model utility and ensure a high level of watermark validity. We can set the parameter M of bottom- M as 100 in most cases. But if a longer or diverse watermark is desired, M can be increased. In this case, we need to ensure that M has no overlap with the top- k evaluation range of the system to avoid the unintentional detection of the watermark in an oracle system.

5 Conclusion

In this paper, we propose Autoregressive Out-of-distribution Watermarking (AOW) for watermarking recommender systems. We generate an entire sequence as the watermark and train the model to memorize it. The watermark is evaluated by doing next-item prediction given the truncated watermark sequences. The watermark sequence is generated autoregressively. We first select an initial item, then query the model to obtain scores for all items, and choose a random item from the bottom- M items with least scores. We conduct extensive experiments to demonstrate the effectiveness of AOW in protecting model ownership as well as preserving model utility, and provide comprehensive analysis on the choice of hyperparameters. Applications for AOW on federated recommendation as a fingerprinting technique would be explored in the future.

Acknowledgments

This research is supported by the Ministry of Education, Singapore, under its Academic Research Fund (Tier 2 Award MOE-T2EP20221-0013, Tier 2 Award MOE-T2EP20220-0011, and Tier 1 Award (RG77/21)). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the Ministry of Education, Singapore. This work is partially supported by Australian Research Council under the streams of Future Fellowship (Grant No. FT210100624), Linkage Project (Grant No. LP230200892), Discovery Project (Grants No. DP240101108).

References

- [1] Yossi Adi, Carsten Baum, Moustapha Cisse, Benny Pinkas, and Joseph Keshet. 2018. Turning your weakness into a strength: Watermarking deep neural networks by backdooring. In *27th USENIX Security Symposium (USENIX Security 18)*. 1615–1631.
- [2] Buse Gul Atli, Yuxi Xia, Samuel Marchal, and N Asokan. 2020. Waffle: Watermarking in federated learning. *arXiv preprint arXiv:2008.07298* (2020).
- [3] Huili Chen, Bitu Darvish Rouhani, Cheng Fu, Jishen Zhao, and Farinaz Koushanfar. 2019. Deepmarks: A secure fingerprinting framework for digital rights management of deep learning models. In *Proceedings of the 2019 on International Conference on Multimedia Retrieval*. 105–113.

- [4] Jinyin Chen, Mingjun Li, and Haibin Zheng. 2023. FedRight: An Effective Model Copyright Protection for Federated Learning. *arXiv preprint arXiv:2303.10399* (2023).
- [5] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph neural networks for social recommendation. In *The world wide web conference*. 417–426.
- [6] Jia Guo and Miodrag Potkonjak. 2018. Watermarking deep neural networks for embedded systems. In *2018 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, 1–8.
- [7] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 639–648.
- [8] Hengrui Jia, Christopher A Choquette-Choo, Varun Chandrasekaran, and Nicolas Papernot. 2021. Entangled watermarks as a defense against model extraction. In *30th USENIX Security Symposium (USENIX Security 21)*. 1937–1954.
- [9] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.
- [10] Erwan Le Merrer, Patrick Perez, and Gilles Trédan. 2020. Adversarial frontier stitching for remote neural network watermarking. *Neural Computing and Applications* 32 (2020), 9233–9244.
- [11] Bowen Li, Lixin Fan, Hanlin Gu, Jie Li, and Qiang Yang. 2022. FedIPR: Ownership verification for federated deep neural network models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 4 (2022), 4521–4536.
- [12] Zheng Li, Chengyu Hu, Yang Zhang, and Shanqing Guo. 2019. How to prove your model belongs to you: A blind-watermark based framework to protect intellectual property of DNN. In *Proceedings of the 35th Annual Computer Security Applications Conference*. 126–137.
- [13] Junchuan Liang and Rong Wang. 2023. FedCIP: Federated Client Intellectual Property Protection with Traitor Tracking. *arXiv preprint arXiv:2306.01356* (2023).
- [14] Xiyao Liu, Shuo Shao, Yue Yang, Kangming Wu, Wenyuan Yang, and Hui Fang. 2021. Secure federated learning model verification: A client-side backdoor triggered watermarking scheme. In *2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2414–2419.
- [15] Jing Long, Tong Chen, Guanhua Ye, Kai Zheng, Quoc Viet Hung Nguyen, and Hongzhi Yin. 2024. Physical Trajectory Inference Attack and Defense in Decentralized POI Recommendation. In *Proceedings of the ACM on Web Conference 2024*. 3379–3387.
- [16] Ryota Namba and Jun Sakuma. 2019. Robust watermarking of neural network with exponential weighting. In *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security*. 228–240.
- [17] Liang Qu, Wei Yuan, Ruiqi Zheng, Lizhen Cui, Yuhui Shi, and Hongzhi Yin. 2024. Towards personalized privacy: User-governed data contribution for federated recommendation. In *Proceedings of the ACM on Web Conference 2024*. 3910–3918.
- [18] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [19] J Ben Schafer, Joseph A Konstan, and John Riedl. 2001. E-commerce recommendation applications. *Data mining and knowledge discovery* 5 (2001), 115–153.
- [20] Shuo Shao, Wenyuan Yang, Hanlin Gu, Jian Lou, Zhan Qin, Lixin Fan, Qiang Yang, and Kui Ren. 2022. FedTracker: Furnishing Ownership Verification and Traceability for Federated Learning Model. *arXiv preprint arXiv:2211.07160* (2022).
- [21] Marwa Sharaf, Ezz El-Din Hemdan, Ayman El-Sayed, and Nirmeen A El-Bahnasawy. 2022. A survey on recommendation systems for financial services. *Multimedia Tools and Applications* 81, 12 (2022), 16761–16781.
- [22] 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.
- [23] Sebastian Szlyler, Buse Gul Atli, Samuel Marchal, and N Asokan. 2021. Dawn: Dynamic adversarial watermarking of neural networks. In *Proceedings of the 29th ACM International Conference on Multimedia*. 4417–4425.
- [24] Yusuke Uchida, Yuki Nagai, Shigeyuki Sakazawa, and Shin'ichi Satoh. 2017. Embedding watermarks into deep neural networks. In *Proceedings of the 2017 ACM on international conference on multimedia retrieval*. 269–277.
- [25] Jiangfeng Wang, Hanzhou Wu, Xinpeng Zhang, and Yuwei Yao. 2020. Watermarking in deep neural networks via error back-propagation. *Electronic Imaging* 2020, 4 (2020), 22–1.
- [26] Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. 2022. Graph neural networks in recommender systems: a survey. *Comput. Surveys* 55, 5 (2022), 1–37.
- [27] Jing Xu, Stefanos Koffas, Oguzhan Ersoy, and Stjepan Picek. 2021. Watermarking graph neural networks based on backdoor attacks. *arXiv preprint arXiv:2110.11024* (2021).
- [28] Wenyuan Yang, Shuo Shao, Yue Yang, Xiyao Liu, Zhihua Xia, Gerald Schaefer, and Hui Fang. 2022. Watermarking in Secure Federated Learning: A Verification Framework Based on Client-Side Backdooring. *arXiv preprint arXiv:2211.07138* (2022).
- [29] Ziqi Yang, Hung Dang, and Ee-Chien Chang. 2019. Effectiveness of distillation attack and countermeasure on neural network watermarking. *arXiv preprint arXiv:1906.06046* (2019).
- [30] Wei Yuan, Chaoqun Yang, Quoc Viet Hung Nguyen, Lizhen Cui, Tiek He, and Hongzhi Yin. 2023. Interaction-level membership inference attack against federated recommender systems. In *Proceedings of the ACM Web Conference 2023*. 1053–1062.
- [31] Wei Yuan, Chaoqun Yang, Liang Qu, Quoc Viet Hung Nguyen, Jianxin Li, and Hongzhi Yin. 2023. Hide Your Model: A Parameter Transmission-free Federated Recommender System. *arXiv preprint arXiv:2311.14968* (2023).
- [32] Wei Yuan, Shilong Yuan, Chaoqun Yang, Nguyen Quoc Viet hung, and Hongzhi Yin. 2023. Manipulating Visually Aware Federated Recommender Systems and Its Countermeasures. *ACM Transactions on Information Systems* 42, 3 (2023), 1–26.
- [33] Wenbin Yue, Zidong Wang, Jieyu Zhang, and Xiaohui Liu. 2021. An overview of recommendation techniques and their applications in healthcare. *IEEE/CAA Journal of Automatica Sinica* 8, 4 (2021), 701–717.
- [34] 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.
- [35] Jialong Zhang, Zhongshu Gu, Jiyong Jang, Hui Wu, Marc Ph Stoecklin, Heqing Huang, and Ian Molloy. 2018. Protecting intellectual property of deep neural networks with watermarking. In *Proceedings of the 2018 on Asia conference on computer and communications security*. 159–172.
- [36] Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. 2019. Deep learning based recommender system: A survey and new perspectives. *ACM computing surveys (CSUR)* 52, 1 (2019), 1–38.
- [37] Sixiao Zhang, Hongzhi Yin, Hongxu Chen, and Cheng Long. 2024. Defense Against Model Extraction Attacks on Recommender Systems. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*. 949–957.
- [38] Shijie Zhang, Wei Yuan, and Hongzhi Yin. 2023. Comprehensive privacy analysis on federated recommender system against attribute inference attacks. *IEEE Transactions on Knowledge and Data Engineering* (2023).
- [39] Xiangyu Zhao, Hanzhou Wu, and Xinpeng Zhang. 2021. Watermarking graph neural networks by random graphs. In *2021 9th International Symposium on Digital Forensics and Security (ISDFS)*. IEEE, 1–6.
- [40] Renjie Zhu, Xinpeng Zhang, Mengte Shi, and Zhenjun Tang. 2020. Secure neural network watermarking protocol against forging attack. *EURASIP Journal on Image and Video Processing* 2020 (2020), 1–12.