# Bayesian Privacy Guarantee for User History in Sequential Recommendation Using Randomised Response

Wenchuan Mu
wenchuan_mu@sutd.edu.sg
Singapore University of Technology and Design
Singapore

Kwan Hui Lim
kwanhui@acm.org
Singapore University of Technology and Design
Singapore

## ABSTRACT

Sequential recommendation systems play an important role in delivering personalised user experiences, yet they rely heavily on detailed user history, raising serious privacy concerns. In this work, we introduce a novel framework that integrates a randomised response mechanism into sequential recommendation to provide strong privacy guarantees while preserving recommendation effectiveness. By obfuscating user history through controlled probabilistic item substitution based on semantic similarity, our approach ensures that released sequences protect individual behaviour with provable Bayesian posterior privacy. We further propose training strategies tailored for privacy-filtered data, including a frequency-based vocabulary expansion method inspired by subword tokenisation. Experiments on four real-world datasets demonstrate that our approach preserves recommendation quality under strong privacy constraints and outperforms existing baselines even without applying privacy filters.

## CCS CONCEPTS

• **Security and privacy** → *Privacy-preserving protocols*; • **Information systems** → *Recommender systems.*

## KEYWORDS

Sequential recommendation, Randomised response

**ACM Reference Format:**
Wenchuan Mu and Kwan Hui Lim. 2025. Bayesian Privacy Guarantee for User History in Sequential Recommendation Using Randomised Response. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management (CIKM '25), November 10–14, 2025, Seoul, Republic of Korea.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3746252.3760856

## 1 INTRODUCTION

Recommender systems play a critical role in providing personalised content [2], where the goal is to suggest relevant items to users from vast item collections [21]. The current focus is on building effective recommender systems that accurately model user preferences based on their historical interaction data [30]. With the progress of deep learning [15], deep sequential recommendation (SR) models have become a leading approach [14, 35] and have shown strong performance across various recommendation baselines [32].

Despite the remarkable progress in SR, most existing methods rely heavily on historical user interaction data. This dependence raises serious concerns about user privacy [39], as collecting and storing such detailed behavioural data can expose sensitive personal information. While recent research has largely focused on improving recommendation effectiveness [12, 13, 17, 20, 28, 33], privacy is also an important concern. As a result, there is a growing need for recommendation approaches that can balance effectiveness with strong privacy protections.

Existing work has explored ways to reduce the reliance of SR systems on user data [3, 8, 23, 27, 31]. For example, federated training can be used to limit exposure to sensitive information [19, 24]. However, such approaches still require detailed user histories at inference time, which introduces privacy risks, particularly when transmitted data are sent to remote servers [36]. Differential privacy [6] has also been explored by Hu and Fang [16] to inject noise into graph neural networks, but this leads to degraded performance [36, 41]. Other efforts [36] focus on obfuscating input data but lack formal privacy guarantees. As a result, there remains a critical gap: how to ensure provable privacy for user history while enabling SR to work effectively on privacy-preserving inputs

In this work, we propose a privacy-preserving framework for sequential recommendation that applies a randomised response [37] mechanism to user history before it is sent to the recommender system. Instead of sharing exact interaction sequences, each item is probabilistically replaced with a semantically similar [25] alternative, providing formal Bayesian privacy guarantees. This obfuscation is applied at the data level and operates independently of the model architecture. To better fit the resulting privacy-filtered sequences, we introduce a training strategy based on frequency-driven vocabulary expansion inspired by subword tokenisation, enabling the model to capture meaningful sequential patterns from perturbed inputs. Experiments on several real-world datasets demonstrate that our approach retains strong performance even under tight privacy constraints and outperforms existing baselines in standard (non-private) settings.

## 2 PRELIMINARIES IN THE SEQUENTIAL RECOMMENDATION PROBLEM

### 2.1 Sequences and Next-Item Recommendation

Essentially, sequential recommendation recommends a sequence and/or takes sequential data as input to make recommendations. For instance, a university library system notices a student reads


*CIKM '25, November 10–14, 2025, Seoul, Republic of Korea*

ACM ISBN 979-8-4007-2040-6/2025/11
https://doi.org/10.1145/3746252.3760856

"Pride and Prejudice" followed by "Jane Eyre". Using sequential recommendation, it suggests "Wuthering Heights" and *then* "Middlemarch", continuing the student's journey through reading classic British literature.

Formally, we let random variable $\mathbf{u}$ denote a user, *e.g.*,, the student who uses the library system. Also, we denote the item set, *e.g.*,, all books, using a (possibly infinite) discrete set $\mathbb{V}$ where each item, *e.g.*,, "Middlemarch", is an element of $\mathbb{V}$. Then, we can define a sequence as follows.

*Definition 2.1 (Sequence).* Let $\mathbb{V}$ denote the set of items. A sequence is a discrete stochastic process $(\mathbf{v}_t)_{t\in\mathbb{Z}}$, where each $\mathbf{v}_t$ is a random variable taking values in $\mathbb{V}$. If a clear start of the sequence exists, we denote this process as $(\mathbf{v}_t)_{t\in\mathbb{N}}$.

*Item recommendation.* The recommendation process can be modelled as a function $f_k : \mathbb{U} \times \mathbb{V}^k \to \mathbb{V}$, if $k$ items in a sequence are given as input. The most straightforward case is when $v_0, v_1, \ldots, v_{k-1}$ are given, whereas the goal is to recommend a good $v_k$ to user $u$. There are also other scenarios like directly recommending $v_{k+1}$ based on $v_0, v_1, \ldots, v_{k-1}$, *i.e.*, $v_k$ is skipped. Recommending a sequence of $m$ items, *e.g.*,, $v_k, v_{k+1}, \ldots, v_{k+m-1}$, can be achieved through $m$ times of item recommendation, *i.e.*, let $f_{k,+1}$ denote the straightforward next-item recommendation function, then we obtain

$$\begin{aligned} v_k &= f_{k,+1}\left(u, v_0, v_1, \ldots, v_{k-1}\right) \\ v_{k+1} &= f_{k+1,+1}\left(u, v_0, v_1, \ldots, v_k\right) \\ &\cdots \\ v_{k+m-1} &= f_{k+m-1,+1}\left(u, v_0, v_1, \ldots, v_{k+m-2}\right). \end{aligned} \tag{1}$$

Intuitively, a fundamental research objective in sequential recommendation is improving the recommender $f_k$ to better fit user preferences, possibly via various objective functions [5, 13, 14, 22, 38].

## 2.2 Machine Learning in Sequential Recommendation

The recommender function $f_k : \mathbb{U} \times \mathbb{V}^k \to \mathbb{V}$ can often be learned. That is, we optimise $f_k$ to let it fit the joint distribution of $k$ known items and the targeted one. For instance, consider a simplest form of Markov chain, which assumes that the next item depends only on $k$ preceding items. Formally, for a user $\mathbf{u}$ and a sequence history $(\mathbf{v}_t)_{t\in(\mathbb{N}\cap[0,k'))}, k'>k$, the Markov assumption is

$$\begin{aligned} \forall t > k.\ & P(\mathbf{v}_t = v \mid \mathbf{u} = u, \mathbf{v}_0 = v_0, \ldots, \mathbf{v}_{t-1} = v_{t-1}) \\ &= P(\mathbf{v}_t = v \mid \mathbf{u} = u, \mathbf{v}_{t-k} = v_{t-k}, \ldots, \mathbf{v}_{t-1} = v_{t-1}). \end{aligned} \tag{2}$$

The Markov property enables efficient user behaviour modelling without requiring the full sequence history [10, 11, 29]. The corresponding recommender function may then take the form

$$\begin{aligned} &f^{\text{Markov}}_{k,+1}(u, v_{t-k}, \ldots, v_{t-1}) \coloneqq \\ &\quad \arg\max_{v\in\mathbb{V}} \hat{P}(\mathbf{v}_t = v \mid \mathbf{u} = u, \mathbf{v}_{t-k} = v_{t-k}, \ldots, \mathbf{v}_{t-1} = v_{t-1}), \end{aligned} \tag{3}$$

where $\hat{P}(\mathbf{v}_t = v \mid \mathbf{u} = u, \mathbf{v}_{t-k} = v_{t-k}, \ldots, \mathbf{v}_{t-1} = v_{t-1})$ is the estimated probability of item $\mathbf{v}_t$ and can be expressed by a parametrised function $g : \Theta \times \mathbb{U} \times \mathbb{V}^{k+1} \to \mathbb{R}$. Here, $\Theta$ denotes the set of all possible parameters, *e.g.*,, a set of neural network models. Intuitively, the recommender $f^{\text{Markov}}_{k,+1}$ can be optimised by minimising the Kullback–Leibler divergence (KL divergence) between the given empirical distribution and the estimated distribution, *i.e.*,

$$\min_{\theta\in\Theta} \quad -\sum_u \sum_{t\geq k} \log g\left(\theta, u, v_{t-k,u}, \ldots, v_{t-1,u}, v_{t,u}\right). \tag{4}$$

*Bidirectional condition.* Sometimes, recommendations are conditioned not only on past items but also on future ones. For example, if a student has read "Pride and Prejudice" and "Jane Eyre", and we know they are later expected to read "Middlemarch", but the book in between is unknown. The system may recommend "Wuthering Heights" in this case.

Here, the recommender $f_k$ can be learned more flexibly than in the purely unidirectional setting [32]. For any sequence containing $k+m$ items, the recommender training can be supervised by any $m$ items in it, rather than just the last $m$ items. Moreover, each of the $m$ supervising items can condition on any other $k$ items in the sequence, rather than $k$ consecutive items. The corresponding recommender function may then take the following form

$$\begin{aligned} &f^{\text{Bi}}_k(u, v_{i_1}, \ldots, v_{i_k}, i_1, \ldots, i_k, i_{\text{target}}) \coloneqq \\ &\quad \arg\max_{v\in\mathbb{V}} \hat{P}(\mathbf{v}_{i_{\text{target}}} = v \mid \mathbf{u} = u, \mathbf{v}_{i_1} = v_{i_1}, \ldots, \mathbf{v}_{i_k} = v_{i_k}). \end{aligned} \tag{5}$$

Similar to the unidirectional case, we could optimise $f^{\text{Bi}}_k$ by minimising the KL divergence between the empirical distribution and the estimated distribution.

# 3 RANDOMISED-RESPONSE SEQUENTIAL RECOMMENDATION

In this section, we introduce the randomised-response sequential recommendation to protect user history privacy while maintaining recommendation effectiveness. We first present the randomised response mechanism that guarantees privacy during sequential recommendation. Then we describe how the recommender can be trained to accommodate this privacy-preserving setting.

## 3.1 Protecting Data Release with Randomised Response

In sequential recommendation, user history is provided to the system to generate the next item. When the system is hosted remotely, this history must be uploaded, even if only a portion (*e.g.*,, in Markov settings) is used. Since the exact user history often contains confidential information, transmitting it poses privacy risks.

Therefore, we aim to investigate whether we can filter out private information in user history, with formal guarantees of privacy protection. Further, if such a guarantee is obtainable, what does the guarantee suggest?

To address this, we propose a randomised response mechanism integrated into the sequential recommendation process, as described in Algorithm 1. Intuitively, Algorithm 1 first assigns an alternative item to each original item (Lines 1 - 10), independently of any user or historical data. Then, each item in the user history is randomly replaced with either its original or alternative value (Lines 11 - 18). The resulting history, disentangled from private information, is then used for recommendation (Line 19).

**Algorithm 1** Randomised-Response Sequential Recommendation

**Require:** $f_{n,+1}$ (Recommender function), $u$ (user), $v_1, v_2, \ldots, v_{n-1}$ (history sequence), $\mathbb{V}$ (item set), $h : \mathbb{V} \to \mathbb{R}^d$ (embedding function), $\epsilon$ (privacy parameter)
**Ensure:** Next item in sequence $v_n$
1: Define mapping of alternatives $\mathcal{A} \leftarrow \{\ \}$
2: **for** $v' \in \mathbb{V}$ **do**
3: $d_{\min} \leftarrow \infty$
4: **for** $v'' \in \mathbb{V} \setminus \{v'\}$ **do**
5: $d \leftarrow \|h(v') - h(v'')\|_2$
6: **if** $d < d_{\min}$ **then**
7: $\mathcal{A}[v'] \leftarrow v''$, also $d_{\min} \leftarrow d$
8: **end if**
9: **end for**
10: **end for**
11: **for** $t = 0, 1, \ldots, n-1$ **do**
12: Sample $q \sim \text{Uniform}(0, 1)$
13: **if** $q > \sqrt[n]{\epsilon/(\epsilon+1)}$ **then**
14: $v_t^{\mathrm{p}} \leftarrow \mathcal{A}[v_t]$
15: **else**
16: $v_t^{\mathrm{p}} \leftarrow v_t$
17: **end if**
18: **end for**
19: $v_n \leftarrow f_{n,+1}\left(u, v_0^{\mathrm{p}}, v_1^{\mathrm{p}}, \ldots, v_{n-1}^{\mathrm{p}}\right)$

*Choose alternative items.* We embed items into dense vectors to capture their semantic and contextual meaning. Then, we select the nearest Euclidean neighbour as the alternative. For example, LLaMA-2 [34] embeddings provide both semantic depth and grounded real-world detail. Intuitively, the alternative item is expected to be closely related to the original one to keep the history sequence reasonable.

*Compute the privacy guarantee.* The privacy is protected when we transform the original history $v_0, v_1, \ldots, v_{n-1}$ to the randomised-response sequence $v_0^{\mathrm{p}}, v_1^{\mathrm{p}}, \ldots, v_{n-1}^{\mathrm{p}}$. For example, if $n = 4$ and the original history is ["Pride and Prejudice", "Wuthering Heights", "Jane Eyre", "Great Expectations"], we apply a randomised process to each item, *e.g.*,, with probability 0.54, an item is kept unchanged, and with probability 0.46, it is replaced by a fixed alternative (*e.g.*,, "Pride and Prejudice" might be replaced with "Persuasion", and "Jane Eyre" with "Jane (April Lindner)"). As a result, the released history might be ["Persuasion", "Wuthering Heights", "Jane (April Lindner)", "Great Expectations"]. In this case, an attacker cannot guess the original sequence with more than a 1 in 10 odds, making the exact reconstruction of reading history highly unlikely. In a general case, a formal Bayesian guarantee is as follows.

Theorem 3.1 (Posterior Privacy). *Let $v_0, v_1, \ldots, v_{n-1} \in \mathbb{V}^n$ be a sequence history of n items. Suppose each value is obfuscated by the randomised response mechanism*

$$\mathbf{R}(v) = \begin{cases} a(v) & \text{with probability } 1 - \left(\frac{\epsilon}{1+\epsilon}\right)^{1/n} \\ v & \text{with probability } \left(\frac{\epsilon}{1+\epsilon}\right)^{1/n}, \end{cases} \tag{6}$$

*where $a : \mathbb{V} \to \mathbb{V}$ is a deterministic injective mapping such that $a(v') \neq v'$ for all $v' \in \mathbb{V}$. Then for any privacy intruding guess $(\hat{\mathbf{v}}_t)_{t=0}^{n-1}$ from sequence $\mathbf{R}(v_0), \ldots, \mathbf{R}(v_{n-1})$, the posterior odds of correctly recovering the original sequence are bounded by $\epsilon$,* i.e.,

$$\frac{P\left((\hat{\mathbf{v}}_t)_{t=0}^{n-1} = v_0, \ldots, v_n \mid \mathbf{R}(v_0), \ldots, \mathbf{R}(v_{n-1})\right)}{P\left((\hat{\mathbf{v}}_t)_{t=0}^{n-1} \neq v_0, \ldots, v_n \mid \mathbf{R}(v_0), \ldots, \mathbf{R}(v_{n-1})\right)} \leq \epsilon. \tag{7}$$

Proof. Let $p < 1/2$ denote the probability parameter of the randomised response Bernoulli trial [40], *i.e.*, the probability of replacing an item with its alternative. Then, the probability of correctly guessing the full sequence is $(1-p)^n$, and the probability of any other sequence is $1-(1-p)^n$. Thus, for any $1-p \leq \sqrt[n]{\epsilon/(1+\epsilon)}$, the posterior odds of the correct sequence given the obfuscated are bounded by $(1-p)^n/(1-(1-p)^n)$, thus upper bounded by $\epsilon$. □

Intuitively, even with knowledge of the mechanism and prior, the probability of recovering the private history is upper-bounded. In practice, given any specified privacy parameter $\epsilon$ and the alternative of each item, we can generate a privacy-filtered sequence for the original history (Lines 11 - 18 in Algorithm 1).

Note that Theorem 3.1 holds regardless of which alternative is chosen for each item. The alternative item choice is a matter of the sequence content. If the alternative is semantically related, the sequence changes less. If the replacement is arbitrary, the result may differ greatly, potentially affecting downstream recommendations.

## 3.2 Training Sequential Recommender for Privacy-filtered Data

Privacy-filtered sequences pose unique challenges for recommender training. Injected randomness disrupts local item continuity, while replacements can introduce semantic drift. To improve generalizability and reduce sparsity, we introduce a frequency-based composition strategy inspired by subword tokenisation methods such as byte pair encoding [7]. Still, adaptation to fit our framework is necessary. Specifically, given the item set $\mathbb{V}$, we iteratively merge frequent adjacent item pairs $(v_t, v_{t+1}) \in \mathbb{V}^2$ into composite units:

$$\mathbb{V}' \leftarrow \mathbb{V} \cup \{v_t\_v_{t+1}\} \tag{8}$$

Intuitively, the item set keeps growing till its size meets a pre-set number, *e.g.*,, 20,000 for Steam.

Combined with bidirectional modelling as in Equation (5), this item composition approach helps the model recover meaningful structure from both directions, even when items are actually replaced with alternative items (*i.e.*, $v_t^{\mathrm{p}} \neq v_t$ with some probability $p$). Note that we do not apply randomised response to labels in order to preserve a reliable training signal for the model, ensuring that supervision remains grounded in true user behaviour.

Compared to existing bidirectional methods such as BERT4Rec [4, 32], our adaptive vocabulary expansion reduces sparsity and captures semantically meaningful segments, enabling the model to operate over higher-level patterns that are more resilient to noise. For example, even if one token in a frequent phrase is replaced, the composed segment may remain familiar and reasonable.

# 4 EXPERIMENT

While we can pre-know how much privacy has been protected by setting the parameter $\epsilon$, we need empirical experiments to answer

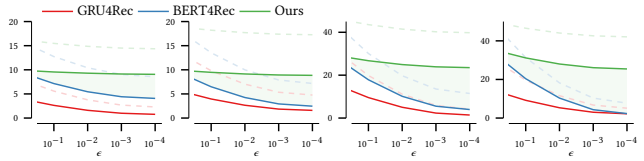


**Figure 1: Performance comparison of next-item prediction with varying privacy constraint. HR at 1 is in solid lines, and NDCG at 5 is in dashed lines. A shaded area illustrates the performance gain of our method when the privacy constraint becomes stricter, *i.e.*, the posterior difference ($\epsilon$) decreases.**

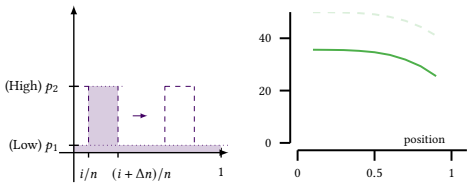


**Figure 2: When the same privacy constraint $\epsilon$ is satisfied, the position of high-privacy region, as illustrated in (a), may still affect the performance, as shown in (b).**

the following research questions. (1) How does the privacy constraint affect the performance of the randomised-response sequential recommender? (2) What is the role of the position of privacy-filtered items within a sequence?

*Essential settings.* We run experiments [1] on four real-world datasets: Amazon Beauty, Steam, MovieLens 1M (ML-1m), and MovieLens 20M (ML-20m) [9], which vary significantly in domain, size, and sparsity. For evaluation, we adopt the widely used leave-one-out protocol [12, 20, 33], *i.e.*, the last interaction is used for testing, the second-to-last for validation, and the rest for training. Each test item is ranked against 100 negative samples [17, 32]. We report Hit Ratio (HR) at 1, and Normalised Discounted Cumulative Gain (NDCG) at 5 [18]. To benchmark performance, we compare against representative baselines, GRU4Rec [14] and BERT4Rec [32].

*RQ1: Sequential recommendation performance with varying privacy constraints.* We vary the privacy parameter $\epsilon$ at a series of values (from 1/10 to 1/10,000) to evaluate the recommendation performance on the privacy-filtered history. At each $\epsilon$ level, we report both HR at 1 and NCDG at 5 for each method.

Figure 1 presents the results of this experiment. The proposed training method consistently outperforms all baselines across the four datasets, achieving at least a 2% improvement. More importantly, our method maintains stable HR and NDCG performance even as the privacy constraint becomes more stringent, approaching $\epsilon = 1/10{,}000$. Specifically, as $\epsilon$ decreases from 1/10 to 1/10,000, our method experiences an average (relative) drop of 16% in HR and 12% in NDCG. This drop is significantly smaller than that of the baselines, which have at least a 66% drop in both metrics. Overall, this result shows that maintaining performance while guaranteeing privacy is realistically achievable.

We also observe that the sequence length of user history affects the stability of performance under varying privacy constraints. For shorter sequences, such as those in Amazon Beauty and Steam (with an average history length of 10 items), the performance drop is relatively small, around 10±0.3% in both HR and NDCG. In contrast, for longer sequences (averaging 145 items), we observe a 24% drop in HR and a 15% drop in NDCG. A similar trend is seen in baseline methods. This may be because in shorter sequences, the next item has weaker dependency on prior history, reflected in their overall lower performance (about 40% of that for longer sequences). As a result, replacing some items with alternatives has a limited impact. In longer sequences, however, models may rely on stronger sequential patterns or shortcuts for next-item prediction, which are more easily disrupted by such changes.

*RQ2: Role of privacy-filtered position.* In some cases, the user-defined privacy constraint allows flexibility in setting the probability for randomised response. That is, an extremely low value of $\sqrt[n]{\epsilon/(\epsilon+1)}$ is not always needed to meet the required privacy level. This opens the possibility of choosing which parts of the sequence history to obfuscate more strongly. For instance, we can define two probabilities, $0 < p_1 < p_2 < 1/2$, where replacing an item with probability $p_2$ provides stronger obfuscation than with $p_1$. If we apply stronger obfuscation (with $p_2$) to $\Delta n$ items in an $n$-length sequence, $p_1$ and $p_2$ should satisfy $(1-p_1)^{n-\Delta n} = \epsilon(1-p_2)^{-\Delta n}/(1+\epsilon)$.

This approach helps us understand how strongly obfuscated parts of the sequence affect next-item prediction performance. In the experiment, we set $p_1 = 0.1$ and applied $p_2$ to $\Delta n = 0.2n$ consecutive items in each sequence. Figure 2 presents the MovieLens 20M result. As the high-obfuscation region shifts from the beginning to the end of the sequence, we observe a gradual increase in performance drop. This is likely due to the nature of real-world sequence data, which often aligns with the Markov assumption, *i.e.*, recent items correlate with the next item more strongly than far-off items do. This finding also suggests that even under the same overall privacy level, the choice of obfuscation strategy can lead to different impacts on model performance.

## 5 CONCLUSION

We introduce a formal Bayesian privacy guarantee with a randomised response mechanism, letting users bound the probability that an attacker can infer their history. This strict constraint ensures all recommendations meet a rigorous privacy standard. Furthermore, to address sparsity from randomised response, we improve bidirectional modelling by merging frequent bi-grams into composite items, improving stability as privacy tightens. Our results confirm the feasibility of preserving user-history privacy without compromising recommendation performance. In future, we will extend our discussion on the Bayesian privacy guarantee compared to differential privacy guarantees on sequential recommendation, as well as investigate its robustness against adversarial attacks [26].

**Acknowledgements**. This research is supported in part by the Ministry of Education, Singapore, under its Academic Research Fund Tier 2 (Award No. MOE-T2EP20123-0015). Any opinions, findings and conclusions, or recommendations expressed in this material are those of the authors and do not reflect the views of the Ministry of Education, Singapore.

## GENAI USAGE DISCLOSURE

Generative AI tools were used for minor language editing, such as grammar/spelling checks and minor changes, on original author-written text.

## REFERENCES


[1] 2025. Code and data for "Bayesian Privacy Guarantee for User History in Sequential Recommendation Using Randomised Response". https://github.com/cestwc/sequential-recommendation-bayesian-privacy.

[2] G. Adomavicius and A. Tuzhilin. 2005. Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering* 17, 6 (2005), 734–749. https://doi.org/10.1109/TKDE.2005.99

[3] Arnaud Berlioz, Arik Friedman, Mohamed Ali Kaafar, Roksana Boreli, and Shlomo Berkovsky. 2015. Applying Differential Privacy to Matrix Factorization. In *Proceedings of the 9th ACM Conference on Recommender Systems* (Vienna, Austria) *(RecSys '15)*. Association for Computing Machinery, New York, NY, USA, 107–114. https://doi.org/10.1145/2792838.2800173

[4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. https://doi.org/10.18653/v1/N19-1423

[5] Tim Donkers, Benedikt Loepp, and Jürgen Ziegler. 2017. Sequential User-based Recurrent Neural Network Recommendations. In *Proceedings of the Eleventh ACM Conference on Recommender Systems* (Como, Italy) *(RecSys '17)*. Association for Computing Machinery, New York, NY, USA, 152–160. https://doi.org/10.1145/3109859.3109877

[6] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 265–284.

[7] Philip Gage. 1994. A new algorithm for data compression. *C Users J.* 12, 2 (Feb. 1994), 23–38.

[8] Jialiang Han, Yun Ma, Qiaozhu Mei, and Xuanzhe Liu. 2021. DeepRec: On-device Deep Learning for Privacy-Preserving Sequential Recommendation in Mobile Commerce. In *Proceedings of the Web Conference 2021* (Ljubljana, Slovenia) *(WWW '21)*. Association for Computing Machinery, New York, NY, USA, 900–911. https://doi.org/10.1145/3442381.3449942

[9] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Trans. Interact. Intell. Syst.* 5, 4, Article 19 (Dec. 2015), 19 pages. https://doi.org/10.1145/2827872

[10] Ruining He, Wang-Cheng Kang, and Julian McAuley. 2017. Translation-based Recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems* (Como, Italy) *(RecSys '17)*. Association for Computing Machinery, New York, NY, USA, 161–169. https://doi.org/10.1145/3109859.3109882

[11] Ruining He and Julian McAuley. 2016. Fusing Similarity Models with Markov Chains for Sparse Sequential Recommendation. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*. 191–200. https://doi.org/10.1109/ICDM.2016.0030

[12] 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* (Perth, Australia) *(WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 173–182. https://doi.org/10.1145/3038912.3052569

[13] Balázs Hidasi and Alexandros Karatzoglou. 2018. Recurrent Neural Networks with Top-k Gains for Session-based Recommendations. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (Torino, Italy) *(CIKM '18)*. Association for Computing Machinery, New York, NY, USA, 843–852. https://doi.org/10.1145/3269206.3271761

[14] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based recommendations with recurrent neural networks. In *International Conference on Learning Representations (ICLR)*. https://arxiv.org/abs/1511.06939

[15] Liang Hu, Longbing Cao, Shoujin Wang, Guandong Xu, Jian Cao, and Zhiping Gu. 2017. Diversifying Personalized Recommendation with User-session Context. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*. 1858–1864. https://doi.org/10.24963/ijcai.2017/258

[16] Wentao Hu and Hui Fang. 2024. Towards Differential Privacy in Sequential Recommendation: A Noisy Graph Neural Network Approach. *ACM Trans. Knowl. Discov. Data* 18, 5, Article 125 (March 2024), 21 pages. https://doi.org/10.1145/3643821

[17] Jin Huang, Wayne Xin Zhao, Hongjian Dou, Ji-Rong Wen, and Edward Y. Chang. 2018. Improving Sequential Recommendation with Knowledge-Enhanced Memory Networks. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval* (Ann Arbor, MI, USA) *(SIGIR '18)*. Association for Computing Machinery, New York, NY, USA, 505–514. https://doi.org/10.1145/3209978.3210017

[18] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Trans. Inf. Syst.* 20, 4 (Oct. 2002), 422–446. https://doi.org/10.1145/582415.582418

[19] Santosh Kabbur, Xia Ning, and George Karypis. 2013. FISM: factored item similarity models for top-N recommender systems. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Chicago, Illinois, USA) *(KDD '13)*. Association for Computing Machinery, New York, NY, USA, 659–667. https://doi.org/10.1145/2487575.2487589

[20] Wang-Cheng Kang and Julian McAuley. 2018. Self-Attentive Sequential Recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. 197–206. https://doi.org/10.1109/ICDM.2018.00035

[21] Donghyun Kim, Chanyoung Park, Jinoh Oh, Sungyoung Lee, and Hwanjo Yu. 2016. Convolutional Matrix Factorization for Document Context-Aware Recommendation. In *Proceedings of the 10th ACM Conference on Recommender Systems* (Boston, Massachusetts, USA) *(RecSys '16)*. Association for Computing Machinery, New York, NY, USA, 233–240. https://doi.org/10.1145/2959100.2959165

[22] 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* (Singapore, Singapore) *(CIKM '17)*. Association for Computing Machinery, New York, NY, USA, 1419–1428. https://doi.org/10.1145/3132847.3132926

[23] Ziqi Liu, Yu-Xiang Wang, and Alexander Smola. 2015. Fast Differentially Private Matrix Factorization. In *Proceedings of the 9th ACM Conference on Recommender Systems* (Vienna, Austria) *(RecSys '15)*. Association for Computing Machinery, New York, NY, USA, 171–178. https://doi.org/10.1145/2792838.2800191

[24] Wu Meihan, Li Li, Chang Tao, Eric Rigall, Wang Xiaodong, and Xu Cheng-Zhong. 2022. FedCDR: Federated Cross-Domain Recommendation for Privacy-Preserving Rating Prediction. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management* (Atlanta, GA, USA) *(CIKM '22)*. Association for Computing Machinery, New York, NY, USA, 2179–2188. https://doi.org/10.1145/3511808.3557320

[25] Wenchuan Mu and Kwan Hui Lim. 2024. Modelling Text Similarity: A Survey. In *Proceedings of the 2023 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining* (Kusadasi, Turkiye) *(ASONAM '23)*. Association for Computing Machinery, New York, NY, USA, 698–705. https://doi.org/10.1145/3625007.3627305

[26] Wenchuan Mu and Kwan Hui Lim. 2025. Get Global Guarantees: On the Probabilistic Nature of Perturbation Robustness. In *Proceedings of the 34rd ACM International Conference on Information and Knowledge Management* (Seoul, Republic of Korea) *(CIKM '25)*. Association for Computing Machinery, New York, NY, USA, 11 pages. https://doi.org/10.1145/3746252.3761039

[27] Jiaqian Ren, Lei Jiang, Hao Peng, Lingjuan Lyu, Zhiwei Liu, Chaochao Chen, Jia Wu, Xu Bai, and Philip S. Yu. 2022. Cross-Network Social User Embedding with Hybrid Differential Privacy Guarantees. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management* (Atlanta, GA, USA) *(CIKM '22)*. Association for Computing Machinery, New York, NY, USA, 1685–1695. https://doi.org/10.1145/3511808.3557278

[28] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence* (Montreal, Quebec, Canada) *(UAI '09)*. AUAI Press, Arlington, Virginia, USA, 452–461.

[29] 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* (Raleigh, North Carolina, USA) *(WWW '10)*. Association for Computing Machinery, New York, NY, USA, 811–820. https://doi.org/10.1145/1772690.1772773

[30] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th International Conference on World Wide Web* (Hong Kong, Hong Kong) *(WWW '01)*. Association for Computing Machinery, New York, NY, USA, 285–295. https://doi.org/10.1145/371920.372071

[31] Hyejin Shin, Sungwook Kim, Junbum Shin, and Xiaokui Xiao. 2018. Privacy Enhanced Matrix Factorization for Recommendation with Local Differential Privacy. *IEEE Transactions on Knowledge and Data Engineering* 30, 9 (2018), 1770–1782. https://doi.org/10.1109/TKDE.2018.2805356

[32] 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* (Beijing, China) *(CIKM '19)*. Association for Computing Machinery, New York, NY, USA, 1441–1450. https://doi.org/10.1145/3357384.3357895

[33] 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* (Marina Del Rey, CA,

USA) *(WSDM '18)*. Association for Computing Machinery, New York, NY, USA, 565–573. https://doi.org/10.1145/3159652.3159656

[34] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, and Shruti Bhosale andand others. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. *CoRR* abs/2307.09288 (2023). https://doi.org/10.48550/ARXIV.2307.09288 arXiv:2307.09288

[35] Shoujin Wang, Liang Hu, Longbing Cao, Xiaoshui Huang, Defu Lian, and Wei Liu. 2018. Attention-Based Transactional Context Embedding for Next-Item Recommendation. *Proceedings of the AAAI Conference on Artificial Intelligence* 32, 1 (Apr. 2018). https://doi.org/10.1609/aaai.v32i1.11851

[36] Wei Wang, Yujie Lin, Pengjie Ren, Zhumin Chen, Tsunenori Mine, Jianli Zhao, Qiang Zhao, Moyan Zhang, Xianye Ben, and Yujun Li. 2025. Privacy-Preserving Sequential Recommendation with Collaborative Confusion. *ACM Trans. Inf. Syst.* 43, 2, Article 50 (Jan. 2025), 25 pages. https://doi.org/10.1145/3707204

[37] Stanley L. Warner. 1965. Randomized Response: A Survey Technique for Eliminating Evasive Answer Bias. *J. Amer. Statist. Assoc.* 60, 309 (1965), 63–69. https://doi.org/10.1080/01621459.1965.10480775 PMID: 12261830.

[38] Feng Yu, Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. 2016. A Dynamic Recurrent Model for Next Basket Recommendation. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Pisa, Italy) *(SIGIR '16)*. Association for Computing Machinery, New York, NY, USA, 729–732. https://doi.org/10.1145/2911451.2914683

[39] Hongyu Zhang, Dongyi Zheng, Xu Yang, Jiyuan Feng, and Qing Liao. 2024. FedDCSR: Federated Cross-domain Sequential Recommendation via Disentangled Representation Learning. In *Proceedings of the 2024 SIAM International Conference on Data Mining (SDM)*. 535–543. https://doi.org/10.1137/1.9781611978032.62

[40] Ruihan Zhang and Jun Sun. 2025. Correct-by-Construction: Certified Individual Fairness through Neural Network Training. *Proc. ACM Program. Lang.* 9, OOPSLA2, Article 329 (oct 2025). https://doi.org/10.1145/3763107

[41] Tianqing Zhu, Gang Li, Wanlei Zhou, and Philip S. Yu. 2017. Differentially Private Data Publishing and Analysis: A Survey. *IEEE Transactions on Knowledge and Data Engineering* 29, 8 (2017), 1619–1638. https://doi.org/10.1109/TKDE.2017.2697856