

Efficiently Maintaining Next Basket Recommendations under Additions and Deletions of Baskets and Items

Benjamin Longxiang Wang
University of Amsterdam & ABN AMRO Bank
wlongxiang1119@gmail.com

Sebastian Schelter
AIRLab, University of Amsterdam
s.schelter@uva.nl

ABSTRACT

Recommender systems play an important role in helping people find information and make decisions in today’s increasingly digitalized societies. However, the wide adoption of such machine learning applications also causes concerns in terms of data privacy. These concerns are addressed by the recent “General Data Protection Regulation” (GDPR) in Europe, which requires companies to delete personal user data upon request, when users enforce their “right to be forgotten”. Many researchers argue that this deletion obligation does not only apply to the data stored in primary data stores such as relational databases, but also requires an update of machine learning models whose trainingset included the personal data to delete. As a consequence, the academic community started to investigate how to unlearn user data from trained machine learning models in an efficient and timely manner.

We explore this direction in the context of a sequential recommendation task called Next Basket Recommendation (NBR), where the goal is to recommend a set of items based on a user’s purchase history. We design efficient algorithms for incrementally and decrementally updating a state-of-the-art next basket recommendation model in response to additions and deletions of user baskets and items. Furthermore, we discuss an efficient, data-parallel implementation of our method in the Spark Structured Streaming system.

We evaluate our implementation on a variety of real-world datasets, where we investigate the impact of our update techniques on several ranking metrics and measure the time to perform model updates. Our results show that our method provides constant update time efficiency with respect to an additional user basket in the incremental case, and linear efficiency in the decremental case where we delete existing baskets. With modest computational resources, we are able to update models with a latency of around 0.2 milliseconds regardless of the history size in the incremental case, and less than one millisecond in the decremental case.

Reference Format:

Benjamin Longxiang Wang and Sebastian Schelter. 2021. Efficiently Maintaining Next Basket Recommendations under Additions and Deletions of Baskets and Items. In *4th Workshop on Online Recommender Systems and User Modeling (ORSUM 2021), in conjunction with the 15th ACM Conference on Recommender Systems, October 2nd, 2021, Amsterdam, The Netherlands (Remote)*.

1 INTRODUCTION

Recommender systems (RS) are ubiquitous in our digitalized societies with applications in e-commerce, social networks and streaming services [32]. However, the wide adoption of such machine

learning applications has been causing concerns in terms of fairness, accountability, privacy, and transparency in the general public and research community alike [5, 30, 38]. It is widely recognised that we need to tackle such responsibility issues at the intersection of machine learning and data management [19, 27, 29, 31].

The right to be forgotten. In this work, we focus on one specific aspect of privacy stipulated by the *General Data Protection Regulation* (GDPR), namely, “the right to be forgotten” [7], which states that the “*data subject shall have the right to [...] the erasure of personal data concerning him or her without undue delay [...] where the data subject withdraws consent*”. When users can enforce their right to be forgotten, companies must delete their personal user data. Many researchers argue that this deletion obligation does not only apply to the data stored in primary data stores such as relational databases, but also requires an update of machine learning models whose trainingset included the personal data to delete. Legal scholars even point out the continual use of AI systems trained on deleted data could be considered illegal under certain interpretations [34]. In some cases, the retained lossy representations in ML models can even be restored by adversarial parties [26].

Towards timely and efficient unlearning. A naive way to realize data deletion is to retrain the model from scratch on the remaining dataset, leaving out the user data to delete. However, this approach is typically expensive in terms of time and computational resources. Furthermore, the GDPR mentions companies should erase data “*without undue delay*” [7] but does not directly specify when deletions should be fully conducted. In practice, data deletion is still a lengthy process. For example, Google could take two to six months to fully delete someone’s data¹; for Facebook, this period is 90 days². We argue that it is an open and important academic question to determine how fast data can be erased from recommendation models, and that *we should design algorithms and systems to empower users to exercise their right to be forgotten as timely as possible*, instead of making them wait for several months.

Efficiently maintaining a TIFU-kNN model for next basket recommendation. In this paper, we explore the direction of efficiently updating a recommendation model in the context of a sequential recommendation task called “Next Basket Recommendation” (NBR, Section 2.1), where the goal is to recommend a set of items based on a user’s purchase history. We focus our efforts on *TIFU-kNN* (Temporal Item Frequency-based User-kNN) [13], a nearest neighbor-based state-of-the-art next basket recommendation model (Section 2.2), which outperforms neural networks in the NBR task, despite its simplicity. We design efficient algorithms for incrementally and decrementally updating a state-of-the-art

ORSUM@ACM RecSys 2021, October 2nd, 2021, Amsterdam, The Netherlands (Remote)
Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://policies.google.com/technologies/retention>

²<https://www.facebook.com/help/250563911970368>

next basket recommendation model in response to additions and deletions of user baskets and items.

We formalize the problem of maintaining a TIFU-kNN model under additions and deletions of baskets and items in Section 3. Next, we design efficient online algorithms that extend TIFU-kNN to learn incrementally as new data arrives (incremental updates Section 4.2) and to forget efficiently as deletion requests come in (decremental updates Section 4.3). We implement our approach in a data-parallel manner in the Apache Spark Structured Streaming system [37] (Section 5). Furthermore, we conduct extensive experiments in Section 6 where we analyze the recommendation performance and update efficiency of our approach, and point out its limitations for an extreme number of deletions in Section 6.3. In summary, this paper provides the following contributions.

- We introduce the problem of incremental/decremental learning for next basket recommendation (Section 3).
- We design efficient online algorithms for incremental and decremental updates of the state-of-the-art NBR model TIFU-kNN (Section 4).
- We discuss a data-parallel implementation of our approach in the Spark Structured Streaming system (Section 5).
- We conduct experiments on three real-world datasets, where we evaluate our approach in terms of predictive performance and update time efficiency. We find that we are able to update the model with constant time efficiency in the incremental case and linear time efficiency in the decremental case with millisecond latency (Section 6).

2 BACKGROUND

2.1 Next Basket Recommendation

We focus on a specific recommendation task called Next Basket Recommendation (NBR), where we aim to recommend a set of items to users based on their historical baskets [25]. NBR is of great interest to the e-commerce and retail industry, where we want to recommend a set of items to fill a user's shopping basket. A basket \mathbf{b} is a set of items, i.e., $\mathbf{b} = \{i_1, i_2, \dots, i_j, \dots, i_{|\mathbf{b}|}\}$, where $i_j \in \mathcal{I}$, and where \mathcal{I} denotes the universe of all items. For a given user, we have access to a sequence of n historical baskets (in increasing chronological order, such that more recent items are at the tail) denoted as $\mathcal{H} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_i, \dots, \mathbf{b}_n]$, where $\mathbf{b}_i \subset \mathcal{I}$. The goal of NBR is then to find a model which takes the historical baskets \mathcal{H} as input and predicts the next basket \mathbf{b}_{n+1} as recommendation.

2.2 TIFU-kNN

Our work centers around adapting the state-of-the-art next basket recommendation model TIFU-kNN [13]. This method builds on the user-based collaborative filtering approach [17]. A target user is first modeled by an embedding vector, and then recommendations are made based on a weighted combination of the aggregated representations of similar users and the target user's own vector embedding. TIFU-kNN calculates user vector representations during the training phase leveraging repeated purchase patterns [1, 2] commonly observed in e-commerce, under the assumption that recent purchases have more predictive power than older ones.

User vector representation. A user vector is calculated to represent a user's past purchase patterns, with a time-decayed window to assign a higher weight to recent items. However, a single time decayed weight is not flexible enough to model another property of temporal dynamics: consecutive steps have small changes while steps far from each other have large changes. Therefore, TIFU-kNN applies a hierarchical structure where user baskets are first divided into groups, baskets within a group are aggregated with time decayed weights to obtain a group vector, and afterwards, these group vectors are aggregated again with time decayed weights to obtain a final user vector. In detail, the algorithm applies the following steps:

Step 1: Multi-hot encoding. Without loss of generality, we denote the historical baskets of our target user as $\mathcal{H} = [\mathbf{b}_1, \dots, \mathbf{b}_i, \dots, \mathbf{b}_n]$ where $\mathbf{b}_i \subset \mathcal{I}$ and $|\mathcal{H}| = n$. Each basket \mathbf{b}_i is converted into a multi-hot encoded vector whose dimension is equal to $|\mathcal{I}|$ with ones in dimensions for items present in the basket, and zeros in the remaining positions. For example, suppose we have a corpus containing four items, i.e., $\mathcal{I} = \{1, 2, 3, 4\}$, and a user's history with two baskets $\mathcal{H} = [\{1, 4\}, \{1, 2, 3\}]$. The corresponding multi-hot vector would then look as follows.

$$[\{1, 4\}, \{1, 2, 3\}] \xrightarrow[\{1, 2, 3, 4\}]{\text{multi-hot encoding}} [[1, 0, 0, 1], [1, 1, 1, 0]]$$

Step 2: Group vector generation. To model the user's behavior shift over time, TIFU-kNN utilizes a hierarchical grouping mechanism to capture the dynamics. Given a fixed group size m as a hyper-parameter, the number of groups is denoted as $k = \lceil \frac{n}{m} \rceil$ where $\lceil \cdot \rceil$ is the ceiling function. The group size dictates the number of baskets in a group. Given a temporal decay rate $0 < r_b \leq 1$ and baskets partitioned into groups of equal length (except for the last group which contains the remaining baskets in case n is not fully divisible by m), a group division example looks as follows.

$$\mathcal{H} = [\overbrace{\mathbf{b}_1, \dots, \mathbf{b}_m}^{\mathbf{v}_{g1}}, \overbrace{\mathbf{b}_{m+1}, \dots, \mathbf{b}_{2m}}^{\mathbf{v}_{g2}}, \dots, \overbrace{\mathbf{b}_{n-1}, \mathbf{b}_n}^{\mathbf{v}_{gk}}]$$

Each group is represented by a group vector \mathbf{v}_g which is calculated by a time-decay weighted average, where the i -th vector within a group is multiplied by a time-decayed weight r_b^{m-i} , where \mathbf{v}_{bi} is a multi-hot encoded vector. Note that the indicator for the group index is left out for brevity.

$$\mathbf{v}_g = \frac{r_b^{m-1}\mathbf{v}_{b1} + r_b^{m-2}\mathbf{v}_{b2} + \dots + \mathbf{v}_{bm}}{m} = \frac{\sum_{i=1}^m r_b^{m-i}\mathbf{v}_{bi}}{m} \quad (1)$$

Step 3: User vector generation. Similar to how a group vector is calculated from basket vectors, TIFU-kNN calculates the final target user vector \mathbf{v}_u representation based on a time-decay weighted average of the group vectors $\mathcal{H} = [\mathbf{v}_{g1}, \mathbf{v}_{g2}, \dots, \mathbf{v}_{gk}]$. Given the group decay rate $0 < r_g \leq 1$, we have

$$\mathbf{v}_u = \frac{r_g^{k-1}\mathbf{v}_{g1} + r_g^{k-2}\mathbf{v}_{g2} + \dots + \mathbf{v}_{gk}}{k} = \frac{\sum_{i=1}^k r_g^{k-i}\mathbf{v}_{gi}}{k} \quad (2)$$

Personalized Collaborative Filtering Recommendation. After obtaining the vector representations of all users, TIFU-kNN generates recommendations inspired by collaborative filtering (CF) [17]. Given a predefined distance measure, it finds the top- k most similar users for each target user. The final prediction combines two parts:

(i) the personalized component, e.g., the user vector representation denoted by $\mathbf{u}_{\text{target}}$, corresponding to the user's temporal repeated purchase pattern; (ii) the collaborative filtering component, which is calculated as the mean of the nearest neighbors' vector representations, denoted by $\mathbf{u}_{\text{neighbors}}$.

The final prediction vector is calculated via a linear combination $\mathbf{p} = \alpha \mathbf{u}_{\text{target}} + (1 - \alpha) \mathbf{u}_{\text{neighbors}}$ of both components, where α is a hyper-parameter used to adjust the weight of each component. Note that the numeric values in the prediction vector \mathbf{p} indicate the strength of preference of the user towards a certain item \mathbf{p} , and the highest scored items are typically leveraged for recommendation.

3 PROBLEM STATEMENT

As outlined in the introduction, we focus on two aspects of the maintenance of TIFU-kNN models: (i) how to incrementally learn user representations in response to additional baskets (*incremental updates*); (ii) how to efficiently delete user data from the learned user representations (*decremental updates*).

Let $\mathcal{D} = \{\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_i, \dots, \mathcal{H}_M\}$ denote a dataset containing historical baskets for M users, and the corresponding user vectors are $\mathcal{U} = \{\mathbf{v}_{u1}, \mathbf{v}_{u2}, \dots, \mathbf{v}_{ui}, \dots, \mathbf{v}_{uM}\}$ where \mathbf{v}_{ui} is a real-valued vector of size $|\mathcal{I}|$. Our goal is to maintain the user vector of a TIFU-kNN model. As each user vector is modeled independently (e.g., only calculated from the baskets of the particular user), we consider only a single user in the following.

Constant time incremental updates. Given the historical baskets $\mathcal{H} = [\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_i, \dots, \mathbf{b}_n]$ and a newly arriving basket \mathbf{b}_{n+1} for user \mathbf{u} , we can always recalculate the corresponding user vector from scratch from the complete history. For incremental updates, we want to improve the asymptotic time complexity from $O(|\mathcal{H}|)$ for retraining from scratch to a constant effort $O(1)$ per additional basket. Given the history \mathcal{H} and the current user vector \mathbf{v}_u , we are interested in an efficient approach $f_{\text{incr}}(\mathbf{v}_u, \mathbf{b}_{n+1}) \rightarrow \mathbf{v}'_u$ to update the current user vector \mathbf{v}_u to \mathbf{v}'_u by incorporating \mathbf{b}_{n+1} .

Efficient decremental updates. In the decremental case, given the current user vector \mathbf{v}_u and a deletion request r (identifying a basket or item to be deleted), we aim to efficiently delete r from the corresponding user vector \mathbf{v}_u . We want to improve the asymptotic complexity from $O(|\mathcal{H}|)$ for retraining from scratch to $O(|\mathcal{H}| - p)$ where p is the position of the basket where a deletion is applied. Formally, we look for a decremental update function $f_{\text{decr}}(\mathcal{H}, \mathbf{v}_u, r)$ which takes three inputs: the historical baskets \mathcal{H} , the existing user vector \mathbf{v}_u and the deletion request r , and outputs an updated user vector \mathbf{v}'_u . We want f_{decr} to satisfy $\forall \mathbf{u}, \forall r : f_{\text{decr}}(\mathcal{H}, \mathbf{v}_u, r) = \mathbf{v}'_u$, e.g., we want the result of a decremental update to be the same as computing the user vector after deletion \mathbf{v}'_u by running TIFU-kNN from scratch on $\mathcal{H} \setminus r$.

4 APPROACH

As mentioned in Section 3, we compare our approach to the baseline approach of retraining the whole model from scratch whenever there is an incremental or decremental update. This baseline has an asymptotic complexity of $O(|\mathcal{H}|)$ per update, be it incremental or decremental. This baseline has several disadvantages: (i) the (re)training phase is expensive in terms of time and resources, especially if it has to be executed repeatedly for a small number of

data deletion requests (ii) the user has to potentially wait long for their data to be actually deleted from the model if the retraining is only executed periodically.

To overcome these challenges, we design online algorithms which update the model with low latency in response to additional data or deletion requests concerning baskets or items. We first discuss general rules for maintaining the decaying average of a series in Section 4.1 and then apply them to derive our approach to incremental updates in Section 4.2 and decremental updates in Section 4.3 under different scenarios.

4.1 Maintaining the Decaying Average of a Series

Calculating decaying averages of vector series is at the heart of TIFU-kNN (Section 2.2) as it is the core computation for both the group vector (Equation 1) and user vector (Equation 2). Without loss of generality, we discuss this problem with a series of real numbers. Let $\mathcal{S} = [x_1, \dots, x_n]$ denote a series of real numbers, where the next item is x_{n+1} . Let \bar{x}_n denote the decaying average of the first n numbers. Given $0 < r \leq 1$, the decaying average of the series is $\bar{x}_n = \frac{1}{n} \sum_{i=1}^n r^{n-i} x_i$. We examine how the decaying average changes in response to incremental and decremental updates. Note that it is straightforward to extend our approach to a series of vectors of real numbers by replacing the enclosed elements by vectors of the same length, which essentially leads to the group vector (Equation 1) and user vector (Equation 2) definitions of TIFU-kNN. We now look at three scenarios: incremental updates, decremental updates, and in-place updates.

Incrementally adding a number. With a newly arriving item x_{n+1} , we update the corresponding decaying average as follows.

$$\begin{aligned} \bar{x}_{n+1} &= \frac{1}{n+1} \sum_{i=1}^n r^{n+1-i} x_i \\ &= \frac{r^n x_1 + r^{n-1} x_2 + \dots + r x_n + x_{n+1}}{n+1} \\ &= \frac{r n \bar{x}_n + x_{n+1}}{n+1} \end{aligned} \quad (3)$$

With this update rule, we can efficiently update the decaying average in constant time using only the current decaying average \bar{x}_n , the count of items n , the additional item x_{n+1} and the decay rate r .

Decrementally deleting a number. If an existing number x_i needs to be deleted from the series, we would like to find an efficient way to update the decaying average without resorting to a full recalculation. Before deletion, the decaying average is:

$$\bar{x}_n = \frac{\overbrace{r^{n-1} x_1 + \dots + r^{n-(i-1)} x_{i-1}}^{\text{preceding deletion}} + \overbrace{r^{n-i} x_i}^{\text{deletion}} + \overbrace{r^{n-(i+1)} x_{i+1} + \dots + x_n}^{\text{succeeding deletion}}}{n}$$

After deletion, the updated decaying average becomes:

$$\bar{x}'_{n-1} = \frac{\overbrace{r^{n-2} x_1 + \dots + r^{n-1-(i-1)} x_{i-1}}^{\text{preceding deletion}} + \overbrace{r^{n-(i+1)} x_{i+1} + \dots + x_n}^{\text{succeeding deletion}}}{n-1}$$

We rearrange the terms such that we express \bar{x}'_{n-1} using \bar{x}_n :

$$\begin{aligned}\bar{x}'_{n-1} &= \frac{n\bar{x}_n + [r^{n-i}(x_{i+1} - x_i) + \dots + r(x_n - x_{n-1}) - x_n]}{(n-1)r} \\ &= \frac{n\bar{x}_n + \mathcal{D}([x_i, \dots, x_n])^T \mathcal{R}(r, n-i)}{(n-1)r}\end{aligned}\quad (4)$$

where $\mathcal{D}(\cdot)$ denotes a vector of the first order difference of the slice of the series *after* the deleted item: $\mathcal{D}([x_i, \dots, x_n]) = [x_{i+1} - x_i, \dots, x_n - x_{n-1}, -x_n]^T$ and $\mathcal{R}(\cdot)$ denotes a decaying vector: $\mathcal{R}(r, n-i) = [r^{n-i}, \dots, r, 1]^T$. We discuss a few observations:

- The decremental update rule (Equation 4) is mathematically symmetrical to its incremental counterpart (Equation 3), as in both cases we update by incorporating the previous value and the change introduced by the additional or removed item. This view provides a unified perspective to deal with incremental and decremental updates.
- To perform a deletion, it is sufficient to access a slice starting from the deleted item. In the worst case when the deleted item is the very first item x_1 , the slice is equivalent to the full series; in the best case when the deleted item is the very last item, we only need to access a single item. Compared to a full recalculation where (re)access to the full history is needed, our update rule reduces the amortized number of items impacted by a deletion from $|\mathcal{S}|$ to $|\mathcal{S}|/2$, assuming a uniform distribution of deletions.
- Lastly the term $\mathcal{D}(\cdot)^T \mathcal{R}(\cdot)$ is a dot product and can be efficiently computed by vectorization, which further improves the efficiency of decremental updates.

Updating an existing number. Suppose x_i is updated to a new value x'_i , then the new decaying average becomes

$$\begin{aligned}\bar{x}'_n &= \frac{r^{n-1}x_1 + \dots + r^{n-i}x'_i + \dots + x_n}{n} \\ &= \frac{r^{n-1}x_1 + \dots + r^{n-i}x'_i + \dots + x_n + r^{n-i}x_i - r^{n-i}x_i}{n} \\ &= \frac{n\bar{x}_n + r^{n-i}(x'_i - x_i)}{n}\end{aligned}\quad (5)$$

Note that for an in-place update, we only need to update the previous decaying average with a term dependent on the difference introduced by the updated item.

The approaches to update decaying averages for the incremental, decremental, and in-place scenarios serve as a basis for the online algorithms for TIFU-kNN developed in the following subsections.

4.2 Incrementally Adding New Baskets

Given an existing user vector \mathbf{v}_u and an additional basket \mathbf{b}_{n+1} , we discuss an algorithm to update the user vector efficiently based on the rules discussed in Section 4.1. Note that it is sufficient just to study the problem for a single user as all user vectors are computed independently. Due to the grouping mechanism in TIFU-kNN (Section 2), adding a basket to the history will either lead to 1) a new group with a single basket (if the last existing group is fully occupied) or 2) an update of the last existing group with one new basket appended. This is best illustrated with an example. Consider a user with four baskets, group size $m = 2$, hence number of groups $k = \lceil \frac{4}{2} \rceil = 2$ which leads to two groups with each group containing

two baskets. Now we add two new baskets one by one. The group composition evolves as follows.

$$\begin{aligned}\mathcal{H} &= [\overbrace{\mathbf{v}_{b1}, \mathbf{v}_{b2}}^{\mathbf{v}_{g1}}, \overbrace{\mathbf{v}_{b3}, \mathbf{v}_{b4}}^{\mathbf{v}_{g2}}] \rightarrow \\ \mathcal{H}' &= [\overbrace{\mathbf{v}_{b1}, \mathbf{v}_{b2}}^{\mathbf{v}_{g1}}, \overbrace{\mathbf{v}_{b3}, \mathbf{v}_{b4}}^{\mathbf{v}_{g2}}, \overbrace{\mathbf{v}_{b5}}^{\mathbf{v}_{g3}}] \rightarrow \\ \mathcal{H}'' &= [\overbrace{\mathbf{v}_{b1}, \mathbf{v}_{b2}}^{\mathbf{v}_{g1}}, \overbrace{\mathbf{v}_{b3}, \mathbf{v}_{b4}}^{\mathbf{v}_{g2}}, \overbrace{\mathbf{v}_{b5}, \mathbf{v}_{b6}}^{\mathbf{v}'_{g3}}]\end{aligned}\quad (6)$$

A group with a single basket is created after adding a new basket to a fully occupied group, and afterwards, the group is updated with the second additional basket. Our update mechanisms differentiate these two scenarios:

Scenario 1: Adding a group with a single basket. When transitioning from \mathcal{H} to \mathcal{H}' (Equation 6), we add an additional group vector (with a single basket) to update the user vector while the existing group vectors remain unchanged. Recall that we work with a group size m , a number of baskets n and a number of groups $k = \lceil \frac{n}{m} \rceil$. Since the new group only contains the newly added basket, the new group vector is equivalent to the new basket vector $\mathbf{v}_{g(k+1)} = \mathbf{v}_{b(n+1)}$. Given the additional group vector $\mathbf{v}_{g(k+1)}$, the existing user vector \mathbf{v}_u as defined in Equation 2, and the group decay rate r_g , we obtain the updated user vector \mathbf{v}'_u by applying the incremental rule derived in Equation 3.

$$\begin{aligned}\mathbf{v}'_u &= \frac{r_g^k \mathbf{v}_{g1} + r_g^{k-1} \mathbf{v}_{g2} + \dots + r_g \mathbf{v}_{gk} + \mathbf{v}_{g(k+1)}}{k+1} \\ &= \frac{kr_g \mathbf{v}_u + \mathbf{v}_{b(n+1)}}{k+1}\end{aligned}\quad (7)$$

Note that we can efficiently update the user vector in constant time regardless of the size of the user's history with access to the current user vector \mathbf{v}_u and the number of groups k .

Scenario 2: Updating an existing group. When transitioning from \mathcal{H}' to \mathcal{H}'' (Equation 6), we have $\tau < m$ where τ is the number of baskets in the last group and m is the group size. In this scenario, the last group vector \mathbf{v}_{gk} needs to be updated to take the added basket into account. Given the number of existing baskets in the last group τ , the basket decay rate r_b , the additional basket denoted as $\mathbf{v}_{b(\tau+1)}$, we update the group vector \mathbf{v}_{gk} by applying the incremental rule derived in Equation 3.

$$\begin{aligned}\mathbf{v}'_{gk} &= \frac{r_b^\tau \mathbf{v}_{b1} + r_b^{\tau-1} \mathbf{v}_{b2} + \dots + r_b \mathbf{v}_{b\tau} + \mathbf{v}_{b(\tau+1)}}{\tau+1} \\ &= \frac{\tau r_b \mathbf{v}_{gk} + \mathbf{v}_{b(n+1)}}{\tau+1}\end{aligned}\quad (8)$$

Analogous to the previous scenario, we only need access to the current group vector \mathbf{v}_{gk} and the number of baskets in the last group τ in order to obtain the updated last group vector efficiently in constant time. Finally, with the updated group vector \mathbf{v}'_{gk} , the target user vector after addition \mathbf{v}'_u is recomputed with access to the last user vector \mathbf{v}_u , the last group vector \mathbf{v}_{gk} and the number

of groups k by applying the in-place update rule in Equation 5.

$$\begin{aligned} \mathbf{v}'_u &= \frac{r_g^{k-1} \mathbf{v}_{g1} + r_g^{k-2} \mathbf{v}_{g2} + \dots + r_g \mathbf{v}_{g(k-1)} + \mathbf{v}'_{gk}}{k} \\ &= \mathbf{v}_u + \frac{\mathbf{v}'_{gk} - \mathbf{v}_{gk}}{k} \end{aligned} \quad (9)$$

4.3 Decrementally Removing Existing Baskets and Items

First, we discuss how a user vector will change if a user requests to delete a specific basket from her basket history. In practical applications, this could happen when a customer on an e-commerce website requests to delete a shopping session for privacy reasons or when a customer made a return. It is sufficient just to look at one basket because we can handle the removal of multiple baskets by applying our algorithm repeatedly. Subsequently, we discuss how to remove a single item from a basket instead of the basket as a whole, which corresponds to a scenario where a customer requests to remove certain sensitive items from a previous shopping cart. Again, we can handle these cases via the update rules discussed in Section 4.1.

Varying group size. Before diving into the algorithmic details, we first discuss a relaxation of the original algorithm to which we refer as *varying group size*. Given a user's historical baskets, if we were to maintain a fixed group size, removing a basket at any location would cause cascading changes in terms of group composition for any subsequent groups following the deleted basket. For users with a large number of baskets, it quickly becomes impractical to regroup and recalculate all group vectors, especially when the basket under deletion is at the beginning. As a consequence, we propose an alternative to avoid expensive regrouping while still maintaining the temporal patterns with recency effects. We allow for a *varying group size*, and thereby restrict the regrouping impact to the enclosing group of the to-be-deleted basket. We demonstrate how the group composition changes under a *varying group size* as follows.

$$\begin{aligned} & \underbrace{[\mathbf{v}_{b1}, \mathbf{v}_{b2}, \mathbf{v}_{b3}, \mathbf{v}_{b4}, \mathbf{v}_{b5}, \mathbf{v}_{b6}, \mathbf{v}_{b7}, \mathbf{v}_{b8}]}_{\substack{\mathbf{v}_{g1} & \mathbf{v}_{g2} & \mathbf{v}_{g3} & \mathbf{v}_{g4}}} \xrightarrow{-\mathbf{v}_{b3}} \\ & \underbrace{[\mathbf{v}_{b1}, \mathbf{v}_{b2}, \mathbf{v}_{b4}, \mathbf{v}_{b5}, \mathbf{v}_{b6}, \mathbf{v}_{b7}, \mathbf{v}_{b8}]}_{\substack{\mathbf{v}_{g1} & \mathbf{v}'_{g2} & \mathbf{v}_{g3} & \mathbf{v}_{g4}}} \xrightarrow{-\mathbf{v}_{b6}} \\ & \underbrace{[\mathbf{v}_{b1}, \mathbf{v}_{b2}, \mathbf{v}_{b4}, \mathbf{v}_{b5}, \mathbf{v}_{b7}, \mathbf{v}_{b8}]}_{\substack{\mathbf{v}_{g1} & \mathbf{v}'_{g2} & \mathbf{v}'_{g3} & \mathbf{v}_{g4}}} \end{aligned}$$

where we show that after removing baskets for \mathbf{v}_{b3} and \mathbf{v}_{b6} , the impacted groups are limited at minimal to \mathbf{v}'_{g2} and \mathbf{v}'_{g3} respectively. We differentiate three different deletion scenarios:

Scenario 1: Deleting a basket from an existing group. If the enclosing group contains more than one basket, deleting a basket from it leads to an updated group vector. Given the to-be-deleted basket vector denoted as \mathbf{v}_{bi} , the current group vector \mathbf{v}_{gj} , the basket decay rate r_b and the number of baskets in the enclosing group τ , we can apply the decremental rule (Equation 4) to obtain

the updated group vector,

$$\mathbf{v}'_{gj} = \frac{\tau \mathbf{v}_{gj} + \mathcal{D}([\mathbf{v}_{bi}, \dots, \mathbf{v}_{b\tau}])^T \mathcal{R}(r_b, \tau - i)}{(\tau - 1)r_b} \quad (10)$$

where we have $\mathcal{D}([\mathbf{v}_{bi}, \dots, \mathbf{v}_{b\tau}]) = [\mathbf{v}_{b(i+1)} - \mathbf{v}_{bi}, \dots, \mathbf{v}_{b\tau} - \mathbf{v}_{b(\tau-1)}, -\mathbf{v}_{b\tau}]^T$ and $\mathcal{R}(r_b, \tau - i) = [r_b^{\tau-i}, \dots, r_b, 1]^T$. Similar to the observations for the decremental rule (Equation 4), we highlight one important property of Equation 10: we only need to access a slice starting from the deleted basket instead of the full history to perform a deletion. Compared to a full retraining, we reduce the amortized number of baskets impacted by a deletion from $|\mathcal{H}|$ to $|\mathcal{H}|/2$ (assuming a uniform distribution of deletes). Finally, following the in-place update rule (Equation 5), we compute the updated user vector \mathbf{v}'_u from the updated group vector.

$$\mathbf{v}'_u = \mathbf{v}_u + \frac{r_g^{k-i} (\mathbf{v}'_{gj} - \mathbf{v}_{gj})}{k} \quad (11)$$

Scenario 2: Deleting a single-basket group. We consider the scenario where a group with a single basket is removed, and the corresponding group vector \mathbf{v}_{gi} will vanish, so we only need to update the user vector \mathbf{v}_u . Given the number of groups k , the vanished group vector \mathbf{v}_{gi} , the group decay rate r_g , the current user vector \mathbf{v}_u , we update the user vector by applying the decremental rule (Equation 4) again, where we have $\mathcal{D}([\mathbf{v}_{gi}, \dots, \mathbf{v}_{gk}]) = [\mathbf{v}_{g(i+1)} - \mathbf{v}_{gi}, \dots, \mathbf{v}_{gk} - \mathbf{v}_{g(k-1)}, -\mathbf{v}_{gk}]^T$ and $\mathcal{R}(r_g, k - i) = [r_g^{k-i}, \dots, r_g, 1]^T$.

$$\mathbf{v}'_u = \frac{k \mathbf{v}_u + \mathcal{D}([\mathbf{v}_{gi}, \dots, \mathbf{v}_{gk}])^T \mathcal{R}(r_g, k - i)}{(k - 1)r_g} \quad (12)$$

Scenario 3: Removing a single item from a basket. Instead of requesting to delete a whole basket, one can also delete an item within a basket, and update the corresponding user vector. Deleting an item for a target basket vector \mathbf{v}_{bi} will either 1) result in an updated basket vector by changing the corresponding item dimension from one to zero in the basket vector or 2) cause the basket vector to vanish if it only contained the item to remove. In the first case, given the updated basket vector \mathbf{v}'_{bi} , the current group vector \mathbf{v}_g , the number of baskets in the group τ and decay rate r_b , we update the group vector by applying the in-place update rule (Equation 5) again.

$$\mathbf{v}'_g = \mathbf{v}_g + \frac{r_b^{\tau-i} (\mathbf{v}'_{bi} - \mathbf{v}_{bi})}{\tau} \quad (13)$$

Now that we have \mathbf{v}'_g , we can apply Equation 11 to obtain the updated user vector. In the second case where the target basket vanishes, we fall back to the previous scenario 1 or 2 depending on if the vanished basket also causes a vanished group.

5 IMPLEMENTATION

In the following, we discuss a data-parallel implementation of our approach in a popular stream processing system. We briefly show the system overview, then discuss how to leverage the parallel nature of TIFU-kNN, and how to combine incremental and decremental updates to fulfill low-latency requirements. Our implementation

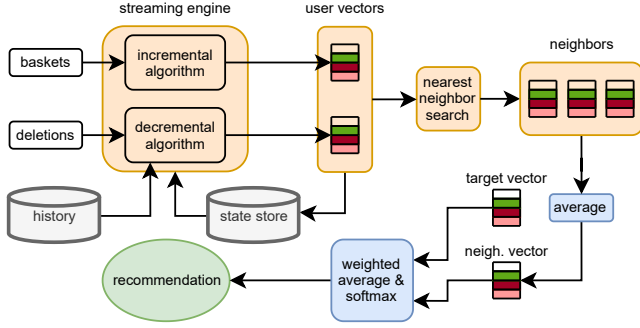


Figure 1: System overview. The streaming engine consists of the incremental algorithm and decremental algorithm which process the newly added baskets and execute deletion requests respectively; the state store maintains the user vectors and provides IO to the streaming engine; next is the nearest neighbor search component which calculates the neighborhood of the target user; finally, the target user vector and neighborhood vector are combined to make recommendations.

is conducted in *Scala* in Apache Spark’s Structured Streaming abstraction [37]. Spark Structured Streaming allows us to express a streaming computation as a state update based on data changes. The Spark engine takes care of running it incrementally and continuously, and of updating the result as new streaming data continues to arrive.

System overview. The system (Figure 1) consists of two stages: 1) the training phase where user baskets and deletion requests are processed and user vectors are produced or updated; 2) the prediction phase where the target user vector and its neighborhood vectors are combined to make recommendations.

Parallelism. Since each user vector is calculated independently, we can leverage this fact to achieve parallelism on the user level, which in turn makes our implementation scale to a large number of users for practical usage. We leverage the `mapGroupsWithState` API³ in Apache Spark for this purpose, because this operation allows us to apply user-defined code on grouped datasets to update user-defined state. We treat each user’s basket stream as a dataset grouped by a unique user identifier, and the user vectors are stored in a state store.

Joint incremental and decremental state updates. The incremental algorithm and the decremental algorithm in the streaming engine are jointly responsible for updating the user vectors while the former processes incoming user data and the latter handles deletion requests. We show a (simplified) version in Algorithm 1 where the f_{incr} (Section 4.2) updates the user vector state without accessing the history; the f_{decr} procedure (Section 4.3) receives the deletion request, the current state and the user’s historical baskets and returns the updated user vector state. In our implementation, the users’ historical baskets are stored in memory indexed with

user identifiers for fast retrieval. Our code is available under an open source license at https://github.com/0xeef/amnesiac_recsys.

Algorithm 1 Joint incremental and decremental state updates

```

function UPDATESTATE(currentState, input)
  Input: the current user vector and the corresponding update
  (either a deletion request or an additional user basket)
  Output: the updated state
  Data: the user historical baskets  $\mathcal{H}$  (required by the deletion
  procedure)
  if input.isDeletion then
    newState  $\leftarrow f_{decr}(\text{currentState}, \text{input}, \mathcal{H})$ 
  else
    newState  $\leftarrow f_{incr}(\text{currentState}, \text{input})$ 
  return newState
  
```

6 EVALUATION

We first perform an evaluation using offline ranking metrics, namely, recall and *Normalized Discounted Accumulative Gain* (NDCG) on real-life datasets to show that our approach achieves a comparable performance to our baseline approach of retraining the model from scratch (Section 2). Next, we examine the efficiency and scalability of our approach by evaluating the time to update a user vector or model.

6.1 Predictive Performance under Incremental/Decremental Updates

Datasets. The state-of-the-art performance of TIFU-kNN has already been shown by the original authors in [13]. We only adapt this algorithm and show that our incremental and decremental algorithms do not introduce major performance regressions while at the same time providing for efficient incremental learning and data deletion capabilities. We base our experiments on the publicly available datasets (TaFeng⁴, Instacart⁵ and ValuedShopper⁶) and the tuned hyper-parameters (listed in Table 1) used for experimentation in the original TIFU-kNN paper [13]. The basic statistics for the datasets are shown in Table 1. These datasets contain time-stamped transactions in which items bought in the same order are treated as a basket.

Experimental setup. For each user, the last basket is held out as a test basket for evaluation. We report the average over all users for the metrics. The baseline, incremental and decremental experiments are implemented according to Section 2, Section 4.2 and Section 4.3 respectively. Recall that the baseline triggers a full retraining for any incremental or decremental updates. For the decremental experiment, as it is difficult to estimate a realistic number of deletion requests, we opt for the setup according to [28]: approximately 1 out of 10,000 users might issue deletion requests in industrial scenarios. Due to the limited number of users in our datasets (Table 1), we opt for a more extreme estimate of 1 out of 1,000 users, which

³<https://spark.apache.org/docs/latest/structured-streaming-programming-guide.html>

⁴<https://www.kaggle.com/chiranjivdas09/ta-feng-grocery-dataset>

⁵<https://www.kaggle.com/c/instacart-market-basket-analysis/data>

⁶<https://www.kaggle.com/c/acquire-valued-shoppers-challenge/data>

Table 1: Dataset basic statistics and hyper-parameters (from left to right: group size m , decay rates r_b and r_g , the number of neighbors k , the combining weight of the target user vector and neighborhood component α).

Data	#users	#items	#baskets	avg basket size	avg #baskets/user	hyper-parameters
TaFeng	13949	11997	79423	6.2	5.7	[7, 0.9, 0.7, 300, 0.7]
Instacart	19935	7999	158933	8.9	8.0	[3, 0.9, 0.7, 900, 0.9]
ValuedShopper	10000	7874	568573	9.1	56.9	[7, 1.0, 0.6, 300, 0.7]

Table 2: Predictive performance under incremental and decremental updates compared to the baseline approach of retraining the model from scratch. Our online algorithms provide equal predictive performance.

Dataset	metric	baseline	incr.	decr.
TaFeng	Recall@10	0.1298	0.1298	0.1297
	NDCG@10	0.0847	0.0847	0.0847
	Recall@20	0.1859	0.1859	0.1856
	NDCG@20	0.0816	0.0816	0.0817
Instacart	Recall@10	0.1709	0.1709	0.1710
	NDCG@10	0.5789	0.5789	0.5788
	Recall@20	0.2291	0.2291	0.2291
	NDCG@20	0.5037	0.5037	0.5035
ValuedShopper	Recall@10	0.1436	0.1436	0.1437
	NDCG@10	0.4058	0.4058	0.4060
	Recall@20	0.2133	0.2133	0.2131
	NDCG@20	0.3465	0.3465	0.3465

is a more difficult experimental setting. We select these users at random and delete 10% of baskets for those users.

Results and discussion. As shown in Table 2, the incremental updates produce exactly the same results and therefore performance metrics as the baseline, which confirms our analysis in Section 4.2. We expect this behavior, as the incremental updates compute the exact user vectors or models given by the baseline approach. We also observe that the decremental algorithm has no significant negative impact on recommendation performance in terms of NDCG and recall metrics. This observation is confirmed by a t-test on metrics reported for all three datasets. As some might see, decremental updates sometimes even boost the performance metric slightly such as Recall@10 for Instacart and NDDG@10 for ValuedShopper. This is possible because decremental updates could remove noise in historical data hence improve predictive performance for certain users.

6.2 Update Efficiency

Next, we focus on the runtime performance and scalability of our approach. We demonstrate that we are able to update the model with low latency in constant time per update in the incremental case, and in linear time per update in the decremental case per update.

Experimental setup. Without loss of generality, we assume that we are working with a single user and a vocabulary of one single item. Thus, the basket history looks like $[\{1\}, \{1\}, \dots]$. This simplified view enables us to gain clearer insights into the efficiency of our approach. This setting applies throughout this section.

Basket additions. The baseline algorithm needs to access the full history for each update, whereas in our incremental algorithm we only need to access the state and the newly added item or basket once, which can be done in constant time, independent of the history size. In this experiment, we follow the setup where we sequentially add baskets and record the update time along the way. As shown in Figure 2a, our incremental algorithm achieves constant update time efficiency, i.e., the time it takes to retrain does not depend on the size of history, whereas the baseline approach suffers from a linear increase in update time as the number of baskets grows.

Basket deletions. The update time for a deletion depends on the location of the deleted basket. The more recent (closer to the end) it is, the less time it takes to retrain the model. To better illustrate this, we investigate three situations: 1) removing baskets one by one starting from the end; 2) removing baskets one by one starting from the very beginning; 3) removing baskets one by one from any random location. In situation (1), the expected update time should be close to constant. As we remove baskets from the end of the historical baskets, the number of subsequent baskets it needs to access to update the user vector is constant. As we can see in Figure 2b, the experiment result validates our assumption though it shows a slight downward trend, which we attribute to regularly triggered maintenance operations in Spark. In situation (2), the expected update time is approximately linear, as we need to access the whole history every time a deletion occurs from the beginning of the history. In situation (3), which is a more realistic scenario, the update time lies in between the previous two situations. In summary, our approach is on par with the baseline for worst-case scenarios and consistently outperforms it in other cases.

6.3 Limitations of Repeated Deletions

In the following, we discuss a limitation of our approach, based on a potential numerical instability issue that could occur under extreme conditions (a large number of repeated deletions for a specific user). We theoretically analyze the numerical properties of our decremental algorithm, and empirically a potential issue in an extreme case scenario.

We can rewrite the user vector update rule (Equation 12) as $\mathbf{u}' = \alpha \mathbf{u} + \mathbf{C}$ where $\alpha = \frac{k}{(k-1)r_g} > \frac{1}{r_g} > 1$. Suppose the initial user vector is $\mathbf{u}_0 = \hat{\mathbf{u}}_0 + \epsilon$, where we use the hat notion $\hat{\mathbf{u}}_i$ to represent the true value, and the ϵ denotes the initial error. After n steps of deletions, the accumulated error becomes $|\mathbf{u}_n - \hat{\mathbf{u}}_n| = \epsilon \alpha^n$. Since $\alpha > 1$, the decremental algorithm is numerically unstable. This conclusion implies that there is a limit in terms of how many deletions we can perform for a given user before the error becomes unbearable.




Figure 2: Updates times of the baseline algorithm (retraining from scratch) and our proposed approach, under accumulated incremental (Figure 2a) and decremental (Figure 2b) updates. Our incremental update approach achieves constant time updates regardless of the number of accumulated updates. Our decremental update approach performs consistently better than the baseline: in the best case (deletions from end), it achieves near-constant time efficiency; in the worst case (deletions from start), it performs on par with the baseline; in the canonical case (deletions at random), it shows a performance between the best and worst cases. The overall latency for updates is as low as 1ms for the canonical decremental case.

Experimental setup. In this experiment, we first build up the user vector through incremental updates, then we continuously conduct decremental updates during which we calculate the true value of the user vector using the baseline algorithm. The hyper-parameters are fixed: $m = 2$, $r_g = 0.7$, $r_b = 0.9$.

Results and discussion. As shown in Figure 2c, the error grows exponentially with the number of decremental updates, which is confirmed by the error fitting curve. While this issue might be theoretically worrying, we argue that its impact in practice may be negligible, as the overall number of deletions is typically low, and these deletions will be interleaved with incremental updates. The incremental algorithm is a numerically stable algorithm that helps deduce the accumulated error. In our experiment, to incur 1% relative error, it takes 180 continuous deletions, which is arguably an extreme case.

7 RELATED WORK

Decremental learning. Contrary to conventional machine learning tasks, decremental learning attempts to remove the impact of training data from the trained model efficiently. This line of research dates back to some early work on incremental and decremental versions of Support Vector Machines [6, 15], while it has recently gained traction due to the rising concerns for privacy in ML models. Ginart et al. [8] study data deletion mechanisms for variations of k-means clustering. Schelter et al. [26] propose decremental learning methods for a selection of non-iterative ML models such as item-based collaborative filtering and KNN. Recent work from Guo et al. [10] studies this problem for neural networks and shows promising theoretical results albeit apparent scalability bottlenecks due to the requirement of reversing Hessian matrix. Bourtole et al. [3] propose a framework in which training samples are divided into shards and trained in silo to obtain a combination of weaker learners, and later on, combined to form a stronger model. When there is a deletion request concerning specific user data, it's sufficient to update the corresponding impacted weaker learner.

Deep learning for NBR. Deep learning [4, 9, 18, 22, 23] gains momentum in RS given its potential in learning powerful latent

representations of users/items [39]. Its ability to model sequential patterns is especially interesting for NBR. Many new methods based on recurrent neural networks and attention mechanisms [33] have been proposed [11, 12, 20, 21, 24, 35, 36]. However, questions for the effectiveness of deep learning in the RS domain, especially next basket recommendation/sequential recommendation, remain open. Other studies show that simple neighborhood-based methods outperform complex neural models easily [13, 14, 16].

8 CONCLUSION

In this work, we focus on efficient data deletion mechanisms for recommender systems in light of privacy protection and introduce the problem of decremental learning to Next Basket Recommendation. We then present efficient and effective algorithms capable of learning and forgetting with millisecond-level latency. Our approach is evaluated on real-world datasets and demonstrates its effectiveness in terms of ranking metrics (recall and NDCG). Moreover, it also proves to have desirable update time efficiency as shown in our experiments. In addition to that, we provide a fully parallel implementation in Apache Spark to facilitate practical use cases. For theoretical completeness, we also study the limitations under extreme conditions, that is, the decremental algorithm could suffer from numerical instability for a large number of continuous deletions. For future work, we intend to investigate: 1) how to incorporate incremental and decremental algorithms for other more complex recommendation scenarios; 2) numerical stability properties for other ML models under similar settings. In the end, we hope for more advances in this area to realize a future of truly “amnesiac” recommender systems.

Acknowledgments. *This work was supported by Ahold Delhaize. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.*

REFERENCES

- [1] Ashton Anderson, Ravi Kumar, Andrew Tomkins, and Sergei Vassilvitskii. 2014. The dynamics of repeat consumption. In *Proceedings of the 23rd international conference on World wide web*. 419–430.
- [2] Austin R Benson, Ravi Kumar, and Andrew Tomkins. 2016. Modeling user consumption sequences. In *Proceedings of the 25th International Conference on World Wide Web*. 519–529.
- [3] Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. 2019. Machine unlearning. *arXiv preprint arXiv:1912.03817* (2019).
- [4] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165* (2020).
- [5] Robin Burke, Nasim Sonboli, and Aldo Ordóñez-Gauger. 2018. Balanced neighborhoods for multi-sided fairness in recommendation. In *Conference on Fairness, Accountability and Transparency*. PMLR, 202–214.
- [6] Gert Cauwenberghs and Tomaso Poggio. 2001. Incremental and decremental support vector machine learning. *Advances in neural information processing systems* (2001), 409–415.
- [7] EU. 2016. <https://gdpr.eu/article-17-right-to-be-forgotten/>.
- [8] A Gınart, M Guan, G Valiant, and J Zou. 2019. Making AI Forget You: Data Deletion in Machine Learning. *Advances in neural information processing systems* (2019).
- [9] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial networks. *arXiv preprint arXiv:1406.2661* (2014).
- [10] Chuan Guo, Tom Goldstein, Awni Hannun, and Laurens Van Der Maaten. 2020. Certified Data Removal from Machine Learning Models. In *International Conference on Machine Learning*. PMLR, 3832–3842.
- [11] 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*. 843–852.
- [12] Balázs Hidasi, Alexandros Karatzoglou, L. Baltrunas, and D. Tikk. 2016. Session-based Recommendations with Recurrent Neural Networks. *CoRR* abs/1511.06939 (2016).
- [13] Haoji Hu, Xiangnan He, Jinyang Gao, and Zhi-Li Zhang. 2020. Modeling personalized item frequency information for next-basket recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1071–1080.
- [14] Dietmar Jannach and Malte Ludewig. 2017. When recurrent neural networks meet the neighborhood for session-based recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*. 306–310.
- [15] Masayuki Karasuyama and Ichiro Takeuchi. 2009. Multiple incremental decremental learning of support vector machines. *Advances in neural information processing systems* 22 (2009), 907–915.
- [16] Barrie Kersbergen and Sebastian Schelter. [n.d.]. Learnings from a Retail Recommendation System on Billions of Interactions at bol. com. ([n. d.]).
- [17] Joseph A Konstan, Bradley N Miller, David Maltz, Jonathan L Herlocker, Lee R Gordon, and John Riedl. 1997. Grouplens: Applying collaborative filtering to usenet news. *Commun. ACM* 40, 3 (1997), 77–87.
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* 25 (2012), 1097–1105.
- [19] Andreas Kunft, Asterios Katsifodimos, Sebastian Schelter, Tilmann Rabl, and Volker Markl. 2017. Blockjoin: Efficient matrix partitioning through joins. *Proceedings of the VLDB Endowment* 10, 13 (2017), 2061–2072.
- [20] Jing Li, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Tao Lian, and Jun Ma. 2017. Neural attentive session-based recommendation. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 1419–1428.
- [21] Qiao Liu, Yifu Zeng, Refuoe Mokhosi, and Haibin Zhang. 2018. STAMP: short-term attention/memory priority model for session-based recommendation. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1831–1839.
- [22] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fiedel, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.
- [23] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-Shot Text-to-Image Generation. *arXiv:2102.12092* [cs.CV].
- [24] Pengjie Ren, Zhumin Chen, Jing Li, Zhaochun Ren, Jun Ma, and Maarten De Rijke. 2019. Repeatnet: A repeat aware neural recommendation machine for session-based recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 4806–4813.
- [25] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing personalized markov chains for next-basket recommendation. In *Proceedings of the 19th international conference on World wide web*. 811–820.
- [26] S. Schelter. 2020. Amnesia - Machine Learning Models That Can Forget User Data Very Fast. In *CIDR*.
- [27] Sebastian Schelter, Felix Biessmann, Tim Januschowski, David Salinas, Stephan Seufert, and Gyuri Szarvas. 2018. On challenges in machine learning model management. *IEEE Data Engineering Bulletin* (2018).
- [28] Sebastian Schelter, Stefan Grafberger, and Ted Dunning. 2021. HedgeCut: Maintaining Randomised Trees for Low-Latency Machine Unlearning. In *Proceedings of the 2021 International Conference on Management of Data*.
- [29] Sebastian Schelter, Juan Soto, Volker Markl, Douglas Burdick, Berthold Reinwald, and Alexandre Evfimievski. 2015. Efficient sample generation for scalable meta learning. In *2015 IEEE 31st International Conference on Data Engineering*. IEEE, 1191–1202.
- [30] Rashmi Sinha and Kirsten Swearingen. 2002. The Role of Transparency in Recommender Systems. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems* (Minneapolis, Minnesota, USA) (CHI EA '02). Association for Computing Machinery, New York, NY, USA, 830a–831. <https://doi.org/10.1145/506443.506619>
- [31] Julia Stoyanovich, Bill Howe, and HV Jagadish. 2020. Responsible data management. *Proceedings of the VLDB Endowment* 13, 12 (2020), 3474–3488.
- [32] Zhu Sun, Qing Guo, Jie Yang, Hui Fang, Guibing Guo, Jie Zhang, and Robin Burke. 2019. Research commentary on recommendations with side information: A survey and research directions. *Electronic Commerce Research and Applications* 37 (2019), 100879.
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, undefinedukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need (*NIPS'17*). Curran Associates Inc., 6000a–6010.
- [34] Eduard Fosch Villaronga, Peter Kieseberg, and Tiffany Li. 2018. Humans forget, machines remember: Artificial intelligence and the right to be forgotten. *Computer Law & Security Review* 34, 2 (2018), 304–313.
- [35] Haochao Ying, Fuzhen Zhuang, Fuzheng Zhang, Yanchi Liu, Guandong Xu, Xing Xie, Hui Xiong, and Jian Wu. 2018. Sequential recommender system based on hierarchical attention network. In *IJCAI International Joint Conference on Artificial Intelligence*.
- [36] 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*. 729–732.
- [37] Matei Zaharia, Reynold S Xin, Patrick Wendell, Tathagata Das, Michael Armbrust, Ankur Dave, Xiangrui Meng, Josh Rosen, Shivaram Venkataraman, Michael J Franklin, et al. 2016. Apache spark: a unified engine for big data processing. *Commun. ACM* 59, 11 (2016), 56–65.
- [38] Bo Zhang, Na Wang, and Hongxia Jin. 2014. Privacy concerns in online recommender systems: influences of control and user data input. In *10th Symposium On Usable Privacy and Security*. 159–173.
- [39] 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.