

# Are We Really Achieving Better Beyond-Accuracy Performance in Next Basket Recommendation?

Ming Li\*

Yuanna Liu\*

University of Amsterdam  
Amsterdam, The Netherlands  
m.li@uva.nl, y.liu8@uva.nl

Andrew Yates

University of Amsterdam  
Amsterdam, The Netherlands  
a.c.yates@uva.nl

Sami Jullien

AIRLab, University of Amsterdam  
Amsterdam, The Netherlands  
s.jullien@uva.nl

Mohammad Aliannejadi

University of Amsterdam  
Amsterdam, The Netherlands  
m.aliannejadi@uva.nl

Mozhdeh Ariannezhad†

Booking.com  
Amsterdam, The Netherlands  
mozhdeh.ariannezhad@booking.com

Maarten de Rijke

University of Amsterdam  
Amsterdam, The Netherlands  
m.derijke@uva.nl

## ABSTRACT

Next basket recommendation (NBR) is a special type of sequential recommendation that is increasingly receiving attention. So far, most NBR studies have focused on optimizing the accuracy of the recommendation, whereas optimizing for beyond-accuracy metrics, e.g., item fairness and diversity remains largely unexplored. Recent studies into NBR have found a substantial performance difference between recommending repeat items and explore items. Repeat items contribute most of the users' perceived accuracy compared with explore items.

Informed by these findings, we identify a potential "short-cut" to optimize for beyond-accuracy metrics while maintaining high accuracy. To leverage and verify the existence of such short-cuts, we propose a plug-and-play *two-step repetition-exploration* (TReX) framework that treats repeat items and explores items separately, where we design a simple yet highly effective repetition module to ensure high accuracy, while two exploration modules target optimizing *only* beyond-accuracy metrics.

Experiments are performed on two widely-used datasets w.r.t. a range of beyond-accuracy metrics, viz. five fairness metrics and three diversity metrics. Our experimental results show that: (i) we can achieve state-of-the-art performance w.r.t. accuracy via the designed repetition module in TReX; and (ii) the simple TReX framework achieves "better" beyond-accuracy performance than existing sophisticated methods. Prima facie, this appears to be good news: we can achieve high accuracy *and* improved beyond-accuracy metrics at the same time. However, we argue that the real-world value of our algorithmic solution, TReX, is likely to be limited and reflect on the reasonableness of the evaluation setup. We end up challenging existing evaluation paradigms, particularly in the context of beyond-accuracy metrics, and provide insights for researchers to

navigate potential pitfalls and determine reasonable metrics to consider when optimizing for accuracy and beyond-accuracy metrics.

## CCS CONCEPTS

• **Information systems** → **Recommender systems**; *Retrieval models and ranking*.

## KEYWORDS

Next basket recommendation; Repetition and exploration; Evaluation

### ACM Reference Format:

Ming Li, Yuanna Liu, Sami Jullien, Mozhdeh Ariannezhad, Andrew Yates, Mohammad Aliannejadi, and Maarten de Rijke. 2024. Are We Really Achieving Better Beyond-Accuracy Performance in Next Basket Recommendation?. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*, July 14–18, 2024, Washington, DC, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3626772.3657835>

## 1 INTRODUCTION

Recommender systems have become an essential instrument for connecting people to the content, services, and products they need. In e-commerce, more and more consumers purchase food and household products online instead of visiting physical retail stores [18]. The COVID-19 pandemic has only accelerated this shift [33]. In this scenario, consumers usually purchase a *set* of items at the same time, a so-called *basket*. *Next basket recommendation* (NBR) is a type of sequential recommendation that caters to this scenario: baskets are the target of recommendation and historical sequential data consists of users' interactions with baskets. NBR has increasingly been attracting attention in recent years [2]. Many methods, based on different machine learning techniques, have been proposed for accurate recommendations, e.g., Markov chain (MC)-based methods [37, 41], frequency and nearest neighbor-based methods [12, 15], RNN-based methods [14, 19, 34, 49], and self-attention methods [9, 39, 50].

**Repetition vs. exploration in NBR.** Recently, Li et al. [24] have assessed the performance of state-of-the-art NBR in terms of repeat and explore items: items that a user has interacted with before and items that they have never interacted with before, respectively. The authors distinguish between the task of repetition recommendation (recommending *repeat items*) and the task of exploration

\*Both authors contributed equally to the paper.

†Work done while the author was with AIRLab, University of Amsterdam.



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

**Table 1: Comparison of the repetition and exploration tasks in NBR.**

Aspect	Repetition	Exploration
Task difficulty	Easy	Difficult
Number of items	Dozens	Thousands
Item interactions	Previous	None
Users' interest	With feedback	Without feedback
Task type	Re-consume	Infer new

recommendation (recommending *explore items*). Repetition and exploration recommendations have different levels of difficulty, where recommending items that are regularly present in a user's baskets is shown to be a far easier task [24]. Building on these findings, repetition-only [1, 17] and exploration-only [21] methods have been proposed to optimize the accuracy of next basket recommendation.

**Accuracy and beyond-accuracy metrics.** Even though accuracy naturally serves as the most important objective of recommendations, it is widely recognized that it should not be the sole focus. Beyond-accuracy metrics such as item fairness [11, 13, 46, 47] and diversity [8, 52, 53] also play crucial roles in evaluating recommendation services. Such beyond-accuracy metrics have gained increasing attention and have been optimized in a range of recommendation scenarios [48, 53]. In the NBR scenario, however, beyond-accuracy metrics have been far less studied than accuracy-based metrics. In this paper, we help to address this knowledge gap. Following the paradigm of multiple-objective recommender systems [16], it is widely recognized that there is a trade-off between accuracy and beyond-accuracy metrics. E.g., diversity goals are reckoned to stand in contrast with accuracy. Put differently, a method achieving a better beyond-accuracy performance while maintaining the same level of accuracy performance is considered to be a success [48, 53]. And how can we achieve a reasonable balance between accuracy and beyond-accuracy metrics in NBR?

**Potential “short-cuts” to balancing accuracy and beyond-accuracy metrics.** Besides the imbalance between repetition and exploration [22–25], Li et al. also found that repeat items contribute most of the accuracy, whereas the explore items in the recommended basket contribute very little to the user's perceived utility. As Table 1 summarizes, there are essential differences between the repetition and exploration tasks, which explain the substantial performance differences between the two tasks.

Inspired by these findings, we hypothesize that there may be a “short-cut” strategy to optimize for both accuracy and beyond-accuracy metrics, which contains two aspects: (i) *accuracy*: Predict repeat items to achieve good accuracy: predicting repeat items is much easier than predicting explore items [24], and (ii) *beyond-accuracy*: Use explore items to improve beyond-accuracy metrics: it is very difficult to recommend quality explore items. Thus, exchange the low accuracy that is typically achieved on such items for beyond-accuracy metrics, i.e., trade accuracy for diversity and item fairness. We call this NBR strategy a short-cut strategy because it avoids making the fundamental trade-off between accuracy and beyond-accuracy metrics.

**TREx framework.** To operationalize our short-cut idea, and check whether the “short-cut” strategy can be made to work, we propose the *two-step repetition-exploration* (TREx) framework. TREx decouples the prediction of repeat items and explore items. Specifically, TREx uses separate models for predicting (a) repeat items, and (b) explore items, and then combines the outcomes of the two prediction models to generate the next basket. In contrast, existing NBR methods usually output the scores/probabilities of all items and then select the top- $k$  items to fill up a basket to be recommended, ignoring the differences between repeat and explore items.

For TREx's repeat item prediction, we propose a simple yet effective probability-based method, which considers the item characteristics and users' repurchase frequency. For exploration recommendations, we design two strategies that cater to the different beyond-accuracy metrics. The flexibility of TREx allows us to design suitable models for repetition and exploration, with the possibility of controlling the proportions of repetition and exploration to investigate the relations between accuracy and various beyond-accuracy metrics.

**Findings and reflections.** We consider two types of widely-used beyond-accuracy metrics, i.e., diversity and item fairness. Specifically, we investigate five fairness metrics (i.e., logEUR, logRUR, EEL, EED, and logDP) [28, 36] and three diversity metrics (i.e., ILD, Entropy, and DS) [48]. To provide an overall understanding of these metrics, we group them according to different levels of connection with accuracy as follows: (i) Strong connection: logRUR, (ii) Weak connection: logEUR, EEL, EED (iii) No connection: logDP, ILD, Entropy, DS. Briefly, the strong connection between logRUR and accuracy stems from the fact that logRUR uses ground truth relevance to discount the exposure, making sure that only correctly predicted items contribute to effective exposure. The connection between logEUR, EEL, and accuracy is weak because they just ensure the exposure distribution across groups of recommended results is close to the group exposure distribution of ground truth, without considering whether the exposure is contributed by correctly predicted items. Since the position weighting model of EED considers ground truth, EED shows a weak connection. There is no connection between accuracy and logDP, ILD, Entropy, and DS because their exposure distributions across groups are designed to reflect a specific distribution. The strength of the connection between a beyond-accuracy metric and accuracy determines whether there is a short-cut towards optimizing both accuracy and the beyond-accuracy metric.

We perform experiments on two brick-and-mortar retailers' NBR datasets, considering six NBR baselines and eight metrics. The experimental results show that: (1) State-of-the-art accuracy can be achieved by only recommending repeat items via the proposed simple yet effective repetition model. (2) Leveraging the “short-cut” using TREx achieves “better” beyond accuracy performance w.r.t. seven out of eight beyond-accuracy metrics. (3) In terms of the item fairness metric having a strong connection with the accuracy (i.e., logRUR), it is more difficult to achieve better beyond-accuracy metrics via the proposed strategy.

**Stepping back.** Instead of blindly claiming TREx with the designed modules as a state-of-the-art method for optimizing both accuracy and various beyond-accuracy metrics, we reflect and challenge our

evaluation paradigm in the definition of success in this setting. The core question is:

*Are we really achieving better beyond-accuracy performance in next basket recommendation?*

Two perspectives offer different ways forward for researchers and practitioners to address this question:

- (1) If we are willing to sacrifice the accuracy of the exploration, then superior beyond-accuracy performance can be achieved by leveraging the “short-cut” strategy via TREx, which is straightforward and efficient. This “short-cut” strategy must be considered before developing more sophisticated and elaborate approaches.
- (2) Conversely, if we believe it is unreasonable to sacrifice the accuracy of exploration [45], the existence of the “short-cut” strategy reveals flaws in our current evaluation paradigm to demonstrate an NBR method’s superiority. A fine-grained analysis (i.e., distinguishing between repetition and exploration) needs to be performed to check whether “better” beyond-accuracy is achieved by triggering the “short-cut” strategy, which would hurt the exploration accuracy after all.

**Our contributions.** The main contributions of the paper are:

- We identify a “short-cut” strategy (i.e., sacrificing accuracy for exploration and using explore items to optimize for beyond-accuracy metrics), which could achieve “better” beyond-accuracy metrics without degrading accuracy.
- We propose a simple repetition recommendation model considering item features and users’ repurchase frequency, which can achieve the state-of-the-art NBR accuracy by only recommending repeat items.
- We propose TREx, a flexible two-step repetition-exploration framework for NBR, which allows us to control the trade-off between accuracy and beyond-accuracy metrics w.r.t. the recommended baskets.
- We conduct experiments on two datasets w.r.t. eight beyond-accuracy metrics, and find that leveraging “short-cuts” via TREx can achieve better performance on a wide range of metrics. We also find that the stronger the connection with accuracy, the more challenging it becomes to utilize a “short-cut” strategy to enhance a beyond-accuracy metric.
- We reflect on, and challenge, existing evaluation paradigms, and find that a fine-grained level analysis can provide a complementary view of a method’s performance.

## 2 RELATED WORK

We summarize related research on next basket recommendation and beyond-accuracy metrics.

**Next basket recommendation.** The NBR problem has been studied for many years. Factorizing personalized Markov chains (FPMC) [37] leverages matrix factorization and Markov chains to model users’ general interest and basket transition relations. HRM [41] applies aggregation operations to learn a hierarchical representation of baskets. RNNs have been adapted to the NBR task to learn long-term trends by modeling the whole basket sequence. E.g., Dream [49] uses max/avg pooling to encode baskets. Sets2Sets [14] adapts

an attention mechanism and adds frequency information to improve performance. Some methods [19, 44] consider the underlying item relations to get a better representation. Yu et al. [50] argue that item-item relations between baskets are important, and leverage GNNs to use these relations. Some methods [3, 20, 39, 42] exploit auxiliary information, including product categories, amounts, prices, and explicit timestamps. TIFUKNN [15] and UP-CF@r [12], frequency-neighbor-based methods, model temporal patterns, and then combine these with neighbor information or user-wise collaborative filtering. Li et al. [24] provide several metrics to evaluate repetition and exploration performance in the NBR task and find that the repetition task is easier than the exploration task. Inspired by this analysis, repetition-only [1, 17] and exploration-only [21] models were proposed for next basket recommendation. Existing NBR work mainly focuses on optimizing accuracy whereas this paper extends to various beyond-accuracy metrics for NBR.

**Beyond-accuracy metrics.** In addition to accuracy, there are various beyond-accuracy metrics (i.e., diversity, fairness, novelty, serendipity, coverage) we need to consider when making recommendations [11]. Diversity is a crucial factor in meeting the diverse demands of users [7, 35, 43, 52]. Recently, empirical and revisitation studies [29, 48] have been conducted to explore the trade-off between accuracy and diversity. The concepts of fairness and item exposure have emerged as crucial considerations since items and producers play pivotal roles within a recommender system and its ecosystem. Related metrics measure whether items receive a fair share of exposure according to different definitions of fairness. Current research on fairness primarily focuses on individual or group fairness, either from the customer’s perspective, adopting a user-centered approach [5], or from the provider’s viewpoint, adopting an item-centered approach [30, 51], or a two-sided approach [31, 46, 47]. Recently, Liu et al. [28] evaluated item fairness on existing NBR methods to investigate the robustness of different fairness metrics. Unlike the work listed above, this paper is not limited to optimizing a specific type of metric. It examines the possibility of leveraging a “short-cut” strategy to seemingly optimize various beyond-accuracy metrics and provides insights w.r.t. evaluation paradigms when extending NBR optimization and evaluation to these beyond-accuracy metrics.

## 3 TASK FORMULATION AND DEFINITIONS

We describe the next basket recommendation problem and formalize the notions of repetition and exploration. Our notation is summarized in Table 2.

**Next basket recommendation.** Given a set of users  $U = \{u_1, u_2, \dots, u_n\}$  and items  $I = \{i_1, i_2, \dots, i_m\}$ ,  $S_u = \{B_u^1, B_u^2, \dots, B_u^t\}$  represents the historical interaction sequence for  $u$ , where  $B_u^t$  is the user’s basket at the time step  $t$ .  $B_u^t$  consists of a set of items  $i \in I$ , and the goal of the *next basket recommendation* task is to predict  $P_u = B_u^{t+1}$ , the following basket of items that the user would probably like, based on the user’s past interactions  $S_u$ , i.e.,

$$P_u = \hat{B}_u^{t+1} = f(S_u), \quad (1)$$

where  $f$  is our basket generation algorithm. We assume that the user’s attention and screen space is limited; hence, like previous studies [24, 28], we recommend fixed-size baskets of sizes 10 or 20.

**Table 2: Notation used in the paper; fairness related notation is adapted from [28, 36].**

Symbol	Description
$u \in U$	Users
$i \in I$	Items
$S_u$	Sequence of historical baskets for $u$
$B_u^t$	$t$ -th basket in $S_u$ , a set of items $i \in I$
$I_{u,t}^{rep}$	Set of repeat items for $u$ up to timestamp $t$
$I_{u,t}^{expl}$	Set of explore items for $u$ up to timestamp $t$
$T_u$	Ground-truth basket for $u$ that we aim to predict
$T_u^{rep}$	Set of <i>repeat items</i> in the ground truth basket $T_u$ for $u$
$T_u^{expl}$	Set of <i>explore items</i> in the ground truth basket $T_u$ for $u$
$P_u$	Predicted basket for $u$
$P_u^{rep}$	Set of <i>repeat items</i> in the predicted basket $P_u$ for $u$
$P_u^{expl}$	Set of <i>explore items</i> in the predicted basket $P_u$ for $u$
$G(P)$	Group alignment matrix for items in $P$
$G^+$	Popular group
$G^-$	Unpopular group
$\mathbf{a}_P$	Exposure vector for items in $P$
$\epsilon_P$	The exposure of groups in $P$ ( $G(P)^T \mathbf{a}_P$ )

**Repetition and exploration.** We assume that the set of items is fixed. Although this might not be the case in real-world settings, modeling the addition and deletion of items in the set of items is out of the scope of this paper. With this assumption in mind, the addition of every new basket to the users’ history, may translate into fewer items left to explore. To differentiate between the items coming from the exploration and repeat consumption behavior, for a user  $u$  and timestamp  $t$ , a set of items  $I_{u,t}^{rep} \subset I$  are considered to be the “repeat items.” The set of explore items  $I_{u,t}^{expl}$  is simply its complement within the overall item set  $I$ . We define  $I_{u,t}^{rep}$  as:

$$I_{u,t}^{rep} = I_{u,t-1}^{rep} \cup B_u^t. \quad (2)$$

This also means that  $I_{u,1}^{rep} \subset \dots \subset I_{u,t-1}^{rep} \subset I_{u,t}^{rep}$ . Conversely, we have  $I_{u,t}^{expl} \subset I_{u,t-1}^{expl} \subset \dots \subset I_{u,1}^{expl}$ .

The task of predicting the next basket for a user  $u$  is equivalent to predicting which items from  $I_{u,t}^{rep}$  and  $I_{u,t}^{expl}$  will appear in  $B_u^{t+1}$ . One way to solve this problem is to decouple it into two subtasks: the repetition subtask that aims to predict which items from  $I_{u,t}^{rep}$  to recommend, and the exploration task that recommends items from  $I_{u,t}^{expl}$ . Table 1 shows the different characteristics w.r.t. the repetition and exploration tasks.

## 4 EVALUATION METRICS

Next, we describe the accuracy and beyond-accuracy metrics (i.e., fairness and diversity) considered in the paper.<sup>1</sup>

**Accuracy.** In terms of accuracy, we use three metrics that are widely used for the NBR task: *Recall@k*, *NDCG@k*, and *PHR@k*. *Recall* measures the ability to find all items that the user will purchase in the next basket; *NDCG* is a ranking metric that also considers the order of the items; *PHR* is a user level measurement which

<sup>1</sup>Due to space limitations, we only provide brief introductions of each metric; more detailed information (e.g., function, responsibility, etc.) can be found in the original papers and relevant survey papers [28, 36, 53].

represents the ratio of users whose recommended basket contains the item in the ground-truth.

**Fairness.** Assume  $\pi(P | u)$  is a user-dependent distribution and  $\rho(u)$  is a distribution over users; overall, the recommended item rankings among all users follow the following distribution:  $\rho(u)\pi(P | u)$ .  $\epsilon_P = G(P)^T \mathbf{a}_P$  is the group exposure within a recommended basket.<sup>2</sup> Its expected value  $\epsilon_\pi = E_{\pi\rho}[\epsilon_P]$  is the group exposure among all the recommended baskets. Following [28, 36], we select a set of well-known fairness metrics and cover two types of fairness considerations as follows:<sup>3</sup>

(1) *Equal opportunity.* Promote equal treatment based on merit or utility, regardless of group membership [28, 36]. (i) *Exposed Utility Ratio (EUR)* [38] quantifies the deviation from the objective that the exposure of each group is proportional to its utility  $Y(G)$ . (ii) *Realized Utility Ratio (RUR)* [38] models actual user engagement, the click-through rates for the groups  $\Gamma(G)$  are proportional to their utility. (iii) *Expected Exposure Loss (EEL)* [10] is the distance between expected exposure and target exposure  $\epsilon^*$ , which is the exposure under the ideal policy.

(2) *Statistical parity.* Ensure comparable exposure among groups. (i) *Expected Exposure Disparity (EED)* [10] measures the inequality in exposure distribution across groups. (ii) *Demographic Parity (DP)* [38] measures the ratio of average exposure given to the two groups. Following [36], we reformulate DP as logDP to tackle the issue of empty-group scenarios and improve interpretability. Exposed Utility Ratio (logEUR) and Realized Utility Ratio (logRUR) are defined in a similar manner.

**Diversity.** Following [48], we consider the following widely-used diversity metrics, which satisfy users’ diversified demands. (i) *Intra-List Distance (ILD)* [6, 7] measures the average distance between every pair of items in the recommendation list ( $P_u$ ), where  $d_{ij}$  is the Euclidean distance between the respective embeddings of categories (ii) *Entropy* [43, 54] quantifies the dispersion of item category distribution in the recommendation list  $P_u$ ; a higher degree of dispersion in the category distribution corresponds to increased diversity. (iii) *Diversity Score (DS)* [27] is calculated as the number of interacted/recommended categories divided by the number of interacted/recommended items. As shown in Table 3, we can group beyond-accuracy metrics according to their connection with accuracy.

## 5 A TWO-STEP REPETITION-EXPLORATION FRAMEWORK

Given the differences depicted in Table 1, we propose a *two-step repetition-exploration* (TReX) framework for NBR. TReX assembles recommendations from a repetition and an exploration module. TReX allows one to easily swap out the sub-algorithms used for repetition and exploration. In the first step, we model the repetition and exploration behavior separately to get candidates from both sources. Then, we generate the recommended basket from those candidates in the second step. The main architectural differences between previous approaches to the NBR problem, which typically

<sup>2</sup>The formula to compute the exposure vector  $\mathbf{a}_P$  using different position weighting models can be found in [28, 36].

<sup>3</sup>Item fairness metric Inequity of Amortized Attention [4] is not used in this paper since some baselines do not have predicted relevance for items.

**Table 3: Summary of fairness and diversity metrics; fairness metrics are adapted from [36].  $\uparrow$  indicates that higher values are better;  $\downarrow$  indicates that lower values are better;  $\circ$  means that the closer the value is to 0, the better the performance.**

Category	Metrics	Goal	Better	Accuracy connection
Equal opportunity	logRUR	Click-through rate proportional to relevance	$\circ$	Strong
	logEUR	Exposure proportional to relevance	$\circ$	Weak
	EEL	Exposure matches ideal (from relevance)	$\downarrow$	Weak
Statistical parity	EED	Exposure well-distributed	$\downarrow$	Weak
	logDP	Exposure equal across groups	$\circ$	None
Diversity	ILD	Average distance between categories for each pair of items in the list	$\uparrow$	None
	Entropy	Entropy of item category distribution in the list	$\uparrow$	None
	DS	Number of categories divided by the number of items in the list	$\uparrow$	None

**Algorithm 1:** TREx Framework

---

**Data:** Basket sequence  $S$ , basket size  $k$ , repetition confidence threshold  $v$

**Result:** Recommended basket  $B_u^{t+1}$  for each user  $u$ ,

- 1 Calculate the repetition feature  $RepI(i)$  for each item;
- 2 **for** each user  $u$  **do**
- 3     Get repeat items  $I_{u,t}^{rep}$ , and explore items  $I_{u,t}^{expl}$ ;
- 4     Calculate the repetition score  $RepS^u(i)$  for each  $i \in I_{u,t}^{rep}$ ;
- 5     Remove items  $i$  from  $I_{u,t}^{rep}$ , when  $RepS^u(i) < v$ ;
- 6     Rank  $I_{u,t}^{rep}$  according to  $RepS^u(i)$  in descending order;
- 7     Initialize next basket  $B_u^{t+1}$ ;
- 8     **if**  $|I_{u,t}^{rep}| < k$  **then**
- 9         Fill  $B_u^{t+1}$  using  $I_{u,t}^{rep}$ ;
- 10          $m \leftarrow k - |I_{u,t}^{rep}|$ ;
- 11         Fill  $m$  empty slots of  $B_u^{t+1}$  using explore items via exploration module;
- 12     **else**
- 13         Fill  $B_u^{t+1}$  using top- $k$  of  $I_{u,t}^{rep}$ ;
- 14     **end**
- 15 **end**

---

consists of a single treatment of all items, and TREx, which treats repeat and explore items differently. The pseudo-code for TREx is given in Algorithm 1. Next, we describe the three modules that make up TREx.<sup>4</sup>

### 5.1 Repetition module

As the repetition task is a much simpler task than exploration, we therefore design a repetition module targeted at improving the accuracy. Intuitively, if a user consumed an item several times in the past, they are likely to repurchase that item in the next basket. Thus, frequency information is a strong signal for repetition prediction [40]. The personal item frequency (PIF) introduced in TIFUKNN [15] and the recency window in UP-CF@r[12] both capture temporal dependencies by focusing more on recent behavior. However, they do not capture the item characteristics w.r.t. repurchasing. For example, a purchase of a bottle of milk and a pan is more likely to be followed by a repurchase of milk rather than a

pan, even if both currently have the same purchase frequency. To consider both item features and user interest simultaneously, we use the repetition score  $RepS^u(i)$  to represent the repurchase score of item  $i$  for user  $u$ . This score is decomposed into two parts, the item-specific repurchase feature  $RepI(i)$  and the user’s interest  $E_i^u$  in item  $i$ . Formally:

$$RepS^u(i) = E_i^u \cdot RepI(i). \quad (3)$$

This corresponds to line 4 in Algorithm 1. Given the items in the dataset  $I = \{i_1, i_2, \dots, i_m\}$ , we need to derive the repurchase feature  $RepI(i)$  for each item in the training set. First, the repurchase frequency  $Rep^F(i)$  can be calculated by gathering the statistical information across users. To mitigate the impact of abnormally high values in some users, we introduce a hyperparameter  $\alpha$  to discount the repurchase frequency of item  $i$ .

$$Rep^F(i) = \frac{\sum_U (\text{item } i \text{ repurchase frequency})^\alpha}{\#\text{users who bought item } i \text{ at least once}}. \quad (4)$$

In addition, some items might only have a few samples, which might lead to low confidence about their repetition feature estimation. We leverage the average estimate  $\overline{Rep^F}$  across all items as supplementary information to help items with a few samples. Then, the final repetition feature is given by:

$$RepI(i) = Rep^F(i) + \frac{\overline{Rep^F}}{N_i}, \quad (5)$$

where  $N_i$  is the number of users who bought item  $i$ . Thus, the average  $\overline{Rep^F}$  will have a small effect on  $RepI(i)$  when we have more samples to compute item-specific features. This corresponds to line 1 in Algorithm 1.

The item frequency in a user’s historical baskets can partially reflect the user’s interest. Yet, user interests can also be dynamic. To model temporal dependencies, we introduce a time-decay factor  $\beta$ , which makes the recent interactions have more impact on the interest  $E_i^u$ . Assume that a specific item  $i$  was purchased by the user  $u$  several times in their historical baskets  $\{B_u^1, B_u^2, \dots, B_u^m\}$ ; the corresponding position set is denoted as  $L_i = \{l_1, l_2, \dots, l_m\}$ ; then  $E_i^u$  is defined as:

$$E_i^u = \sum_{j=1}^m \beta^{T-l_j}, \quad (6)$$

where  $T$  represents the length of the user’s basket. TREx’s repeat recommendation model takes item features, user interests, and the temporal order of baskets into consideration. We treat the items in baskets independently and calculate the repetition score  $RepS$  for all items that appeared in the previous baskets for each user, which will be used in the final basket generation process.

<sup>4</sup>Theoretically, TREx allows us to choose or design the suitable repetition and exploration modules both targeted at the accuracy to achieve state-of-the-art performance. However, we aim to investigate the “short-cut” and relationship between accuracy and various beyond-accuracy metrics.

**Table 4: Statistics of the processed datasets.**

Dataset	#items	#users	Avg. basket size	Avg. #baskets per user	Repeat ratio	Explore ratio
Instacart	29,399	19,210	10.06	15.91	0.60	0.40
Dunnhumby	37,162	2,482	10.07	43.17	0.43	0.57

## 5.2 Exploration module

As it is more challenging than repetition, exploration is also an important aspect of NBR. To complement the repetition module, we design different exploration modules, targeting item fairness and diversity, respectively. For each user  $u$ , the exploration candidates  $I_{u,t}^{expl}$  are the set of items that the user never bought before.

**Item fairness.** According to [24], we find that NBR methods usually have varying degrees of popularity bias, which means they recommend more popular items compared to the ground truth and harm item fairness. Thus, we recommend unpopular items  $i \in G^-$  for the exploration module for the sake of approaching the distribution of ground truth and decreasing the exposure gap between the popular and the unpopular groups. Specifically, we randomly sample explore items based on a sampling probability, which is calculated from the purchase frequency of unpopular items.

**Diversity.** Diversity optimizes for more dispersed categories in the predicted basket. For each user, we record categories of repetition candidates, rank exploration candidates according to their popularity, and select explore items to fill in the  $B_u^{t+1}$  in turn. The category of each explore item differs from the categories already in  $B_u^{t+1}$ .

## 5.3 Basket generation module

To construct the final basket to be recommended by TReX for the accuracy objective, we adopt a repetition greedy approach and first consider the item candidates generated by the repetition module and fill the remaining slots via the exploration module.  $TReX_{Fairness}$  and  $TReX_{diversity}$  denote TReX with the exploration module targeted at fairness and diversity, respectively. For a user  $u$ , we get their repetition score  $RepS^u(i)$ , where  $i \in I_{u,t}^{rep}$  (Algorithm 1, lines 3–4). First, we define a confidence threshold  $v$  for the repetition score and repetition items are removed from the  $i \in I_{u,t}^{rep}$  when the corresponding  $RepS^u(i) < v$  (line 5).<sup>5</sup> Then,  $I_{u,t}^{rep}$  can be seen as the repetition candidates set. If the number of repetition candidates exceeds the basket size, the items with a high score will have priority to fill the basket (Algorithm 1, line 13). If the number of repetition candidates is smaller than the basket size, the basket is first filled with all items in the repetition candidates set  $I_{u,t}^{rep}$ . Then, we fill up the basket using the explore items via the exploration module, where  $m$  represents the number of empty slots (lines 9–11).

# 6 EXPERIMENTS


## 6.1 Experimental setup

**Datasets.** We conduct experiments on two widely-used datasets: (i) Instacart,<sup>6</sup> which includes a large number of grocery orders from

<sup>5</sup>The confidence threshold  $v$  controls the proportion of repeat items and explore items in the recommendation, as well as the accuracy and beyond-accuracy trade-off in this paper. We sweep repetition confidence bound  $v$  to get TReX variants with different accuracy and beyond-accuracy metrics performance.

<sup>6</sup><https://www.kaggle.com/c/instacart-market-basket-analysis/data>

users; following [28, 32], ~20000 users are randomly selected to conduct experiments; and (ii) Dunnhumby,<sup>7</sup> which contains two years' household-level transactions of 2500 frequent shoppers at a retailer. Following [1, 28], we sample users who have at least three baskets and remove items that appeared less than five times. The two datasets vary in the repeat ratio, i.e., the proportion of repeat items in the ground-truth baskets [24]. We focus on the fixed size (10 or 20) NBR problem. The statistics of the processed datasets are shown in Table 4. In our experiments, each dataset is partitioned according to [1, 12, 28, 32]. The training baskets encompass all user baskets except the last one. In cases where users have over 50 baskets in the training data, only their last 50 baskets are considered for inclusion in the training set. The final baskets of all users are then divided equally between a 50% validation set and a 50% test set. Figure 1 shows the distribution of users across repeat ratios, which is the proportion of *repeat items* in the ground-truth basket.



**Figure 1: Distribution of users across different repeat ratios for Instacart and Dunnhumby.**

**NBR baselines.** We compare TReX with 8 representative baselines, which we select based on their characteristics in the analysis performed in [24, 28], divided into three groups:


**6.1.1 Simple baselines.** (i) **G-TopFreq** uses the  $k$  most popular items in the dataset to form the recommended next basket. (ii) **P-TopFreq** is a personalized TopFreq method, which treats the most frequent  $k$  items in historical records of the user as the next basket. (iii) **GP-TopFreq** [24] is a simple combination of P-TopFreq and G-TopFreq, which first use P-TopFreq to fill the basket, then use G-TopFreq to fill the remaining slots.

**6.1.2 Nearest neighbor-based methods.** (i) **TIFUKNN** [15] is a state-of-art method that models the temporal dynamics of frequency information of users' past baskets to introduce Personalized Frequency Information (PIF), then it uses KNN-based method on the PIF. (ii) **UP-CF@r** [12] is a combination of recency aware user-wise popularity and user-wise collaborative filtering.

**6.1.3 Neural network-based methods.** (i) **Dream** [49] models users' global sequential basket behavior for NBR using recurrent neural network (RNN). (ii) **DNNTSP** [50] is a state-of-art method that leverages a GNN and self-attention techniques. It encodes item-item relations via a graph and employs a self-attention mechanism to capture temporal dependencies of users' basket sequences. (iii) **ReCANet** [1] is a repeat-only model for NBR, which uses user-item representations with historical consumption patterns via RNN.

**Configurations.** To assess group fairness (Section 4), we follow configurations from previous research [26, 28]; the group of items is determined by their popularity (i.e., the number of purchases

<sup>7</sup><https://www.dunnhumby.com/source-files/>



**Figure 2: Performance of TREx-Rep when we add a time-decay factor  $\beta$  (+T), add both  $\beta$  and item-specific repetition feature  $Repl(i)$  (+T+RF).**


recorded in the historical baskets of the dataset). The top 20% of items with the highest purchase frequency as the popular group ( $G^+$ ), while the remaining 80% of items are assigned to the unpopular group ( $G^-$ ). For the baseline methods, a grid search is performed to find the optimal hyper-parameters via the validation set. For TIFUKNN, the number of neighbors  $k$  is tuned on  $\{100, 300, 500, 900, 1100, 1300\}$ , the number of groups  $m$  is tuned on  $\{3, 7, 11, 15, 19, 23\}$ , the within-basket time-decayed ratio  $r_b$  and the group time-decayed ratio  $r_g$  are selected from  $\{0.1, 0.2, \dots, 0.9, 1\}$ , and the fusion weight  $\alpha$  is selected from  $\{0, 0.1, \dots, 0.9, 1\}$ . For UP-CF@r, recency window  $r$  is tuned on  $\{1, 5, 10, 25, 100, \infty\}$ , locality  $q$  is tuned on  $[1, 5, 10, 50, 100, 1000]$ , and asymmetry  $\alpha$  is tuned on  $\{0, 0.25, 0.5, 0.75, 1\}$ . For Dream, DNNTSP, and ReCANet, the item and user embedding size is tuned on  $\{16, 32, 64, 128\}$ . As to TREx, for the repetition module,  $\alpha$  is selected from  $\{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$ , and the time-decay factor  $\beta$  is selected from  $\{0.7, 0.75, 0.8, 0.85, 0.9, 0.95, 1.0\}$ . To facilitate reproducibility, we release the source code and all hyper-parameters in an online repository: <https://github.com/lynEcho/TREx>.

## 6.2 Overall accuracy performance


By decoupling the repetition and exploration tasks, TREx-Rep optimizes for the repeat items prediction and accounts for the accuracy of the NBR performance. Table 5 shows the experimental results for TREx-Rep and the baselines. We observe that TREx-Rep surpasses two complex deep learning-based methods (i.e., Dream and DNNTSP) by a large margin on the Dunnhumby and Instacart datasets, and TREx-Rep always achieves or matches the SOTA accuracy on both datasets across different accuracy metrics. Note that, TREx-Rep achieves a competitive accuracy performance by only using part of the available slots in the basket.<sup>8</sup> Compared to the deep learning methods with complex architectures that try to learn basket representations and model temporal relations, TREx-Rep is very efficient due to its simplicity.

To investigate the effect of the repetition features and the improvement in repetition performance in NBR. We conduct experiments on TREx-Rep by gradually adding the time-decay factor  $\beta$  and item-specific repetition feature  $Repl(i)$ . The results are shown in Figure 2. The accuracy increases when we gradually integrate different factors into TREx-Rep, which indicates that both the time-decay factor  $\beta$  and the item-specific repetition feature  $Repl(i)$  contribute

<sup>8</sup>As TREx-Rep only recommends repeat items, the basket could not be fulfilled when the number of user’s repeat items (historical items) is smaller than the basket size. ReCANet also only recommends repeat items, however, it is a complex neural-based model, which is much slower than the proposed TREx-Rep module.



**Figure 3: The recall improvement of (+T+RF) over (+T) when the training sample ratio changes from 0.2 to 1.**



**Figure 4: Performance of TREx<sub>Diversity</sub> at different  $v$  values, compared with different NBR methods in terms of different diversity metrics. The red + marker indicates the direction with both high accuracy and diversity.**

to the accuracy performance of TREx-Rep. Significant improvements over only using the time-decay factor  $\beta$  can be observed on the Dunnhumby dataset when the item-specific repetition feature  $Repl(i)$  is also adopted to compute the repetition score  $Reps^u(i)$ . Note that the improvement of adding  $Repl(i)$  to TREx-Rep on the Instacart dataset is relatively small. We conjecture that items in the Instacart dataset are more regular products, that have little difference in repetition feature with each other. Figure 3 shows the performance when using different amounts of training samples, the improvement in recall resulting from adding  $Repl(i)$  increases when we use more training data since we have more samples for estimating the repetition feature  $Repl(i)$ .

## 6.3 Beyond-accuracy performance

We conduct experiments to verify whether TREx with the designed models (i.e., TREx<sub>Diversity</sub> and TREx<sub>Fairness</sub>) could achieve better performance on representative diversity and item fairness metrics. Note that, the recommended basket remains fixed for a specific user in existing baselines, resulting in fixed performance regarding both accuracy and beyond-accuracy metrics on each dataset. In contrast, TREx provides the flexibility to adjust the trade-off between accuracy and beyond-accuracy metrics by adjusting the repetition confidence bound  $v$ . This allows for a more nuanced control over the recommendation process compared to traditional baselines.

**Table 5: Comparison of TREx-Rep (repetition-module only) against baselines and two types of state-of-art methods; boldface indicates the maximum; underlining indicates the second best performing method. † indicates that TREx-Rep results achieve the same level of performance as SOTA baselines (paired t-test).**

Dataset	Metric	G-TopFreq	P-TopFreq	GP-TopFreq	UP-CF@r	TIFUKNN	Dream	DNNTSP	ReCANet	TREx-Rep
Instacart	Recall@10	0.0704	0.3143	0.3150	0.3377	0.3456	0.0704	0.3295	<b>0.3490</b>	0.3476†
	NDCG@10	0.0817	0.3339	0.3343	0.3582	0.3657	0.0817	0.3434	<b>0.3699</b>	0.3661†
	PHR@10	0.4600	0.8447	0.8460	0.8586	0.8639	0.4600	0.8581	<b>0.8668</b>	0.8655†
	Recall@20	0.0973	0.4138	0.4168	0.4405	<u>0.4559</u>	0.0979	0.4339	<b>0.4562</b>	0.4557†
	NDCG@20	0.0962	0.3889	0.3902	0.4161	<u>0.4271</u>	0.0968	0.4018	<b>0.4303</b>	0.4269†
	PHR@20	0.5302	0.8921	0.8959	0.9045	<b>0.9098</b>	0.5346	0.9033	<u>0.9097</u>	0.9092†
Dunnhumby	Recall@10	0.0897	0.1628	0.1628	0.1699	<u>0.1763</u>	0.0896	0.0871	0.1730	<b>0.1815</b> †
	NDCG@10	0.0798	0.1562	0.1562	0.1639	<u>0.1683</u>	0.0759	0.0792	0.1625	<b>0.1689</b> †
	PHR@10	0.3795	0.5399	0.5399	0.5536	<u>0.5729</u>	0.3873	0.4303	0.5655	<b>0.5761</b> †
	Recall@20	0.1046	0.2075	0.2075	0.2168	0.2227	0.1081	0.1442	<u>0.2252</u>	<b>0.2257</b> †
	NDCG@20	0.0877	0.1787	0.1787	0.1885	<u>0.1917</u>	0.0853	0.1021	0.1879	<b>0.1921</b> †
	PHR@20	0.4392	0.6116	0.6116	0.6326	0.6342	0.4558	0.5378	<u>0.6377</u>	<b>0.6390</b> †

**Diversity.** The experimental results w.r.t. the accuracy and different diversity metrics (i.e., ILD, Entropy, and DS) are shown in Figure 4.<sup>9</sup> We have the following observations: (1) Compared to methods (i.e., TIFUKNN and ReCANet) with the best accuracy, TREx<sub>Diversity</sub> can achieve better performance in terms of all three diversity metrics while preserving the same level of accuracy on both datasets. (2) In contrast to other baseline methods (excluding TIFUKNN and ReCANet), TREx<sub>Diversity</sub> showcases the ability to recommend baskets with enhanced accuracy and diversity simultaneously.

**Item fairness.** The experimental results regarding the accuracy and five fairness metrics (LogRUR, logEUR, logDP, EEL, and EED) are depicted in Figure 5. Based on our analysis, we make the following observations: (i) On the Dunnhumby dataset, TREx<sub>Fairness</sub> demonstrates superior fairness w.r.t. logDP and logEUR while maintaining the same level of accuracy performance as the best-performing baselines (i.e., TIFUKNN and ReCANet). Similarly, on Dunnhumby, TREx<sub>Fairness</sub> showcases enhanced fairness across four fairness metrics (logDP, logEUR, EEL, and EED) while achieving accuracy performance comparable to the best-performing baselines. (ii) TREx<sub>Fairness</sub> demonstrates its capability to recommend baskets with improved accuracy and fairness w.r.t. logDP and logEUR concurrently, when compared to complex baselines such as Dream, UP-CF@r, and DNNTSP. (iii) In terms of logRUR, TREx<sub>Fairness</sub> exhibits inferior performance in fairness while maintaining similar accuracy levels compared to several existing baselines. Moreover, as both accuracy and fairness decrease simultaneously, a win-win and lose-lose scenario is evident rather than a conventional trade-off relationship in this fairness evaluation.

**Connections with accuracy.** To get a better understanding of the possibility of leveraging the “short-cut” via TREx to improve beyond-accuracy metrics, we conduct an analysis by categorizing these beyond-accuracy metrics into different groups based on their connections with accuracy (see Section 4 and Table 3).

We can observe that TREx can easily achieve better performance w.r.t. beyond-accuracy metrics have no connections with the accuracy (i.e., ILD, Entropy, DS, and logDP) on two datasets. When

beyond-accuracy metrics (e.g., logEUR, EEL, and EED) exhibit weak associations with accuracy, TREx outperforms alternative methods in some instances (4 out of 6). However, in cases where beyond-accuracy metrics are strongly correlated with accuracy (e.g., logRUR), TREx struggles to achieve superior performance. Since only accurate predictions contribute to improvements in logRUR fairness, leveraging the exploration module to optimize such beyond-accuracy metrics is very challenging.

## 6.4 Reflections and discussions

The above results verify our hypothesis and demonstrate the effectiveness of leveraging a “short-cut” strategy to achieve better beyond-accuracy under the current evaluation paradigms.

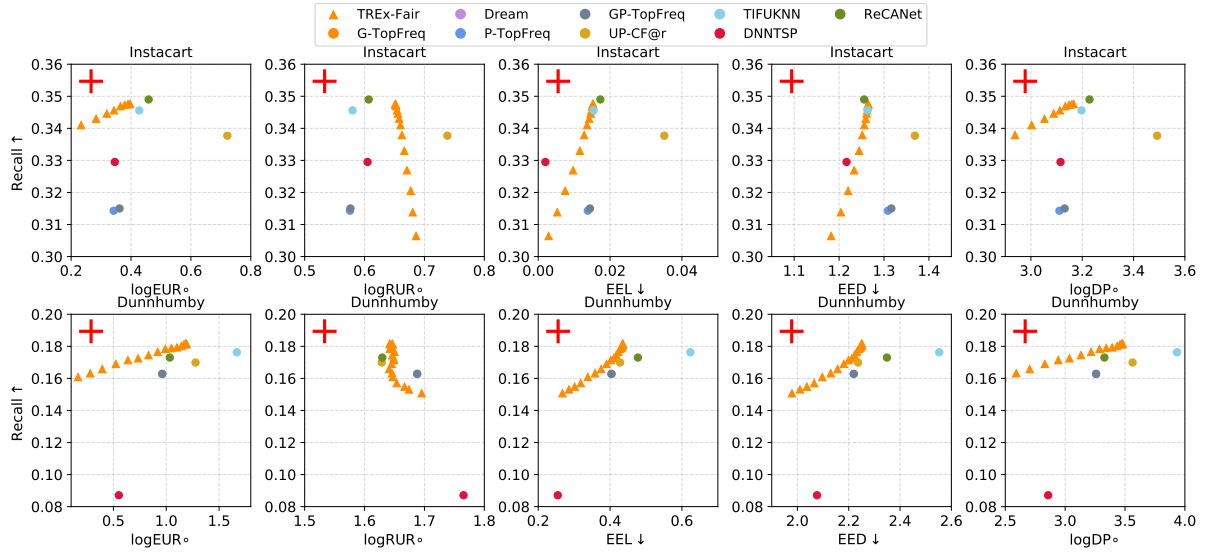
It is controversial to use this “short-cut” strategy in real-world scenarios when NBR practitioners consider beyond-accuracy metrics. In scenarios where the accuracy of exploration is not important to practitioners and only overall accuracy is of concern, the “short-cut” strategy proves to be a straightforward and efficient means to achieve better performance w.r.t. various beyond-accuracy metrics. TREx must be considered or serve as a baseline before designing more sophisticated methods, such as including multi-objective loss functions [7, 20], integer programming [53], and so on.

However, in some scenarios, it is unreasonable to sacrifice the exploration accuracy despite it being low. Therefore, the existence of the “short-cut” strategy reveals the potential flaws of the existing evaluation paradigms (i.e., using overall metrics to define success). We look into the exploration accuracy [24] of TREx<sub>Diversity</sub> when it outperforms several existing baselines in terms of both overall accuracy and diversity (i.e., success according to existing evaluation paradigm). Table 6 shows the huge decrease in the accuracy of exploring items in the recommended basket of TREx<sub>Diversity</sub>, compared to these baselines, since the designed module in TREx<sub>Diversity</sub> is mainly designed for improving diversity and does not consider accuracy. In this sense, we can not simply claim the superiority of TREx<sub>Diversity</sub> compared to these baselines just depends on the overall performance.

Note that, the fundamental reason for the existence of this “short-cut” is that predicting accurate *explore items* is much more difficult

<sup>9</sup>G-TopFreq and Dream exhibit low recall, fairness, and diversity, which prevents them from being visible in Figures 4 and 5.





**Figure 5: Performance of  $TREx_{Fairness}$  at different  $v$  values, compared with different NBR methods in terms of different fairness metrics. The red + marker indicates the direction with both high accuracy and fairness.**

than predicting *repeat items*, and exploration prediction only accounts for a limited user’s overall accuracy [22–25]. Given that exploration prediction contributes only minimally to the overall accuracy of users, it becomes feasible to allocate resources toward optimizing other beyond-accuracy metrics instead of accuracy itself.

Therefore, beyond using the overall performance to measure accuracy and beyond-accuracy metrics, a fine-grained level evaluation could help to provide a more rigid identification of the success when considering beyond-accuracy metrics.

**Table 6: Exploration accuracy [24] of  $TREx_{Diversity}$  compared with NBR methods that are inferior to it within existing evaluation paradigms.**

Dataset	Metric	TIFUKNN	Dream	DNNTSP	$TREx_{Div}$
Instacart	$Recall_{expl}@10$	0.0014	0.0322	0.0014	0.0002
	$PHR_{expl}@10$	0.0037	0.1431	0.0040	0.0009
	$Recall_{expl}@20$	0.0077	0.0526	0.0072	0.0008
	$PHR_{expl}@20$	0.0198	0.2120	0.0217	0.0031
Dunnhumby	$Recall_{expl}@10$	0.0042	0.0111	0.0017	0.0000
	$PHR_{expl}@10$	0.0139	0.0521	0.0085	0.0019
	$Recall_{expl}@20$	0.0069	0.0214	0.0028	0.0016
	$PHR_{expl}@20$	0.0232	0.1045	0.0115	0.0065

## 7 CONCLUSION

We have expanded the research objectives of NBR to go beyond sole accuracy to encompass both accuracy and beyond-accuracy metrics. We have recognized a potential “short-cut” strategy to optimize beyond-accuracy metrics while preserving high accuracy levels. To capitalize on and validate the presence of such “short-cuts,” we have introduced a plug-and-play framework called *two-step repetition-exploration* (TREx) considering the differences between repetition and exploration tasks. This framework treats repeat items and explore items as distinct entities, employing a straightforward yet highly effective repetition module to uphold accuracy standards.

Concurrently, two exploration modules have been devised to target the optimization of beyond-accuracy metrics. We have conducted experiments on two publicly available datasets w.r.t. eight representative beyond-accuracy metrics, including item fairness (i.e., logEUR, LogRUR, logDP, EEL, and EED) and diversity (i.e., ILD, Entropy, and DS).

Our experimental results demonstrate the effectiveness of our proposed “short-cut” strategy, which can achieve better beyond-accuracy performance w.r.t. several fairness and diversity metrics on different datasets. Additionally, we group beyond-accuracy metrics according to the strength of their connection with accuracy. Our analysis reveals that the stronger the connection with accuracy, the more difficult it becomes to employ a “short-cut” strategy to optimize these beyond-accuracy metrics, favoring the metrics with a stronger connection to avoid such short-cuts.

As to the broader implications of our work, we have discussed the reasonableness of leveraging the “short-cut” strategy to trade the accuracy of exploration for beyond-accuracy metrics in various scenarios. The presence of this “short-cut” highlights a potential flaw in the definition of success within existing evaluation paradigms, particularly in scenarios where exploration accuracy is important despite being low [45]. A fine-grained level evaluation should be performed in NBR to offer a more precise identification of achieving “better” performance in such a scenario.

Despite the simplicity of the “short-cut” strategy and TREx, our paper sheds light on the research direction of considering both accuracy and beyond-accuracy metrics in NBR. Rather than blindly embracing sophisticated methods in NBR, follow-up research should realize the existence of the “short-cut” and potential flaws of existing evaluation paradigms in this research direction.

**Acknowledgements.** This work is partially supported by the Dutch Research Council (NWO), under project numbers 024.004.022, NWA.1389.20.183, KICH3.LTP.20.006, and VILvidi.223.166. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

## REFERENCES

- [1] Moshdeh Ariannezhad, Sami Jullien, Ming Li, Min Fang, Sebastian Schelter, and Maarten de Rijke. 2022. ReCANet: A Repeat Consumption-Aware Neural Network for Next Basket Recommendation in Grocery Shopping. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 1240–1250.
- [2] Moshdeh Ariannezhad, Ming Li, Sami Jullien, and Maarten de Rijke. 2023. Complex Item Set Recommendation. In *SIGIR 2023: 46th international ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 3444–3447.
- [3] Ting Bai, Jian-Yun Nie, Wayne Xin Zhao, Yutao Zhu, Pan Du, and Ji-Rong Wen. 2018. An Attribute-aware Neural Attentive Model for Next Basket recommendation. In *The 41st International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1201–1204.
- [4] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international ACM SIGIR conference on research & development in information retrieval*. 405–414.
- [5] Jesús Bobadilla, Raúl Lara-Cabrera, Ángel González-Prieto, and Fernando Ortega. 2020. Deepfair: Deep Learning for Improving Fairness in Recommender Systems. *arXiv preprint arXiv:2006.05255* (2020).
- [6] Yukuo Cen, Jianwei Zhang, Xu Zou, Chang Zhou, Hongxia Yang, and Jie Tang. 2020. Controllable Multi-interest Framework for Recommendation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2942–2951.
- [7] Wanyu Chen, Pengjie Ren, Fei Cai, Fei Sun, and Maarten de Rijke. 2020. Improving end-to-end sequential recommendations with intent-aware diversification. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 175–184.
- [8] Wanyu Chen, Pengjie Ren, Fei Cai, Fei Sun, and Maarten De Rijke. 2021. Multi-interest Diversification for End-to-end Sequential Recommendation. *ACM Transactions on Information Systems* 40, 1 (2021), 1–30.
- [9] Yongjun Chen, Jia Li, Chenghao Liu, Chenxi Li, Markus Anderle, Julian McAuley, and Caiming Xiong. 2021. Modeling Dynamic Attributes for Next Basket Recommendation. *arXiv preprint arXiv:2109.11654* (2021).
- [10] Fernando Diaz, Bhaskar Mitra, Michael D Ekstrand, Asia J Biega, and Ben Carterette. 2020. Evaluating Stochastic Rankings with Expected Exposure. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 275–284.
- [11] Michael D. Ekstrand, Robin Burke, and Fernando Diaz. 2019. Fairness and Discrimination in Retrieval and Recommendation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1403–1404.
- [12] Guglielmo Faggioli, Mirko Polato, and Fabio Aioli. 2020. Recency Aware Collaborative Filtering for Next Basket Recommendation. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*. 80–87.
- [13] Yingqiang Ge, Shuchang Liu, Ruoyuan Gao, Yikun Xian, Yunqi Li, Xiangyu Zhao, Changhua Pei, Fei Sun, Junfeng Ge, Wenwu Ou, and Yongfeng Zhang. 2021. Towards Long-Term Fairness in Recommendation. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. 445–453.
- [14] Haoji Hu and Xiangnan He. 2019. Sets2Sets: Learning from Sequential Sets with Neural Networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1491–1499.
- [15] 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.
- [16] Dietmar Jannach. 2022. Multi-Objective Recommender Systems: Survey and Challenges. In *MORS workshop held in conjunction with the 16th ACM Conference on Recommender Systems (RecSys), 2022*.
- [17] Ori Katz, Oren Barkan, Noam Koenigstein, and Nir Zabari. 2022. Learning to Ride a Buy-Cycle: A Hyper-Convolutional Model for Next Basket Repurchase Recommendation. In *Proceedings of the 16th ACM Conference on Recommender Systems*. 316–326.
- [18] Anuj Kumar and Kartik Hosanagar. 2019. Measuring the Value of Recommendation Links on Product Demand. *Information Systems Research* 30, 3 (2019), 819–838.
- [19] Duc-Trong Le, Hady W Lauw, and Yuan Fang. 2019. Correlation-sensitive Next-basket Recommendation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. 2808–2814.
- [20] Youfang Leng, Li Yu, Jie Xiong, and Guanyu Xu. 2020. Recurrent Convolution Basket Map for Diversity Next-Basket Recommendation. In *International Conference on Database Systems for Advanced Applications*. 638–653.
- [21] Ming Li, Moshdeh Ariannezhad, Andrew Yates, and Maarten de Rijke. 2023. Masked and Swapped Sequence Modeling for Next Novel Basket Recommendation in Grocery Shopping. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 35–46.
- [22] Ming Li, Moshdeh Ariannezhad, Andrew Yates, and Maarten De Rijke. 2023. Who Will Purchase this Item Next? Reverse Next Period Recommendation in Grocery Shopping. *ACM Transactions on Recommender Systems* 1, 2 (2023), 1–32.
- [23] Ming Li, Jin Huang, and Maarten de Rijke. 2023. Repetition and Exploration in Offline Reinforcement Learning-based Recommendations. In *4th Workshop on Deep Reinforcement Learning for Information Retrieval at CIKM 2023*. ACM.
- [24] Ming Li, Sami Jullien, Moshdeh Ariannezhad, and Maarten de Rijke. 2023. A next basket recommendation reality check. *ACM Transactions on Information Systems* 41, 4 (2023), 1–29.
- [25] Ming Li, Ali Vardasbi, Andrew Yates, and Maarten de Rijke. 2023. Repetition and Exploration in Sequential Recommendation. In *SIGIR 2023: 46th international ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2532–2541.
- [26] Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, Juntao Tan, Shuchang Liu, and Yongfeng Zhang. 2022. Fairness in recommendation: A survey. *ACM Transactions on Intelligent Systems and Technology* (2022).
- [27] Yile Liang, Tiejun Qian, Qing Li, and Hongzhi Yin. 2021. Enhancing Domain-level and User-level Adaptivity in Diversified Recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 747–756.
- [28] Yuanna Liu, Ming Li, Moshdeh Ariannezhad, Masoud Mansoury, Mohammad Alianejadi, and Maarten de Rijke. 2024. Measuring Item Fairness in Next Basket Recommendation: A Reproducibility Study. In *ECIR 2024: 46th European Conference on Information Retrieval*. Springer, 210–225.
- [29] Malte Ludewig and Dietmar Jannach. 2018. Evaluation of Session-based Recommendation Algorithms. *User Modeling and User-Adapted Interaction* 28 (2018), 331–390.
- [30] Marco Morik, Ashudeep Singh, Jessica Hong, and Thorsten Joachims. 2020. Controlling Fairness and Bias in Dynamic Learning-to-Rank. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 429–438.
- [31] Mohammadmehdi Naghiaei, Hossein A. Rahmani, and Yashar Deldjoo. 2022. CPFair: Personalized Consumer and Producer Fairness Re-Ranking for Recommender Systems. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 770–779.
- [32] Sergey Naumov, Marina Ananyeva, Oleg Lashinin, Sergey Kolesnikov, and Dmitry I Ignatov. 2023. Time-Dependent Next-Basket Recommendations. In *European Conference on Information Retrieval*. Springer, 502–511.
- [33] OECD. 2020. E-commerce in the Time of COVID-19. <https://www.oecd.org/coronavirus/policy-responses/e-commerce-in-the-time-of-covid-19-3a2b78e8/>.
- [34] Yuqi Qin, Pengfei Wang, and Chenliang Li. 2021. The World is Binary: Contrastive Learning for Denoising Next Basket Recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 859–868.
- [35] Massimo Quadrana, Paolo Cremonesi, and Dietmar Jannach. 2018. Sequence-Aware Recommender Systems. *Comput. Surveys* 51, 4 (2018), 1–36.
- [36] Amifa Raj and Michael D Ekstrand. 2022. Measuring fairness in ranked results: An analytical and empirical comparison. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 726–736.
- [37] 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.
- [38] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of Exposure in Rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2219–2228.
- [39] Leilei Sun, Yansong Bai, Bowen Du, Chuanren Liu, Hui Xiong, and Weifeng Lv. 2020. Dual Sequential Network for Temporal Sets Prediction. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1439–1448.
- [40] Mengting Wan, Di Wang, Jie Liu, Paul Bennett, and Julian McAuley. 2018. Representing and Recommending Shopping Baskets with Complementarity, Compatibility and Loyalty. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 1133–1142.
- [41] Pengfei Wang, Jiafeng Guo, Yanyan Lan, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2015. Learning Hierarchical Representation Model for Next Basket Recommendation. In *Proceedings of the 38th International ACM SIGIR conference on Research and Development in Information Retrieval*. 403–412.
- [42] Pengfei Wang, Yongfeng Zhang, Shuzi Niu, and Jiafeng Guo. 2019. Modeling Temporal Dynamics of Users' Purchase Behaviors for Next Basket Prediction. *Journal of Computer Science and Technology* 34, 6 (2019), 1230–1240.
- [43] Shoujin Wang, Liang Hu, Yan Wang, Quan Z Sheng, Mehmet Orgun, and Longbing Cao. 2019. Modeling Multi-purpose Sessions for Next-item Recommendations via Mixture-channel Purpose Routing Networks. In *International Joint Conference on Artificial Intelligence*.

- [44] Shoujin Wang, Liang Hu, Yan Wang, Quan Z Sheng, Mehmet Orgun, and Longbing Cao. 2020. Intention Nets: Psychology-inspired User Choice behavior Modeling for Next-basket Prediction. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*. 6259–6266.
- [45] Patti Williams, Iris W. Hung, Anirban Mukhopadhyay, Rik Pieters, Xinyue Zhou, Tim Wildschut, Constantine Sedikides, Kan Shi, Cong Feng, Cassie Mogilner, Jennifer Aaker, Sepandar D. Kamvar, Fabrizio Di Muro, and Kyle B. Murray. 2014. Emotions and Consumer Behavior. *Journal of Consumer Research* 40, 5 (2014), viii–xi.
- [46] Haolun Wu, Bhaskar Mitra, Chen Ma, Fernando Diaz, and Xue Liu. 2022. Joint Multisided Exposure Fairness for Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 703–714.
- [47] Yao Wu, Jian Cao, Guandong Xu, and Yudong Tan. 2021. TFROM: A Two-Sided Fairness-Aware Recommendation Model for Both Customers and Providers. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1013–1022.
- [48] Qing Yin, Hui Fang, Zhu Sun, and Yew-Soon Ong. 2023. Understanding Diversity in Session-Based Recommendation. *ACM Transactions on Information Systems* 42, 1 (2023), 1–34.
- [49] 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.
- [50] Le Yu, Leilei Sun, Bowen Du, Chuanren Liu, Hui Xiong, and Weifeng Lv. 2020. Predicting Temporal Sets with Deep Neural Networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1083–1091.
- [51] Meike Zehlike and Carlos Castillo. 2020. Reducing Disparate Exposure in Ranking: A Learning To Rank Approach. In *Proceedings of The Web Conference 2020*. 2849–2855.
- [52] Mi Zhang and Neil Hurley. 2008. Avoiding Monotony: Improving the Diversity of Recommendation Lists. In *Proceedings of the 2008 ACM Conference on Recommender Systems*. 123–130.
- [53] Yuying Zhao, Yu Wang, Yunchao Liu, Xueqi Cheng, Charu Aggarwal, and Tyler Derr. 2023. Fairness and Diversity in Recommender Systems: A Survey. *arXiv preprint arXiv:2307.04644* (2023).
- [54] Yu Zheng, Chen Gao, Liang Chen, Depeng Jin, and Yong Li. 2021. DGCN: Diversified Recommendation with Graph Convolutional Networks. In *Proceedings of the Web Conference 2021*. 401–412.