Investigating the Robustness of Sequential Recommender Systems Against Training Data Perturbations [*]

Filippo Betello[1] , Federico Siciliano[1] ,
Pushkar Mishra[2] , and Fabrizio Silvestri[1]

[1] Sapienza University of Rome, Rome, Italy
{betello, siciliano, fsilvestri}@diag.uniroma1.it
[2] AI at Meta, London, UK pushkarmishra@meta.com

Abstract. Sequential Recommender Systems (SRSs) are widely employed to model user behavior over time. However, their robustness in the face of perturbations in training data remains a largely understudied yet critical issue. A fundamental challenge emerges in previous studies aimed at assessing the robustness of SRSs: the Rank-Biased Overlap (RBO) similarity is not particularly suited for this task as it is designed for infinite rankings of items and thus shows limitations in real-world scenarios. For instance, it fails to achieve a perfect score of 1 for two identical finite-length rankings. To address this challenge, we introduce a novel contribution: Finite Rank-Biased Overlap (FRBO), an enhanced similarity tailored explicitly for finite rankings. This innovation facilitates a more intuitive evaluation in practical settings. In pursuit of our goal, we empirically investigate the impact of removing items at different positions within a temporally ordered sequence. We evaluate two distinct SRS models across multiple datasets, measuring their performance using metrics such as Normalized Discounted Cumulative Gain (NDCG) and Rank List Sensitivity. Our results demonstrate that removing items at the end of the sequence has a statistically significant impact on performance, with NDCG decreasing up to 60%. Conversely, removing items from the beginning or middle has no significant effect. These findings underscore the criticality of the position of perturbed items in the training data. As we spotlight the vulnerabilities inherent in current SRSs, we fervently advocate for intensified research efforts to fortify their robustness against adversarial perturbations.

Keywords: Recommender Systems · Evaluation of Recommender Systems · Model Stability · Input Data Perturbation

1   Introduction

Recommender systems have become ubiquitous in our daily lives [1], playing a key role in helping users navigate the vast amounts of information available online. Thanks to the global spread of e-commerce services, social media and streaming platforms, recommender systems have become increasingly important for personalized content delivery and user engagement [45]. In recent years, Sequential Recommender Systems (SRSs) have emerged as a popular approach to modeling user behavior over time [29], leveraging the temporal dependencies in users' interaction sequences to make more accurate predictions.

However, despite their success, the robustness of SRSs against perturbations in the training data remains an open research question [22]. In real-world scenarios, disruptions may occur when users employ different services for the same purpose. Data becomes fragmented and divided between a service provider and its competitors in such cases. Nevertheless, the provider must train a recommender system with such incomplete data while ensuring robustness to perturbations. This challenge is accentuated when we scrutinize previous attempts [24] to assess the robustness of SRSs: Rank-Biased Overlap (RBO) [41], designed explicitly for infinite lists, reveals limitations when applied to real-world scenarios.

Our experiments revolve around the following research questions:

- **RQ1**: Do changes in the training seed heavily impact rankings?
- **RQ2**: How does the type of removal influence the model's performance?
- **RQ3**: Does more item removed significantly decrease in performance?

The contribution of this study is two-fold. Firstly, we propose the novel Finite Rank-Biased Overlap (FRBO) measure. Unlike RBO, which is only suited for infinite rankings, FRBO is specifically designed to assess the robustness of SRSs within finite ranking scenarios, aligning seamlessly with real-world settings. Secondly, we empirically assess the impact of item removal from user interaction sequences on SRS performance. Our investigation shows that the most recent user interaction sequence items are critical for accurate recommendation performance. When these items are removed, there is a significant drop in all metrics.

2   Related works

2.1   Sequential Recommender Systems

SRSs are algorithms that leverage a user's past interactions with items to make personalized recommendations over time and they have been widely used in various applications, including e-commerce [31,18], social media [14,3], and music streaming platforms [32,33,2]. Compared to traditional recommender systems, SRSs consider the order and timing of these interactions, allowing for more accurate predictions of a user's preferences and behaviors [39].

Various techniques have been proposed to implement SRSs. Initially, Markov Chain models were employed in Sequential Recommendation [13,12], but they

struggle to capture complex dependencies in long-term sequences. In recent years, Recurrent Neural Networks (RNNs) have emerged as one of the most promising approaches in this field [9,16,30]. These methods encode the users' historical preferences into a vector that gets updated at each time step and is used to predict the next item in the sequence. Despite their success, RNNs may face challenges dealing with long-term dependencies and generating diverse recommendations. Another approach comes from the use of the attention mechanism [38]: two different examples are SASRec [20] and BERT4Rec [36] architectures. This method dynamically weighs the importance of different sequence parts to better capture the important features and improve the prediction accuracy. Recently, Graph Neural Networks have become popular in the field of recommendations system [43,28], especially in the sequential domain [7,11].

2.2   Robustness in Sequential Recommender Systems

Robustness is an important aspect of SRSs as they are vulnerable to noisy and incomplete data. Surveys on the robustness of recommender systems [17,6] discussed the challenges in developing robust recommender systems and presented various techniques for improving the robustness of recommender systems.

   Many tools have been developed to test the robustness of different algorithms: [25], provided both a formalisation of the topic and a framework; the same was done more recently by [26], who developed a toolkit called RGRecSys which provides a unified framework for evaluating the robustness of SRSs.

   Few recent work has focused on studying the problem and trying to increase robustness: [37] focuses on robustness in the sense of training stability while [24] investigated the robustness of SRSs to interaction perturbations. They showed that even small perturbations in user-item interactions can lead to significant changes in the recommendations. They proposed Rank List Sensitivity (RLS), a measure of the stability of rankings produced by recommender systems.

   Our work expands these, making a more accurate investigation of the effect different types of perturbation can have on the models' performance. While [24] perturbs a single interaction in the whole dataset, we perturb the sequences of all users and analyze the performance as the number of perturbations changes. We also provide a theoretical contribution in a new sensitivity evaluation measure for finite rankings, presented in Sec. 3.2.

3   Methodology

3.1   Setting

In Sequential Recommendation, each user u is represented by a temporally ordered sequence of items $S_u = (I_1, I_2, ..., I_j, ..., I_{L_u-1}, I_{L_u})$ with which it has interacted, where L_u is the length of the sequence for user u.

   User-object interactions in real-world scenarios are often fragmented across services, resulting in a lack of comprehensive data. For example, in the domains

of movies and TV shows, a single user may interact with content on TV, in a movie theater, or across multiple streaming platforms. To mimic this real-world scenario in our training data perturbations, we considered three different cases, each removing n items at a specific position in the sequence:

- **Beginning**: $S_u = (I_{n+1}, \ldots, I_{L_u-1})$. This represents a user who signs up for a new service, so all his past interactions, i.e., those at the beginning of the complete sequence, were performed on other services.
- **Middle**: $S_u = (I_1, \ldots, I_{\lfloor \frac{L_u-1-n}{2} \rfloor}, I_{\lfloor \frac{L_u-1+n}{2} \rfloor} \ldots, I_{L_u-1})$. This represents a user who takes a break from using the service for a certain period and resumes using it. Still, any interactions they had during the considered period are not available to the service provider.
- **End**: $S_u = (I_1, \ldots, I_{L_u-1-N})$. This represents a user who has stopped using the service, so the service provider loses all the subsequent user interactions. The service provider still has an interest in winning the user back through their platform or other means, such as advertising. Thus, it is essential to have a robust model to continue providing relevant items to the user.

with $n \in \{1, 2, ..., 10\}$. In practice, the data is first separated into training, validation and test set (always composed by I_{L_u}). Subsequently, only the training data are perturbed, with a methodology dependent on the scenario considered and the model is then trained on these. The models, trained on data perturbed in a different manner, are therefore always tested on the same data.

3.2   Metrics

To evaluate the performance of the models, we employ traditional evaluation metrics used for Sequential Recommendation: Precision, Recall, MRR and NDCG.

Moreover, to investigate the stability of the recommendation models, we employ the Rank List Sensitivity (RLS) [24]: it compares two lists of rankings $\mathcal{X}$ and $\mathcal{Y}$, one derived from the model trained under standard conditions and the other derived from the model trained with perturbed data.

Therefore, having these two rankings, and a similarity function sim between them, we can formalize the RLS measure as:

$$\mathbf{RLS} = \frac{1}{|\mathcal{X}|} \sum_{k=1}^{|\mathcal{X}|} \text{sim}(R^{X_k}, R^{Y_k}) \tag{1}$$

where X_k and Y_k represent the k-th ranking inside $\mathcal{X}$ and $\mathcal{Y}$ respectively. RLS's similarity measure can be chosen from two possible options:

- **Jaccard Similarity (JAC)** [19] is a normalized measure of the similarity of the contents of two sets. A model is stable if its Jaccard score is close to 1.

$$\mathbf{JAC(X, Y)} = \frac{|X \cap Y|}{|X \cup Y|} \tag{2}$$

- **Rank-Biased Overlap (RBO)** [41] measures the similarity of orderings between two rank lists. Higher values indicate that the items in the two lists are arranged similarly:

$$\mathbf{RBO(X, Y)} = (1 - p) \sum_{d=1}^{+\infty} p^{d-1} \frac{|X[1:d] \cap Y[1:d]|}{d} \tag{3}$$

In the domain of recommendation systems, it is customary to compute metrics using finite-length rankings, typically denoted by appending "@k" to the metric's name, such as NDCG@k. While traditional metrics (e.g. NDCG, MRR, etc.) readily adapt to finite-length rankings, maintaining their core meaning, the same behaviour does not extend to RLS when employing RBO. The reason lies in Equation 3, which exhibits a notable limitation: it fails to converge to one, even when applied to identical finite-length lists. To overcome this limitation, we introduce the Finite Rank-Biased Overlap (FRBO) similarity, denoted as FRBO@k, which represents a novel formulation engineered to ensure convergence to a value of 1 for identical lists and a value of 0 for entirely dissimilar lists.

Theorem 1. *Given a set of items $I = \{I_1, ..., I_{N_I}\}$, two rankings $X = (x_1, ..., x_k)$ and $Y = (y_1, ..., y_k)$, such that $x_i, y_i \in I$, and $k \in \mathbb{N}^+$*

$$\mathbf{FRBO(X, Y)@k} = \frac{\mathrm{RBO(X,Y)@k} - min_{X,Y}\mathrm{RBO@k}}{max_{X,Y}\mathrm{RBO@k} - min_{X,Y}\mathrm{RBO@k}} \tag{4}$$

$$\min_{X,Y} \mathrm{FRBO(X,Y)@k} = 0, \ \max_{X,Y} \mathrm{FRBO(X,Y)@k} = 1$$

Proof. This follows from the fact that given a function $f : A \to [a, b]$ and another function $g = \frac{f-a}{b-a}$, then $g : A \to [0, 1]$, where A is any set. □

To normalize RBO, we need to identify its minimum and maximum values when the summation is carried out up to the top-k elements of the ranking, simultaneously proving that these values are not naturally constrained to be 0 and 1.

Lemma 1. *Given a set of items $I = \{I_1, ..., I_{N_I}\}$, two rankings $X = (x_1, ..., x_k)$ and $Y = (y_1, ..., y_k)$, such that $x_i, y_i \in I$, and $k \in \mathbb{N}^+$, the following holds:*

$$\min_{X,Y} \mathrm{RBO@k} = \begin{cases} 0, & \text{if } k \leq \lfloor \frac{N_I}{2} \rfloor \\ (1-p)\left(2\frac{p^{\lfloor \frac{N_I}{2} \rfloor} - p^{N_I}}{1-p} - N_I \ell\right) & \text{otherwise} \end{cases} \tag{5}$$

where $\mathrm{RBO(X,Y)@k} = (1-p) \sum_{d=1}^{k} p^{d-1} \frac{|X[1:d] \cap Y[1:d]|}{d}$

and $\ell = p^{\lfloor \frac{N_I}{2} \rfloor} \Phi(p, 1, \lfloor \frac{N_I}{2} \rfloor + 1) - p^{N_I} \Phi(p, 1, N_I + 1)$

Proof. For the first part, it suffices to consider two rankings X, Y that share no common elements, so that $|X[1:d] \cap Y[1:d]| = 0$. This leads to:

$$\min_{X,Y} \mathrm{RBO@k} = (1-p) \sum_{d=1}^{k} p^{d-1} \frac{0}{d} = 0$$

However, since the number of items I is not infinite, if $k > \lfloor \frac{N_I}{2} \rfloor$, at least one element must necessarily be in common between the two rankings:

$$|X[1:k] \cap Y[1:k]| = 0 \iff |X[1:k] \cup Y[1:k]| = 2k \leq N_I \iff k \leq \lfloor \frac{N_I}{2} \rfloor$$

Consequently, the similarity can't assume a value of 0 if $k > \lfloor \frac{N_I}{2} \rfloor$.

Given that $\mathrm{RBO(X,Y)@}k+1 \geq \mathrm{RBO(X,Y)@}k \; \forall k \in \mathbb{N}^+$, similarity it's minimized when intersections between the two rankings occur as far down the list as possible. When $d > \lfloor \frac{N_I}{2} \rfloor$, there are at least $2d - N_I$ intersections, because items in d-th position in one ranking are necessarily contained in the other.

$$\min_{X,Y} \mathrm{RBO@}k = (1-p)\left(0 + \sum_{d=\lfloor \frac{N_I}{2} \rfloor+1}^{N_I} p^{d-1} \frac{2d - N_I}{d} \right)$$

$$= (1-p)\left(2 \sum_{d=\lfloor \frac{N_I}{2} \rfloor+1}^{N_I} p^{d-1} - N_I \sum_{d=\lfloor \frac{N_I}{2} \rfloor+1}^{N_I} \frac{p^{d-1}}{d} \right)$$

The first series can be regarded as a finite geometric series, for which we can apply the formula for the sum of a geometric series: $\sum_{n=1}^{k} ar^{n-1} = \frac{a(1-r^k)}{1-r}$:

$$\sum_{d=\lfloor \frac{N_I}{2} \rfloor+1}^{N_I} p^{d-1} = p^{\lfloor \frac{N_I}{2} \rfloor} \sum_{d=1}^{N_I - \lfloor \frac{N_I}{2} \rfloor} p^{d-1} = \frac{p^{\lfloor \frac{N_I}{2} \rfloor} - p^{N_I}}{1-p}$$

Using Lerch transcendent function $\Phi(z,s,\alpha) = \sum_{n=0}^{+\infty} \frac{z^n}{(n+\alpha)^s}$ in second series:

$$\sum_{d=\lfloor \frac{N_I}{2} \rfloor+1}^{N_I} \frac{p^{d-1}}{d} = \sum_{d=0}^{N_I - \lfloor \frac{N_I}{2} \rfloor - 1} \frac{p^{d+\lfloor \frac{N_I}{2} \rfloor}}{d + \lfloor \frac{N_I}{2} \rfloor + 1}$$

$$= \sum_{d=0}^{+\infty} \frac{p^{d+\lfloor \frac{N_I}{2} \rfloor}}{d + \lfloor \frac{N_I}{2} \rfloor + 1} - \sum_{d=N_I - \lfloor \frac{N_I}{2} \rfloor}^{+\infty} \frac{p^{d+\lfloor \frac{N_I}{2} \rfloor}}{d + \lfloor \frac{N_I}{2} \rfloor + 1}$$

$$= p^{\lfloor \frac{N_I}{2} \rfloor} \Phi(p, 1, \lfloor \frac{N_I}{2} \rfloor + 1) - p^{N_I} \Phi(p, 1, N_I + 1)$$

$\square$

Lemma 2. *Given a set of items* $I = \{I_1, ..., I_{N_I}\}$*, two rankings* $X = (x_1, ..., x_k)$ *and* $Y = (y_1, ..., y_k)$*, such that* $x_i, y_i \in I$*, and* $k \in \mathbb{N}^+$*, the following holds:*

$$\max_{X,Y} \mathrm{RBO@}k = 1 - p^k \tag{6}$$

Proof. RBO@k reaches its maximum value when the two rankings are identical:

$$|X[1:d] \cap Y[1:d]| = d \ \ \forall d \in \{1, \ldots, k\}$$

Referring to the result for the geometric series, we can compute:

$$max_{X,Y} \text{RBO@k} = (1-p) \sum_{d=1}^{k} p^{d-1} \frac{d}{d} = (1-p) \sum_{d=1}^{k} p^{d-1} = 1 - p^k$$

□

Corollary 1. *Given a set of items* $I = \{I_1, ..., I_{N_I}\}$ *and* $k \in \mathbb{N}^+$:

$$\min_{X,Y} \text{RBO(X, Y)@k} \neq 0, \ \max_{X,Y} \text{RBO(X, Y)@k} \neq 1$$

$$\exists X, Y \in \{(x_1, \ldots, x_k) | x_i \in I \wedge x_i \neq x_j \forall i \neq j\}$$

$$s.t. \ \text{RBO(X, Y)@k} \neq \text{FRBO(X, Y)@k}$$

Proof. This follows from Theorems 1 and 2.   □

It's worth considering that in real-world scenarios, the number of possible items N_I is significantly larger compared to the length k of the rankings used to compute the metrics, i.e., $N_I >> k$. In this context, we can safely omit the minimum value from Eq. 4, resulting in:

$$\mathbf{FRBO(X, Y)@k} = \frac{\text{RBO(X, Y)@k}}{\max_{X,Y} \text{RBO@k}} = \frac{1-p}{1-p^k} \sum_{d=1}^{k} p^{d-1} \frac{|X[1:d] \cap Y[1:d]|}{d}$$

In this section, we have shown that RBO is not an adequate similarity score when dealing with finite-length rankings. So, we have derived expressions that quantify the minimum and maximum values of RBO, allowing us to compute a normalized version of RBO.

Kendall's Tau [21] assumes that two rankings contain precisely the same items. However, this assumption may not hold for finite top-k ranked lists. In addition, average overlap [42,10] has a peculiar property of monotonicity in depth, where greater agreement with a deeper ranking does not necessarily lead to a higher score, and less agreement does not necessarily lead to a lower score [41].

4 Experiments

4.1 Datasets

We use four different datasets:

Table 1: Dataset statistics after preprocessing

Dataset	Users	Items	Interactions	Average $\frac{\text{Actions}}{\text{User}}$	Median $\frac{\text{Actions}}{\text{User}}$
MovieLens 1M	6040	3952	1M	165	96
MovieLens 100K	943	1682	100K	106	65
Foursquare Tokyo	2293	61858	537703	250	153
Foursquare New York	1083	38333	227428	210	173

MovieLens [15] → This benchmark dataset is often used to test recommender systems. In this work, we use the 100K version and 1M version.

Foursquare [44] → This dataset contains check-ins from New York City and Tokyo collected over a period of approximately ten months.

The statistics for all the datasets are shown in Table 1.

We select datasets widely used in the literature and with a high number of interactions per user. The limitation in dataset selection arises from our intention to assess the robustness against the removal of up to 10 elements. Therefore, the dataset must satisfy the following constraint: $L_u > 10 \quad \forall u \in U$, where L_u is the number of interactions of user u, i.e. the length of the sequence S_u of interactions, and U is the set of all users in the dataset. If the condition is not met, we delete all the items for a user with less than ten interactions. In this case, we cannot train the model on this particular user. We have, thus, decided to exclude datasets such as Amazon [23] for they cannot meet the previous criteria.

4.2   Architectures

In our study, we use two different architectures to validate the results:

- **SASRec** [20] uses self-attention processes to determine the importance of each interaction between the user and the item.
- **GRU4Rec** [16] is a recurrent neural network architecture that uses gated recurrent units (GRUs) [8] to improve the accuracy of the prediction.

We choose to use these two models because both have demonstrated excellent performance in several benchmarks and have been widely cited in the literature. Furthermore, as one model employs attention while the other utilizes RNN, their network functioning differs, which makes evaluating their behavior under training perturbations particularly interesting. We use the models' implementation provided by the RecBole Python library [46], with their default hyperparameters.

4.3   Experimental Setup

All the experiments are performed on a single NVIDIA RTX 3090. The batch size is fixed to 4096. Adam optimizer is used with a fixed learning rate of $5*10^{-4}$.

Table 2: **Variation of metrics between two seeds.** Metrics for the 4 datasets considered for GRU4Rec and SASRec. For Precision (Prec.), Recall, MRR and NDCG, it is shown the percentage variation between the obtained performance using two different initialization seeds for the models. For each metric, the value corresponding to the dataset where a model is less robust is highlighted in **bold**. For each dataset, the value corresponding to the model that is the least robust of the two given a metric is underlined.

	SASRec						GRU4Rec					
	Prec.	Recall	MRR	NDCG	FRBO	JAC	Prec.	Recall	MRR	NDCG	FRBO	JAC
ML 100k	0.5%	0.5%	0.5%	<u>0.3%</u>	.466	.489	**<u>1.7%</u>**	**<u>1.7%</u>**	1.2%	0.1%	<u>.337</u>	<u>.413</u>
ML 1M	<u>0.3%</u>	<u>0.4%</u>	**0.5%**	**0.5%**	.549	.569	0.2%	0.2%	<u>3.5%</u>	2.5%	<u>.311</u>	<u>.347</u>
FS NYC	<u>0.9%</u>	<u>0.9%</u>	0.0%	<u>0.3%</u>	**.398**	**.273**	0.1%	0.1%	<u>0.6%</u>	0.1%	**<u>.110</u>**	**<u>.083</u>**
FS TKY	**<u>1.4%</u>**	**<u>1.4%</u>**	0.4%	0.02%	.418	.267	0.8%	0.8%	**<u>4.7%</u>**	**<u>3.4%</u>**	<u>.210</u>	<u>.165</u>

The number of epochs is set to 300, but in order to avoid overfitting, we stop the training when the NDCG@20 does not improve for 50 epochs. The average duration of each run is 1.5 hours. The RecBole library was utilized for conducting all the experiments, encompassing data preprocessing, model configuration, training, and testing. This comprehensive library ensures the reproducibility of the entire process. All the evaluation metrics are calculated with a cut-off K of 20. To validate the effective degradation in performance and rankings similarity, we employed the paired Student's t-test [35], after testing normality of distributions with Shapiro–Wilk test [34], and a significance level of 10^{-3}.

5 Results

5.1 Intrinsic Models Instability (RQ1)

To measure the inherent robustness of the models, i.e. in the baseline case (without removal of items), we train the model twice using different initialization seeds and compute the percentage discrepancy between the Precision, Recall, MRR and NDCG obtained by the two rankings. The results are shown in Table 2: it can be seen that in general the discrepancy is negligible, almost always less than 1%. On the other hand, the two RLS, calculated using FRBO and Jaccard respectively, show us the similarity between the two rankings produced with different initialization seeds. The results deviate significantly from the ideal value of 1, indicating considerably different rankings. These combined results indicate to us that the models converge to an adequate performance beyond the initialization seed, but that the actual rankings produced are heavily influenced by it. The **bold** represent the dataset with the least robust result for each metric. No datasets stands out significantly, yet Foursquare Tokyo seems to give more problems regarding standard evaluation metrics, while for Foursquare New York

City it seems more difficult to produce stable rankings. The underlined values instead compare, for each metric and dataset, which of the two models is the least robust. If we consider metrics that do not consider the position of the positive item in the ranking, i.e. Precision and Recall, GRU4Rec seems more robust than SASRec. If, on the other hand, we look at metrics that penalize relevant items in positions too low in the ranking, we see that the opposite happens. This suggests to us that GRU4Rec can return a better set of results, but in a less relevant order than SASRec does. As proof of this, if we check the RLS values, we see that GRU4Rec is always the least robust model as the initialization seed changes.

5.2   Comparison of the position of removal (RQ2)

Table 3 compares performance and stability when ten items are removed from the sequence versus retaining all items (reference value) in the training set with a consistent initialization seed. It's observed that discarding items from the beginning or middle of the sequence does not significantly impair the model's performance; only a minor decline is noted, potentially attributable to the marginally reduced volume of total training data. Instead, it can be observed how removing items from the end of the sequence leads to a drastic reduction in metrics: in the case of SASRec applied to the MovieLens 1M dataset, the NDCG more than halves. Finally, we can see how the difference between the three settings, although maintaining the same trend, is less marked for GRU4Rec applied to the Foursquare NYC dataset. This may be due to the generally higher performance of GRU4Rec [20]. A more in-depth analysis is presented in Section 5.4.

Table 3 also shows the RLS values, computed using FRBO and Jaccard similarity, on the same model-dataset pairs: removing items at the end of the sequence leads to considerable variation in the rankings produced by the models. The values approach 0, meaning that the produced rankings share almost no items. Our results are in contrast to those of [24], which instead claim that an initial perturbation of a user sequence leads to a higher impact on the RLS. However, their experimental setting is different than ours, as explained in Section 2.

5.3   Effect of the number of elements removed (RQ3)

As we discussed in the previous section, removing elements that are at the beginning or in the middle of the temporally ordered sequence has no effect on performance. This is also confirmed by Figure 1, where we can also see, however, that for the above-mentioned cases there is no variation as the number of removed elements increases. The deviation of the RLS displayed in Figure 1d will be analyzed in more detail in Section 5.4. For the remaining setting, the one where we remove items at the end of the sequence, the effect of the number of items removed is evident: the metrics drop drastically as the number of items removed increases. This result holds true for both models considered and for all four datasets tested.

Table 3: **Variations in metrics for ten-item removal.** Metrics for the 3 scenarios considered for SASRec on ML-1M and GRU4Rec on FS-NYC. For Precision, Recall, MRR, and NDCG, it is shown the percentage variation between removing ten items and the reference value. For each metric, in **bold** it is highlighted the value representing the less robust model. † indicates a statistically significant result.

Model	Removal	Prec.	Recall	MRR	NDCG	FRBO	JAC
SASRec ML-1M	Beginning	-0.23%	-0.23%	-0.15%	-0.07%	$.399^{\dagger}$	$.368^{\dagger}$
	Middle	-0.35%	-0.29%	-1.09%	-0.73%	$.385^{\dagger}$	$.356^{\dagger}$
	End	$\mathbf{-15.5\%^{\dagger}}$	$\mathbf{-15.3\%^{\dagger}}$	$\mathbf{-45.9\%^{\dagger}}$	$\mathbf{-56.0\%^{\dagger}}$	$\mathbf{.080^{\dagger}}$	$\mathbf{.106^{\dagger}}$
GRU4Rec FS-NYC	Beginning	-0.23%	-0.23%	-1.47%	-0.94%	$.105^{\dagger}$	$.075^{\dagger}$
	Middle	-0.93%	-0.93%	-0.18%	-0.44%	$.110^{\dagger}$	$.074^{\dagger}$
	End	$\mathbf{-4.92\%^{\dagger}}$	$\mathbf{-4.86\%^{\dagger}}$	$\mathbf{-8.42\%^{\dagger}}$	$\mathbf{-7.39\%^{\dagger}}$	$\mathbf{.089^{\dagger}}$	$\mathbf{.062^{\dagger}}$
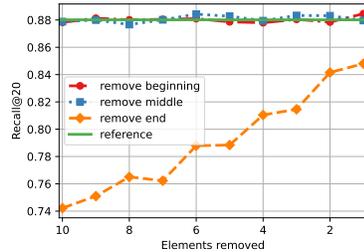
(a) HR@20 SASRec ML-100k
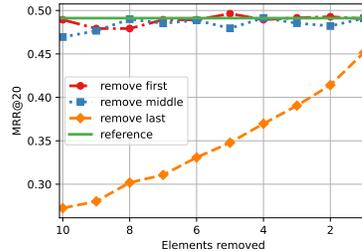
(b) MRR GRU4Rec ML-1M
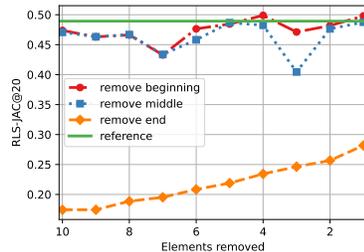
(c) RLS-JAC@20 SASRec ML-100k
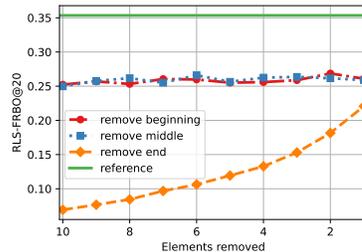
(d) RLS-FRBO@20 GRU4Rec ML-1M

Fig. 1: Plots of various metrics for the ML-100k and ML-1M datasets as the number of removed elements increases. The baseline is shown as a horizontal solid line, while dashed lines show the metrics as the number of items removed changes for the three scenarios considered.

5.4   Differences between the datasets

Figures 2 and 3 show the performance for the three different settings for the SAS-Rec model applied to all datasets. From Figures 2a, 2b, 2c, 2d we see that the
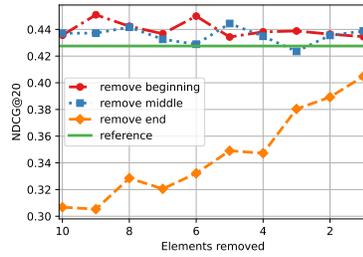
(a) NDCG@20 SASRec ML-100k
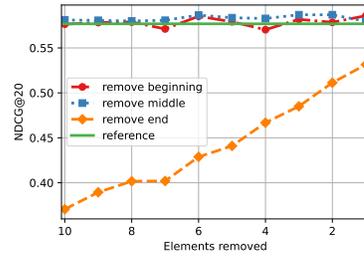
(b) NDCG@20 SASRec ML-1M
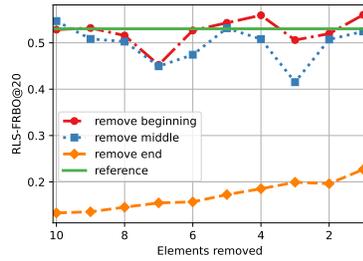
(c) FRBO@20 SASRec ML-100k
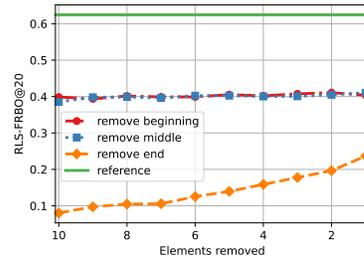
(d) FRBO@20 SASRec ML-1M

Fig. 2: Plots of NDCG and FRBO for SASRec on the ML-100K and ML-1M datasets. The baseline is shown as a horizontal solid line, while dashed lines show the metrics as the number of items removed changes for the three scenarios considered.

downward trend of the metric when removing items at the end of the sequence is a characteristic of the MovieLens dataset: both NDCG@20 and RLS-FRBO@20 show a decrease when increasing the number of removed items. We hypothesize that this is happening because the average number of actions per user and the number of items (see Table 1) are not that large compared to the number of items removed. On the other hand, the Foursquare datasets (3a, 3b, 3c, 3d) do not suffer major performance degradation, probably due to the higher average number of actions per user and the number of items (see Table 1) than Movie-Lens. In addition to this, and probably for the same motivation, the degradation of the RLS is lower with respect to that displayed in the MovieLens datasets. Finally, it is interesting to note that on MovieLens 1M, there is a consistent performance degradation even when elements at the beginning and in the middle of the temporally ordered sequence are removed. This can be observed as a small decrease in the NDCG@20 (Figure 2b), but a sharp decrease in the value of the RLS-FRBO (Figures 1d, 2d). This means that even if the model performs approximately the same, the rankings produced vary greatly. The cause may be the fact that the MovieLens 1M dataset, among those considered, has the largest number of users and interactions.

(a) NDCG@20 SASRec FS-TKY

(b) NDCG@20 SASRec FS-NYC

(c) FRBO@20 SASRec FS-TKY

(d) FRBO@20 SASRec FS-NYC
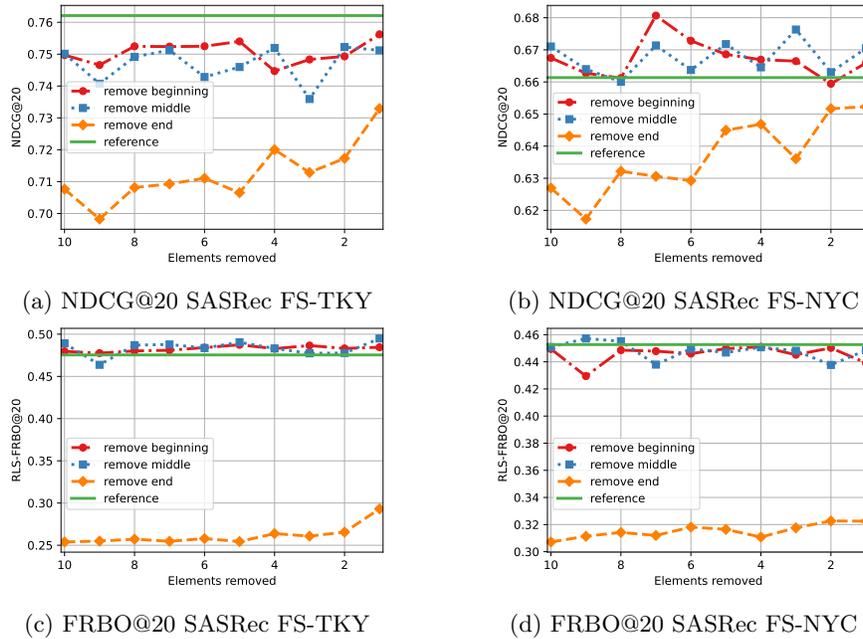
Fig. 3: Plots of NDCG and FRBO for SASRec on FS-TKY and FS-NYC datasets. The baseline is shown as a horizontal solid line, while dashed lines show the metrics as the number of items removed changes for the three scenarios considered.

6   Conclusion

In this work, we have analyzed the importance of the position of items in a temporally ordered sequence for training SRSs. For this purpose, we introduced Finite RBO, a version of RBO for finite-length ranking lists and proved its normalization in [0,1]. Our results demonstrate the importance of the most recent elements in users' interaction sequence: when these items are removed from the training data, there is a significant drop in all evaluation metrics for all case studies investigated and this reduction is proportional to the number of elements removed. Conversely, this reduction is not as pronounced when elements at the beginning and middle of the sequence are removed. We validated our hypothesis using four different datasets and two different models, using traditional evaluation metrics such as NDCG, Recall, but also RLS, a measure specifically designed to measure Sensitivity. Future work in this direction could first extend our results to more models and more datasets, and then investigate a way to make the models robust to the removal of training data. We hypothesize that the solution may lie in using different training strategies [27], robust loss functions [5,40], or different optimization objectives [4].

References

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

2. Afchar, D., Melchiorre, A., Schedl, M., Hennequin, R., Epure, E., Moussallam, M.: Explainability in music recommender systems. AI Magazine **43**(2), 190–208 (2022)

3. Amato, F., Moscato, V., Picariello, A., Sperlí, G.: Recommendation in social media networks. In: 2017 IEEE Third International Conference on Multimedia Big Data (BigMM). pp. 213–216. IEEE (2017)

4. Bacciu, A., Siciliano, F., Tonellotto, N., Silvestri, F.: Integrating item relevance in training loss for sequential recommender systems. arXiv preprint arXiv:2305.10824 (2023)

5. Bucarelli, M.S., Cassano, L., Siciliano, F., Mantrach, A., Silvestri, F.: Leveraging inter-rater agreement for classification in the presence of noisy labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3439–3448 (2023)

6. Burke, R., O'Mahony, M.P., Hurley, N.J.: Robust collaborative recommendation. Recommender systems handbook pp. 961–995 (2015)

7. Chang, J., Gao, C., Zheng, Y., Hui, Y., Niu, Y., Song, Y., Jin, D., Li, Y.: Sequential recommendation with graph neural networks. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. p. 378–387. SIGIR '21, Association for Computing Machinery, New York, NY, USA (2021). `https://doi.org/10.1145/3404835.3462968`, `https://doi.org/10.1145/3404835.3462968`

8. Cho, K., van Merrienboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder-decoder for statistical machine translation (2014)

9. Donkers, T., Loepp, B., Ziegler, J.: Sequential user-based recurrent neural network recommendations. In: Proceedings of the Eleventh ACM Conference on Recommender Systems. p. 152–160. RecSys '17, Association for Computing Machinery, New York, NY, USA (2017). `https://doi.org/10.1145/3109859.3109877`, `https://doi.org/10.1145/3109859.3109877`

10. Fagin, R., Kumar, R., Sivakumar, D.: Comparing top k lists. SIAM Journal on discrete mathematics **17**(1), 134–160 (2003)

11. Fan, Z., Liu, Z., Zhang, J., Xiong, Y., Zheng, L., Yu, P.S.: Continuous-time sequential recommendation with temporal graph collaborative transformer. In: Proceedings of the 30th ACM International Conference on Information & Knowledge Management. p. 433–442. CIKM '21, Association for Computing Machinery, New York, NY, USA (2021). `https://doi.org/10.1145/3459637.3482242`, `https://doi.org/10.1145/3459637.3482242`

12. Fouss, F., Faulkner, S., Kolp, M., Pirotte, A., Saerens, M., et al.: Web recommendation system based on a markov-chainmodel. In: ICEIS (4). pp. 56–63 (2005)

13. Fouss, F., Pirotte, A., Saerens, M.: A novel way of computing similarities between nodes of a graph, with application to collaborative recommendation. In: The 2005 IEEE/WIC/ACM International Conference on Web Intelligence (WI'05). pp. 550–556. IEEE (2005)

14. Guy, I., Zwerdling, N., Ronen, I., Carmel, D., Uziel, E.: Social media recommendation based on people and tags. In: Proceedings of the 33rd International

ACM SIGIR Conference on Research and Development in Information Retrieval. p. 194–201. SIGIR '10, Association for Computing Machinery, New York, NY, USA (2010). https://doi.org/10.1145/1835449.1835484, https://doi.org/10.1145/1835449.1835484

15. Harper, F.M., Konstan, J.A.: The movielens datasets: History and context. ACM Trans. Interact. Intell. Syst. **5**(4) (dec 2015). https://doi.org/10.1145/2827872, https://doi.org/10.1145/2827872

16. Hidasi, B., Karatzoglou, A., Baltrunas, L., Tikk, D.: Session-based recommendations with recurrent neural networks (2016)

17. Hurley, N.J.: Robustness of recommender systems. In: Proceedings of the Fifth ACM Conference on Recommender Systems. p. 9–10. RecSys '11, Association for Computing Machinery, New York, NY, USA (2011). https://doi.org/10.1145/2043932.2043937, https://doi.org/10.1145/2043932.2043937

18. Hwangbo, H., Kim, Y.S., Cha, K.J.: Recommendation system development for fashion retail e-commerce. Electronic Commerce Research and Applications **28**, 94–101 (2018)

19. Jaccard, P.: The distribution of the flora in the alpine zone. 1. New Phytologist **11**(2), 37–50 (1912)

20. Kang, W.C., McAuley, J.: Self-attentive sequential recommendation. In: 2018 IEEE International Conference on Data Mining (ICDM). pp. 197–206. IEEE (2018)

21. Kendall, M.G.: Rank correlation methods. (1948)

22. Li, Y., Chen, H., Fu, Z., Ge, Y., Zhang, Y.: User-oriented fairness in recommendation. In: Proceedings of the Web Conference 2021. p. 624–632. WWW '21, Association for Computing Machinery, New York, NY, USA (2021). https://doi.org/10.1145/3442381.3449866, https://doi.org/10.1145/3442381.3449866

23. Ni, J., Li, J., McAuley, J.: Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In: Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP). pp. 188–197 (2019)

24. Oh, S., Ustun, B., McAuley, J., Kumar, S.: Rank list sensitivity of recommender systems to interaction perturbations. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management. p. 1584–1594. CIKM '22, Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3511808.3557425, https://doi.org/10.1145/3511808.3557425

25. O'Mahony, M., Hurley, N., Kushmerick, N., Silvestre, G.: Collaborative recommendation: A robustness analysis. ACM Transactions on Internet Technology (TOIT) **4**(4), 344–377 (2004)

26. Ovaisi, Z., Heinecke, S., Li, J., Zhang, Y., Zheleva, E., Xiong, C.: Rgrecsys: A toolkit for robustness evaluation of recommender systems. In: Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining. pp. 1597–1600 (2022)

27. Petrov, A., Macdonald, C.: Effective and efficient training for sequential recommendation using recency sampling. In: Proceedings of the 16th ACM Conference on Recommender Systems. pp. 81–91 (2022)

28. Purificato, A., Cassarà, G., Liò, P., Silvestri, F.: Sheaf neural networks for graph-based recommender systems (2023)

29. Quadrana, M., Cremonesi, P., Jannach, D.: Sequence-aware recommender systems. ACM Computing Surveys (CSUR) **51**(4), 1–36 (2018)

30. Quadrana, M., Karatzoglou, A., Hidasi, B., Cremonesi, P.: Personalizing session-based recommendations with hierarchical recurrent neural networks. In: Proceedings of the Eleventh ACM Conference on Recommender Systems. p. 130–137. RecSys '17, Association for Computing Machinery, New York, NY, USA (2017). `https://doi.org/10.1145/3109859.3109896`, `https://doi.org/10.1145/3109859.3109896`

31. Schafer, J.B., Konstan, J.A., Riedl, J.: E-commerce recommendation applications. Data Mining and Knowledge Discovery **5**, 115–153 (2001)

32. Schedl, M., Knees, P., McFee, B., Bogdanov, D., Kaminskas, M.: Music recommender systems. Recommender Systems Handbook pp. 453–492 (2015)

33. Schedl, M., Zamani, H., Chen, C.W., Deldjoo, Y., Elahi, M.: Current challenges and visions in music recommender systems research. International Journal of Multimedia Information Retrieval **7**, 95–116 (2018)

34. Shapiro, S.S., Wilk, M.B.: An analysis of variance test for normality (complete samples). Biometrika **52**(3/4), 591–611 (1965)

35. Student: The probable error of a mean. Biometrika **6**(1), 1–25 (1908)

36. Sun, F., Liu, J., Wu, J., Pei, C., Lin, X., Ou, W., Jiang, P.: Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer. In: Proceedings of the 28th ACM International Conference on Information and Knowledge Management. p. 1441–1450. CIKM '19, Association for Computing Machinery, New York, NY, USA (2019). `https://doi.org/10.1145/3357384.3357895`, `https://doi.org/10.1145/3357384.3357895`

37. Tang, J., Drori, Y., Chang, D., Sathiamoorthy, M., Gilmer, J., Wei, L., Yi, X., Hong, L., Chi, E.H.: Improving training stability for multitask ranking models in recommender systems. arXiv preprint arXiv:2302.09178 (2023)

38. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. p. 6000–6010. NIPS'17, Curran Associates Inc., Red Hook, NY, USA (2017)

39. Wang, S., Hu, L., Wang, Y., Cao, L., Sheng, Q.Z., Orgun, M.: Sequential recommender systems: challenges, progress and prospects. arXiv preprint arXiv:2001.04830 (2019)

40. Wani, F.A., Bucarelli, M.S., Silvestri, F.: Combining distance to class centroids and outlier discounting for improved learning with noisy labels (2023)

41. Webber, W., Moffat, A., Zobel, J.: A similarity measure for indefinite rankings. ACM Trans. Inf. Syst. **28**(4) (nov 2010). `https://doi.org/10.1145/1852102.1852106`, `https://doi.org/10.1145/1852102.1852106`

42. Wu, S., Crestani, F.: Methods for ranking information retrieval systems without relevance judgments. In: Proceedings of the 2003 ACM symposium on Applied computing. pp. 811–816 (2003)

43. Wu, S., Sun, F., Zhang, W., Xie, X., Cui, B.: Graph neural networks in recommender systems: a survey. ACM Computing Surveys **55**(5), 1–37 (2022)

44. Yang, D., Zhang, D., Zheng, V.W., Yu, Z.: Modeling user activity preference by leveraging user spatial temporal characteristics in lbsns. IEEE Transactions on Systems, Man, and Cybernetics: Systems **45**(1), 129–142 (2014)

45. Zhang, S., Yao, L., Sun, A., Tay, Y.: Deep learning based recommender system: A survey and new perspectives. ACM Comput. Surv. **52**(1) (feb 2019). `https://doi.org/10.1145/3285029`, `https://doi.org/10.1145/3285029`

46. Zhao, W.X., Mu, S., Hou, Y., Lin, Z., Chen, Y., Pan, X., Li, K., Lu, Y., Wang, H., Tian, C., Min, Y., Feng, Z., Fan, X., Chen, X., Wang, P., Ji, W., Li, Y., Wang, X.,

Wen, J.R.: Recbole: Towards a unified, comprehensive and efficient framework for recommendation algorithms (2021)