

UniRec: A Dual Enhancement of Uniformity and Frequency in Sequential Recommendations

Yang Liu
liuyangfd22@m.fudan.edu.cn
School of Computer Science
Fudan University
Shanghai, China

Yitong Wang*
yitongw@fudan.edu.cn
School of Computer Science
Fudan University
Shanghai, China

Chenyue Feng
cyfeng22@m.fudan.edu.cn
School of Computer Science
Fudan University
Shanghai, China

ABSTRACT

Representation learning in sequential recommendation is critical for accurately modeling user interaction patterns and improving recommendation precision. However, existing approaches predominantly emphasize item-to-item transitions, often neglecting the time intervals between interactions, which are closely related to behavior pattern changes. Additionally, broader interaction attributes, such as item frequency, are frequently overlooked. We found that both sequences with more uniform time intervals and items with higher frequency yield better prediction performance. Conversely, non-uniform sequences exacerbate user interest drift and less-frequent items are difficult to model due to sparse sampling, presenting unique challenges inadequately addressed by current methods. In this paper, we propose UniRec, a novel bidirectional enhancement sequential recommendation method. UniRec leverages sequence uniformity and item frequency to enhance performance, particularly improving the representation of non-uniform sequences and less-frequent items. These two branches mutually reinforce each other, driving comprehensive performance optimization in complex sequential recommendation scenarios. Additionally, we present a multidimensional time module to further enhance adaptability. To the best of our knowledge, UniRec is the first method to utilize the characteristics of uniformity and frequency for feature augmentation. Comparing with eleven advanced models across four datasets, we demonstrate that UniRec outperforms SOTA models significantly. The code is available at <https://github.com/Linxi000/UniRec>.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Sequential Recommendation, Sequence Uniformity, Item Frequency, Feature Enhancement

ACM Reference Format:

Yang Liu, Yitong Wang, and Chenyue Feng. 2024. UniRec: A Dual Enhancement of Uniformity and Frequency in Sequential Recommendations. In *Proceedings of 33rd ACM International Conference on Information and*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM'24, October 21–25, 2024, Boise, ID

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/18/06
<https://doi.org/XXXXX.XXXXX>

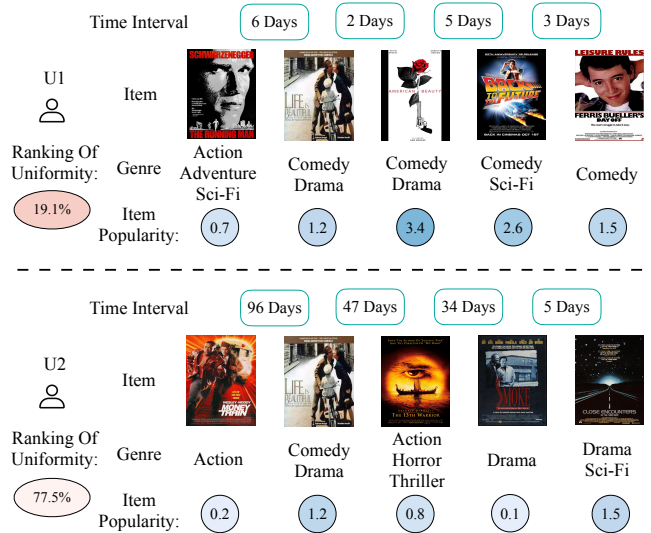


Figure 1: An example of uniform and non-uniform sequences in a real dataset.

Knowledge Management (CIKM'24). ACM, New York, NY, USA, 10 pages.
<https://doi.org/XXXXX.XXXXX>

1 INTRODUCTION

Sequential recommendation systems have become increasingly prevalent due to their ability to effectively model user preferences [40, 31, 10, 39]. Such systems utilize the sequential order of user interactions over time to predict future interests [21, 36, 7]. Incorporating temporal information into these algorithms has proven effective, as it provides significant insights into user behavioral patterns [25, 45, 4, 9, 38, 6]. **Current approaches primarily focus on modeling explicit timestamps [25, 32] or capturing cyclic patterns [4], but they often overlook time intervals, which reveal user characteristics and convey critical information within user interaction sequences.** Yizhou Dang et al. propose that variations in the time intervals between sequential interactions can serve as indicators of shifts in user preferences [5]. Building on this premise, they designed data augmentation operators to improve the uniformity of sequences. However, this direction still lacks full study and holds potential significance, as sequence uniformity is a common phenomenon across various datasets. **Additionally, the effectiveness of a model in capturing item characteristics is influenced by the frequency of these items.** While considerable research has focused on

enhancing the recommendation performance for long-tail items [22, 27], the utilization of item frequency to enhance model performance remains an area requiring further exploration.

Figure 1 illustrates segments of the interaction of uniform sequence versus non-uniform sequences from different users, encompassing items of both high and low frequency. The "Ranking of Uniformity" sorts interaction sequences by the variance of their time intervals in ascending order, with lower percentages indicating greater uniformity. For example, U1 with a ranking of 19.1% is more uniform than 80.9% of the sequences. "Item Popularity" is defined as the proportion of an item's occurrences relative to the number of all interactions, thus quantifying the frequency of item appearances within the dataset. This figure illustrates that time intervals within uniform sequences are typically shorter and more stable, indicating steadier user interests. In contrast, non-uniform sequences exhibit more variable time intervals, reflecting more frequent changes in user interests. Furthermore, the intensity of the color within the circles signifies the model's effectiveness in learning the representations of the corresponding users or items, with darker colors indicating higher effectiveness.

We first analyze the performance of sequences with different intervals and item frequencies in section 2 and validate that sequences with higher uniformity and items with greater frequency tend to exhibit better performance. Following this, we implement a dual enhancement approach UniRec in section 3. For sequences, we generate non-uniform subsets from uniform sequences by incorporating less-frequent items to simulate fluctuating user interests, thereby enhancing the modeling of non-uniform sequence representations later. For items, we train a neighbor aggregation mechanism on frequent items and extend it to less-frequent items using curriculum learning to improve their representations and transfer this knowledge to sequence modeling. This dual-branch approach is simple and effective, providing a new perspective for feature enhancement in sequential recommendation. Additionally, we integrate the temporal characteristics of both uniform and non-uniform sequences to conduct multidimensional temporal modeling.

In summary, the contributions of this paper are as follows:

- We propose a novel dual enhancement architecture that leverages sequence uniformity and item frequency. This architecture comprises two independent yet mutually reinforced branches, collectively driving comprehensive performance optimization.
- We improve the model's ability to handle non-uniform sequences and less-frequent items and provide a new perspective for feature enhancement in sequential recommendation.
- We conduct extensive experiments on 4 real-world datasets, demonstrating significant improvements over 11 competing models, including 6 cutting-edge models that incorporate temporal modeling in their sequential recommendation systems.

2 PRELIMINARY STUDY

In subsection 2.2, we demonstrate that uniform sequences and frequent items consistently perform better across various datasets. In subsection 2.3, we further validate this by demonstrating that,

Table 1: Performance of sequential recommendation models on different subsets.

Dataset	Strategy	SASRec			Bert4Rec			LightSANS		
		NDCG	Hit	MRR	NDCG	Hit	MRR	NDCG	Hit	MRR
ML-1M	all	0.1584	0.3449	0.1058	0.1779	0.3770	0.1218	0.1779	0.3770	0.1218
	\mathbb{I}_f	0.1714	0.3707	0.1151	0.1923	0.4025	0.1331	0.1923	0.4025	0.1331
	\mathbb{I}_l	0.0846	0.1980	0.0530	0.0958	0.2323	0.0573	0.0958	0.2323	0.0573
	Impr.	102.6%	87.22%	117.17%	100.73%	73.27%	132.29%	100.73%	73.27%	132.29%
	\mathbb{S}_u	0.1958	0.4145	0.1340	0.2171	0.4501	0.1511	0.2001	0.4222	0.1374
	\mathbb{S}_n	0.1024	0.2405	0.0636	0.1191	0.2674	0.0778	0.1030	0.2363	0.0658
	Impr.	91.21%	72.35%	110.69%	82.28%	68.32%	94.22%	94.27%	78.67%	108.81%
Gowalla	all	0.1214	0.1950	0.0999	0.0982	0.1639	0.0791	0.1310	0.2090	0.1082
	\mathbb{I}_f	0.1502	0.241	0.1235	0.1207	0.1998	0.0977	0.1522	0.2406	0.1261
	\mathbb{I}_l	0.0821	0.1376	0.0661	0.0490	0.0856	0.0384	0.0869	0.1436	0.0703
	Impr.	82.95%	75.15%	86.84%	146.33%	133.41%	154.43%	75.14%	67.55%	79.37%
	\mathbb{S}_u	0.1466	0.2341	0.1208	0.1215	0.2013	0.0983	0.1568	0.2503	0.1291
	\mathbb{S}_n	0.1026	0.1657	0.0842	0.0807	0.1360	0.0647	0.1118	0.1782	0.0925
	Impr.	42.88%	41.28%	43.47%	50.56%	48.01%	51.93%	40.25%	40.46%	39.57%

regardless of the partitioning thresholds, uniformity and frequency consistently lead to better performance.

2.1 Symbol Description

We distinguish the uniformity and non-uniformity of sequences by adopting the classification method proposed by TiCoSeRec [5], which evaluates and ranks all sequences by calculating the variance of time intervals. Sequences with smaller variances are considered more uniform. Based on this, sequences are divided into two subsets: \mathbb{S}_u and \mathbb{S}_n . The former includes sequences with consistent time intervals, while the latter contains sequences with significant fluctuations in intervals. Similarly, we rank each item based on the frequency of its occurrence across all user interactions. Define \mathbb{I}_f as the set of frequently occurring items and \mathbb{I}_l as the set of less-frequently occurring items.

2.2 Generality Analysis

2.2.1 Task. In this experiment, we aim to investigate the comparative recommendation performance on uniform versus non-uniform sequences as well as frequent versus less-frequent items, within the context of different datasets. To achieve balance and fairness, we ensured that subsets \mathbb{S}_u and \mathbb{S}_n , as well as \mathbb{I}_f and \mathbb{I}_l , were balanced by equating the interaction numbers as much as possible. Following this division criterion, we assigned "uniformity" and "frequency" labels to each interaction sequence and item, recording the overall evaluation results of the model and the experimental outcomes for data with different labels.

2.2.2 Experimental Configuration. TiCoSeRec [5] has already demonstrated on several Amazon datasets and Yelp that uniform sequences significantly outperform non-uniform sequences. Here, we extend these findings to both frequent and less-frequent items by testing on two additional datasets, MovieLens 1M (ML-1M) [12] and Gowalla [3]. The ML-1M dataset, a publicly available movie ratings database, comprises 999,611 ratings from 6,040 users on 3,416 movies, with a sparsity of 95.16%. The Gowalla dataset, representing check-in data from a location-based social network, contains 6,442,892 check-ins at 1,280,970 unique locations by 107,093 users, with a sparsity of 99.99%. We utilized three classical sequential recommendation

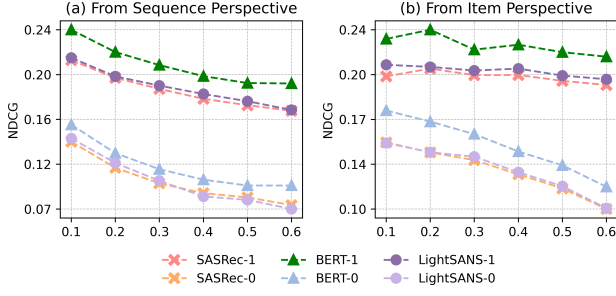


Figure 2: The performance of models under different subset partition ratios, with the X-axis representing the percentage of data classified as uniform and frequent.

baselines—SASRec [21], BERT4Rec [36], and LightSANS [7] for our analysis. The evaluation metrics include Normalized Discounted Cumulative Gain (NDCG), Hit Rate (HR), and Mean Reciprocal Rank (MRR) at top 20. The evaluation strategy employed is full ranking, which involves evaluating the model on the entire set of items.

2.2.3 Results Analysis. Table 1 shows the performance of various baselines across two datasets, comparing uniform and non-uniform sequences, as well as frequent and less-frequent items. In the table, "all" represents results tested on the entire dataset, while \mathbb{S}_u and \mathbb{S}_n , along with \mathbb{I}_f and \mathbb{I}_l , represent results tested on these specific subsets. The experimental results show that performance on subsets \mathbb{S}_u and \mathbb{I}_f is the best, also "all" exceed those on \mathbb{S}_n and \mathbb{I}_l . For the Gowalla dataset, the Bert4Rec model shows up to a 146.33% improvement in NDCG@20 when predicting \mathbb{I}_f instead of \mathbb{I}_l . Similarly, LightSANS improves by up to 94.27% in NDCG@20 for the ML-1M dataset when transitioning from \mathbb{S}_n to \mathbb{S}_u . This phenomenon, where performance on \mathbb{I}_f substantially exceeds that on \mathbb{I}_l , corroborates the hypothesis that frequent items, benefiting from a larger volume of interaction data, are more predictable. Additionally, models generally exhibit superior performance on \mathbb{S}_u compared to \mathbb{S}_n , suggesting that models more effectively learn from stable user preferences present in uniform sequences.

2.3 Invariance Analysis

We further explore the impact of different partitioning ratios on model performance using the ML-1M dataset. Specifically, we analyze the effects of varying the ratios for both \mathbb{S}_u and \mathbb{S}_n and \mathbb{I}_f and \mathbb{I}_l using three classical baseline models.

Figure 2a displays the experimental results on \mathbb{S}_u and \mathbb{S}_n . In this figure, the "-1" suffix attached to each model indicates the performance on the \mathbb{S}_u , whereas the "-0" suffix indicates the performance on the \mathbb{S}_n . Figure 2b presents the results on \mathbb{I}_f and \mathbb{I}_l , where "-1" and "-0" similarly denote the performance on \mathbb{I}_f and \mathbb{I}_l , respectively. The performance trends on MRR@20 and HR@20 are very similar to those observed with NDCG@20.

The results indicate a noticeable decline in the performance of sequential recommendation models as the partitioning thresholds shift from uniform to non-uniform sequences and from frequent to less-frequent items. This trend highlights the models' sensitivity to the variability in user behavior patterns and item frequencies.

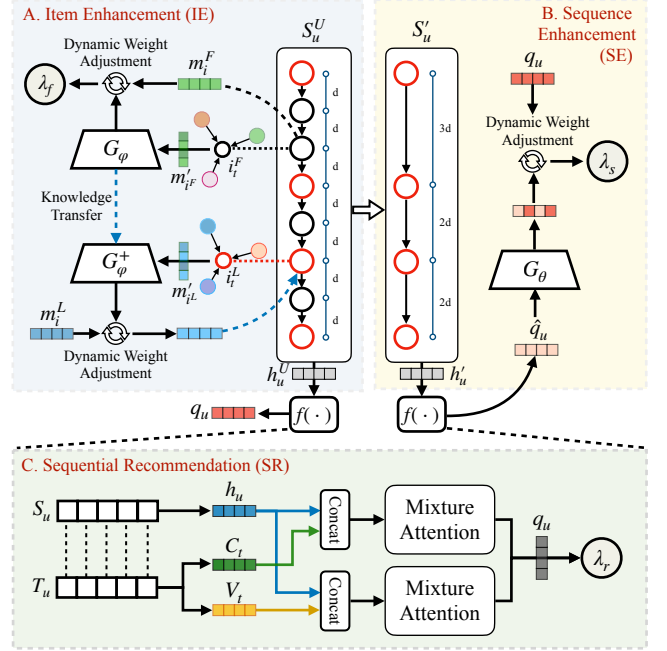


Figure 3: Overview framework of Item Enhancement (A), Sequence Enhancement (B), and Multidimensional Time Modeling in Sequential Recommendation (C), using a uniform sequence as an example.

3 METHODOLOGY

This section provides a detailed exposition of UniRec. First, we address the dual enhancement architecture, which comprises the sequences branch (subsection 3.2) and the items branch (subsection 3.3). Subsequently, a Multidimensional Time mixture attention module (subsection 3.4) is designed to accommodate different uniformity sequences. Lastly, subsection 3.5 describes the inference process of the model. Figure 3 illustrates the overall architecture of the UniRec framework.

3.1 Problem Formulation

Let \mathcal{U} denote the set of all users and \mathcal{I} represent the set of all items. For each user $u \in \mathcal{U}$, we formulate the interactions in chronological order, expressed as $\mathcal{S}_u^{s\text{-type}} = (i_1^{i\text{-type}}, \dots, i_t^{i\text{-type}}, \dots, i_N^{i\text{-type}})$. Here, $i_t^{i\text{-type}} \in \mathcal{I}$ specifies the item with which the user interacted at timestamp t . The term "s-type" distinguishes a sequence as uniform or non-uniform, denoted as \mathcal{S}_u^U and \mathcal{S}_u^N ; "i-type" identifies an item as frequent or less-frequent as i_t^F and i_t^L , respectively. N signifies the sequence length, which is fixed. For sequences shorter than N , we employ the padding operation to fill the missing parts and for those longer than N we truncate the excess part. Define $M_I \in \mathbb{R}^{\mathcal{I} \times d}$ as a learnable matrix of all items' embedding, d is a positive integer denoting the latent dimension. By performing a lookup table operation on M_I , we can retrieve every single item embedding $m_i \in \mathbb{R}^d$, to form the user embedding $h_u = [m_1, \dots, m_t, \dots, m_N] \in \mathbb{R}^{N \times d}$.

3.2 Sequence Enhancement

Sequences with smaller variances are considered more uniform and sequences are divided into two subsets: \mathbb{S}_u and \mathbb{S}_n . Each sequence is classified based on a predefined time variance threshold into either \mathcal{S}_u^U or \mathcal{S}_u^N , where $\mathcal{S}_u^U \in \mathbb{S}_u$ and $\mathcal{S}_u^N \in \mathbb{S}_n$. Similarly, item is categorized based on their frequency of occurrence in interactions into i_t^F or i_t^L , where $i_t^F \in \mathbb{I}_f$ and $i_t^L \in \mathbb{I}_l$. For each uniform sequence \mathcal{S}_u^U , we generate a corresponding non-uniform sub-sequence \mathcal{S}'_u to emulate the irregular patterns observed in real-world datasets, thereby enhancing the capability to model complex user behaviors. The generation process retains all items from \mathbb{I}_l within \mathcal{S}_u^U , and if the count of $i_t^L \in \mathcal{S}_u^U$ is fewer than M , additional i_t^L are randomly sampled from \mathcal{S}_u^U , where M is the hyper-parameter of the minimum length of \mathcal{S}'_u :

$$\mathcal{S}'_u = \begin{cases} \text{if } \text{count}(\mathcal{S}_u^U, \mathbb{I}_l) < M : \\ \quad \{i_t^L : i_t^L \in \mathcal{S}_u^U\} \cup \{\text{Sampled } i_t^L \in \mathcal{S}_u^U\} \\ \text{otherwise :} \\ \quad \{i_t^L : i_t^L \in \mathcal{S}_u^U\} \end{cases} \quad (1)$$

the variance of time intervals increases from the sequence \mathcal{S}_u^U to \mathcal{S}'_u , and there is a substantial rise in the relative composition of i_t^L within \mathcal{S}'_u .

We utilize \mathcal{S}'_u to enhance the model's learning capability with respect to \mathcal{S}'_u . First, we generate the initial embeddings for \mathcal{S}_u^U and \mathcal{S}'_u , denoted as $h_u^U \in \mathbb{R}^{N \times d}$ and $h'_u \in \mathbb{R}^{N \times d}$ respectively. For each sequence, we employ a sequence encoder $f(\cdot)$, which is the sequential recommendation modeling process:

$$q_u = f(h_u^U), \quad \hat{q}_u = f(h'_u) \quad (2)$$

where $q_u \in \mathbb{R}^{N \times 2d}$ and $\hat{q}_u \in \mathbb{R}^{N \times 2d}$ are the representations for \mathcal{S}_u^U and \mathcal{S}'_u . The specifics of $f(\cdot)$ will be detailed in subsection 3.4. Next, the objective is to bring q_u and \hat{q}_u as close as possible in the feature space to enhance the model's ability to handle the temporal dynamics of non-uniform sequences, thereby minimizing \tilde{x} through a generative model G_θ , which consists of a feed-forward layer:

$$\tilde{x} = q_u - G_\theta(\hat{q}_u) \quad (3)$$

Meanwhile, a **curriculum learning strategy** is adopted, which mimics the human learning process: from simple to complex. This strategy gradually increases the training samples' complexity. Specifically, the model initially learns predominantly from more uniform sequences, while sequences with more complex user interest drifts are introduced later in the training. This process is managed with a dynamically weighted loss function λ_s guiding the progression:

$$\lambda_s = w_s \|\tilde{x}\|^2 \quad (4)$$

$$w_s = \sin\left(\frac{\pi}{2} \cdot \frac{e - e_b}{e_{all}} + \frac{\pi}{2} \cdot \frac{V_{max} - V_u}{V_{max} - V_{min}}\right) \quad (5)$$

where w_s represents a dynamic weight coefficient, e denotes the current epoch number, e_b denotes the epoch at which this loss function starts to contribute to the training process, and e_{all} denotes the total number of training epochs. For each $\mathcal{S}_u^U \in \mathbb{S}_u$, the variance of the time intervals is defined as V_u . V_{max} is the maximum time interval variance among all sequences, while V_{min} is the minimum. This design allows w_s to dynamically change its value during the training process based on the uniformity of sequences and phases of

training progress. This task serving as an auxiliary task, parallel to the main task of sequential recommendation, specifically enhances the model's performance on \mathcal{S}'_u , thereby implicitly improving the model's adaptability and prediction accuracy on \mathbb{S}_n .

3.3 Item Enhancement

Given that the generated \mathcal{S}'_u are predominantly composed of i_t^L , together with a general prevalence of i_t^L in \mathcal{S}_u^N , enhancing model performance on i_t^L has become critical. The proposed item enhancement approach operates from two aspects: utilizing the information from neighboring items and leveraging the knowledge transferred from i_t^F to i_t^L . Leveraging neighbors for enhancement involves two steps: candidate neighbor generation and representation aggregation.

Initially, the candidate neighbor generation process is conducted for each item. For each center item $i_c \in \mathcal{I}$, a potential candidate neighbor set \mathbb{N}_{i_c} is identified. A bunch of score $s(i_c, j)$ is calculated for i_c against every other item j (where $j \in \mathcal{I} \setminus \{i_c\}$). These scores are then ranked, and the items with higher scores are chosen to constitute the neighbor set \mathbb{N}_{i_c} . $s(i_c, j)$ integrated three factors: the temporal interval T between i_c and j , the popularity H of item j , and the similarity S between i_c and j . Both H and S are normalized to ensure consistency in the scoring mechanism. $s(i_c, j)$ is defined as:

$$s(i_c, j) = g(T) + \phi(T, H) + \phi(T, S) \quad (6)$$

$$g(T) = \frac{1}{1 + \log(1 + T)} \quad (7)$$

$$\phi(T, x) = \frac{T + \Theta}{e^{(T+\Theta)/\Gamma x}} \quad (8)$$

where Θ and Γ are constants, determined based on dataset specifics. As T increases, $g(T)$ gradually decreases. Similarly, an increase in T or a decrease in x results in a lower value of $\phi(T, x)$. This scoring framework adeptly manages the temporal dynamics among items, accounting for factors such as the popularity and similarity of potential neighboring items. In each training batch, K neighbors are randomly sampled from \mathbb{N}_{i_c} , where K is a hyper-parameter.

Then we aggregate these K candidate neighbors to enhance i_c with a simple attention mechanism. We generate the initial embedding for i_c , denoting as $m_c \in \mathbb{R}^d$, as well as the embedding $m_o \in \mathbb{R}^d$ for these K neighbors, $o \in \{1, 2, \dots, K\}$. The aggregation process is as follows:

$$m_n = \sum_{k=1}^K \frac{\exp(m_c^T m_k)}{\sum_{j=1}^K \exp(m_c^T m_j)} \quad (9)$$

m_n represents the aggregated embedding from the neighbors. We then concatenate m_n and m_c to form the updated representation $m'_c = [m_c \parallel m_n] \in \mathbb{R}^{2d}$, where \parallel denotes the concatenation operation. As a result, m'_c contains more information related to i_c than m_c .

Meanwhile, to enable $i_t^L \in \mathbb{I}_l$ to better utilize the related information from \mathbb{N}_{i_c} , we transfer the knowledge learned from \mathbb{I}_f on neighbor aggregation representation to \mathbb{I}_l . Define the embedding of i_t^F obtained from M_I as m_i^F . Define the updated embedding m'_c of i_t^F as m'_i . We train the aggregation mechanism on i_t^F by minimizing the following loss function:

$$\lambda_f = w_i \|m_i^F - G_\varphi(m_{i:F}')\|^2 \quad (10)$$

$$w_i = \sin\left(\frac{\pi}{2} \cdot \frac{e - e_b}{e_{all}} + \frac{\pi}{2} \cdot \frac{F - F_{min}}{F_{max} - F_{min}}\right) \quad (11)$$

where G_φ is a fully connected layer that aligns the dimensions of $m_{i:F}'$ and m_i^F to be consistent. w_i is a dynamic parameter used to adjust the magnitude of the loss function across different items. F represents the frequency score of the current item across all interactions. F_{min} is the minimum F of $i_t^F \in \mathbb{I}_t$, while F_{max} is the maximum. A curriculum learning strategy, analogous to the sequence branch, is also employed. In the initial training phase, high-frequency items are prioritized, with a gradual shift towards less-frequent items in the later stages.

Finally, update the embeddings of all i_t^L after a certain epoch e_t of training by minimizing the following loss:

$$\lambda_l = \eta \|m_i^L - G_\varphi^+(m_{i:L}')\|^2 \quad (12)$$

$$\eta = \sin\left(\frac{\pi}{2} \cdot \frac{e - e_t}{e_{all}}\right) \quad (13)$$

where m_i^L is the representation of i_t^L obtained from M_I , $m_{i:L}'$ is the updated representation m_c' of i_t^L , and η is a parameter that dynamically increases with the increase of the training epoch. G_φ^+ represents the G_φ trained after $(e - e_b)$ epochs and is static. By refining i_t^L representation through the auxiliary task before the main task training, the accuracy and performance of the model concerning i_t^L are improved.

3.4 Multidimensional Time Modeling

Given the varying dependencies on temporal information, where S_u^U has a lower reliance on time and S_u^N requires richer temporal details, we propose a multidimensional time modeling module to accommodate these differing needs. As demonstrated in subsection 4.6, utilizing time interval information is more effective for S_u^U , while employing comprehensive temporal context proves more effective for S_u^N . Therefore, we design this module to better leverage the appropriate temporal information.

For each S_u we define its corresponding timestamp sequence as $\mathcal{T}_u = (t_1, t_2, \dots, t_N)$. The corresponding time interval sequence is defined as $\mathcal{T}_{intv} = (\tau_1, \tau_2, \dots, \tau_{N-1})$, where each $\tau_k = t_{k+1} - t_k$ denotes the interval between the k^{th} and $(k+1)^{\text{th}}$ interactions. Each τ_k is encoded by an embedding matrix, resulting in a time interval embedding $v_k \in \mathbb{R}^d$. For temporal context modeling, we adopted the approach proposed by Xu et al. [43], which specifically uses a self-attention mechanism based on time representation learning, and models temporal information such as year, month, and day separately. Subsequently, this information is aggregated through a linear layer to form the final temporal context embedding $c_i \in \mathbb{R}^d$ for each interaction i . In a word, for each S_u , we obtain its item sequence embedding $h_u \in \mathbb{R}^{N \times d}$, along with the temporal context representation $C_t = [c_1, c_2, \dots, c_N] \in \mathbb{R}^{N \times d}$, and the time interval embeddings $V_t = [0, v_1, v_2, \dots, v_{n-1}] \in \mathbb{R}^{N \times d}$, 0 represents a $1 \times d$ zero vector.

Next, recognizing that sequences with different uniformity require varying levels of temporal information, we integrate h_u with C_t and V_t respectively using a mixture attention mechanism. This

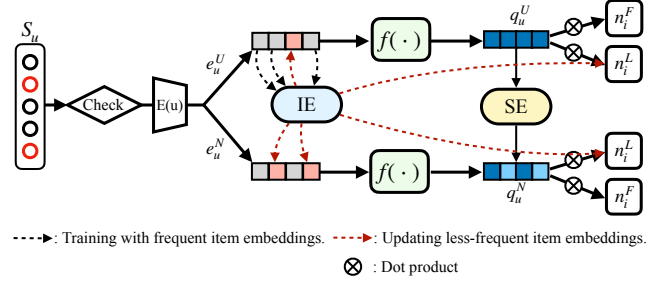


Figure 4: Overview of inference phase.

serves as the sequence encoder $f(\cdot)$, generating q_u , the embedding of the user u 's interaction sequence, tailored to the specific needs of each sequence. Integrate h_u with C_t and V_t in the same way, taking the application of mixture attention on h_u and C_t as an example. First, concatenate h_u and C_t to obtain the initial embedding of a sequence as $e_u = h_u \parallel C_t$. Next, we preprocess the input X for mixture attention, which is defined as $X = e_u + P$, where $P \in \mathbb{R}^{N \times 2d}$ is the position encoding matrix. The mixture attention mechanism can be mathematically described as:

$$\text{MixATT}(X) = \text{FFL}(\text{SAL}(X)) \quad (14)$$

$$\text{FFL}(X) = \text{ReLU}(XW_F + b_F)W_{F'} + b_{F'} \quad (15)$$

$$\text{SAL}(X) = \text{Concat}(H_1, \dots, H_H) \quad (16)$$

where $\text{MixATT}(X)$ represents a composite model that integrates a self-attention mechanism $\text{SAL}(X)$ and a feed-forward layer $\text{FFL}(X)$. FFL involves two linear transformations with weight matrices W_F and $W_{F'}$, and bias terms b_F and $b_{F'}$. SAL combines the outputs H_j from each attention head $j \in \{1, \dots, H\}$. Each H_j is given by $\text{softmax}(A_j / \sqrt{d_V})W_j^O$, where $\sqrt{d_V}$ is a scaling factor to stabilize learning, and W_j^O is the output projection matrix for the j^{th} head. A_j is the attention score matrix proposed by Viet-Anh Tran et al. [38], combining Gaussian distribution to mix two types of input data. $A_j = \sum_{k \in \{m, c\}} p_{kj} \mathcal{N}(A; Q_k^T, \sigma^2 I)$ is approximated by a mixture model. The non-negative mixture weights p_{kj} sum to one, indicating the contribution of each context type. Q_k is obtained by projecting the input context X_k using matrix W_k . The Gaussian distribution's variance parameter is σ^2 , and I is the identity matrix.

The loss function for the recommendation task can be defined as follows:

$$\lambda_r = q_u n_i^T \quad (17)$$

where q_u is the output of the FFL and $n_i = [m_i | c_i]$ is the embedding of the next item to be predicted. Similarly, the mixture attention mechanism is also applied to h_u and V_t . The outputs processed through the mixture attention mechanism, are mutually supervised within a multi-task learning framework.

3.5 Inference Process

Figure 4 shows how the integrated components—IE (Item Enhancement), SE (Sequence Enhancement), and $f(u)$ (Sequential Recommendation)—work together to provide robust and contextually rich recommendations. For a given input sequence S_u , we first determine whether it is S_u^U or S_u^N . S_u^U is initialized with embedding e_u^U ,

Table 2: Performance comparison over four datasets. Numbers in bold indicate the best performance, those underlined denote the second best, and numbers marked with an asterisk represent the third best. Models marked with † are data-augmented methods based on SASRec.

Dataset	Metric	Non-Time-Aware							Time-Aware					
		GRU4Rec	Caser	STAMP	SASRec	BERT4Rec	LightSANS	DuoRec [†]	TiSASRec	Meantime	TiCoSeRec [†]	FEARrec	MOJITO	UniRec
ML-1M	NDCG@10	0.5758	0.5447	0.5302	0.5801	0.5658	0.5671	0.5816	0.5849*	0.5804	0.5732	0.5515	<u>0.5929</u>	0.6261
	HR@10	0.7856	0.7692	0.7444	0.8098*	0.7780	0.7719	0.7971	0.7893	0.8098*	0.7904	0.7594	<u>0.8197</u>	0.8347
	MRR@10	0.5093	0.4751	0.4629	0.5149	0.4987	0.5021	0.5044	<u>0.5205</u>	0.5079	0.4967	0.4856	0.5157*	0.5613
Beauty	NDCG@10	0.3014	0.2805	0.2809	0.2944	0.3140	0.3349*	0.3123	0.2908	0.3201	0.3188	0.3382	<u>0.3392</u>	0.3693
	HR@10	0.4593	0.4394	0.4225	0.4418	0.4629	0.5042*	0.4828	0.4323	0.4655	0.4862	0.4863	<u>0.5087</u>	0.5313
	MRR@10	0.2525	0.2313	0.2372	0.2490	0.2682	<u>0.2986</u>	0.2733	0.2471	0.2753	0.2841	0.2922*	0.2865	0.3201
Books	NDCG@10	0.5811	0.5356	0.4816	0.6072	0.5616	0.6049	0.5942	0.6045	0.6073*	0.5875	0.5745	<u>0.6171</u>	0.6309
	HR@10	0.7951	0.7674	0.7078	0.8216	0.7868	0.8176	0.8130	0.8218*	0.8169	0.8023	0.7994	<u>0.8513</u>	0.8617
	MRR@10	0.5134	0.4626	0.4107	0.5394*	0.4905	0.5378	0.5325	0.5357	0.5135	0.5278	0.5036	<u>0.5428</u>	0.5752
Toys	NDCG@10	0.2779	0.2173	0.2446	0.3118	0.2327	<u>0.3364</u>	0.2794	0.3224	0.3187	0.2754	0.3121	0.3323*	0.3609
	HR@10	0.4432	0.3752	0.3886	0.4626	0.3887	0.4934*	0.4643	0.4766	0.4762	0.4431	0.4655	<u>0.5087</u>	0.5260
	MRR@10	0.2270	0.1688	0.2004	0.2652	0.1848	<u>0.2877</u>	0.2687	0.2747	0.2711	0.2421	0.2647	0.2775*	0.3103

while S_u^N is initialized with e_u^N . Within each S_u^U, i_t^F are utilized to train G_φ through the loss function λ_f in the IE module. Conversely, for both S_u^U and S_u^N, i_t^F are updated based on the output from G_φ^+ using the loss λ_l . After processing the sequence through $f(u)$, we train its embedding via the primary task loss λ_r . The sequence embedding is then refined by the SE module to further enhance the sequence representation using the loss λ_s . Finally, the sequence embedding and the embedding of the item to be predicted are scored by calculating their dot product.

4 EXPERIMENT

4.1 Experimental Settings

4.1.1 Datasets. In addition to the ML-1M [12] dataset used in section 2, we also use datasets from e-commerce platforms, including those for books, beauty products, and toys, as detailed below:

- (1) The Amazon Book [13] dataset consists of 6,275,735 interactions of users rating a book. This dataset includes 79,713 users and 91,465 books, with a density of 0.00086, indicating the sparsity of user-item interactions.
- (2) The Amazon Beauty [30] dataset comprises 198,502 interactions involving 22,363 users and 12,101 beauty products, with a density of 0.00073.
- (3) The Amazon Toys [30] dataset includes 167,597 interactions from 19,412 users and 11,924 toys, with a sparse density of 0.00072.

For each dataset, we adopt the k-core filtering [34] as a pre-processing step, which iteratively removes users and items whose interactions are fewer than k , until each user and item in the dataset has at least k interactions. Specifically, for the ML-1M, we set $k_{\text{item}} = 5$ and $k_{\text{user}} = 10$; for the Beauty and Toy, we set $k_{\text{item}} = 5$ and $k_{\text{user}} = 5$; and for the Books, the settings are $k_{\text{user}} = 30$ and $k_{\text{item}} = 20$.

4.2 Evaluation Settings

We arrange the dataset in chronological order and allocate the last item as the validation set and the penultimate item as the test set,

using the remaining data to construct the training set. To ensure fair evaluation, for each positive item in the test set, we pair it with 100 negative items sampled uniformly, and the model’s performance is assessed based on these pairs. We primarily utilize three metrics for performance evaluation based on top-10 recommendation results: NDCG, HR, and MRR. Specifically, NDCG assesses the ranking quality of recommended items, HR measures the presence of at least one relevant item, and MRR evaluates the rank of the top relevant item.

4.2.1 Comparison Methods. We conduct a comprehensive comparison of UniRec with 11 baseline models. These include six classic sequential recommendation models: GRU4Rec [19], Caser [37], STAMP [26], SASRec [21], BERT4Rec [36], and LightSANS [7]. Additionally, we evaluate five time-aware models: TiSASRec [25], Meantime [4], TiCoSeRec[5], FEARrec [6], and MOJITO [38], all of which leverage temporal information to improve performance.

4.2.2 Implementation Details. All models are trained for up to 200 epochs utilizing the Adam optimizer [23]. Early stopping is implemented with a patience threshold of 20 epochs. We assign a value of 64 to the parameter d , utilize a batch size of 512, and set the learning rate to 0.01. The length of the sequence is fixed at 50. Both hyper-parameters M and K are set to 3. The mixture attention mechanism is configured with 2 heads. We test the partitioning ratios for uniform and non-uniform users within the range of {0.3, 0.4, 0.5, 0.6, 0.7, 0.8}, and for frequent and less-frequent items within the range of {0.4, 0.5, 0.6, 0.7, 0.8, 0.9}, across each dataset.

4.3 Overall Performance

Table 2 presents the experimental results of UniRec and 11 baselines across four datasets, several conclusions can be drawn. First, time-aware models generally outperform non-time-aware sequential recommendation models across various datasets. This highlights the critical importance of incorporating temporal dynamics into the recommendation process, as it substantially enhances the relevance and accuracy of the recommendations. Second, UniRec significantly




Figure 5: Ablation performance with various enhancements across different subsets from ML-1M.

outperforms other comparative models across all datasets and evaluation metrics, confirming its effectiveness. The bidirectional enhancement strategy for sequences and items adopted by UniRec, along with the multidimensional time modeling, greatly enhances the precision in modeling user interests and item characteristics. For instance, on the ML-1M dataset, UniRec achieves improvements of 3.32% in NDCG@10 and 4.08% in MRR@10 compared to the existing SOTA techniques. Third, UniRec demonstrates exceptional performance across datasets with varying sparsity and scale, whether in the lower-sparsity, smaller-scale ML-1M dataset or in the larger, more sparse Amazon datasets. This proves its adaptability and robustness to different levels of sparsity and data sizes. For example, on the Books dataset, UniRec increases MRR@10 by 3.24%, and on the Beauty dataset, it raises NDCG@10 by 3.01%. Lastly, compared to TiCoSeRec, which enhances data by improving sequence uniformity, UniRec enhances the utilization of sequence uniformity by incorporating item frequency more effectively. This demonstrates the potential of enhancing sequential recommendations from both perspectives of item frequency and sequence uniformity.

4.4 Ablation Experiment

To understand the impact of various components in our model, we conduct an ablation study. We divide the model into the following parts for evaluation: Multidimensional Time Modeling (A), Sequence Enhancement (B), Item Enhancement (C), and Item Popularity & Similarity (D). Specifically, w/o A refers to the replacement of multidimensional time modeling with a single-dimensional time modeling structure, utilizing only time interval modeling and disregarding contextual time information. w/o B refers to removing the sequence enhancement task, while w/o C refers to removing the item enhancement task. w/o D refers to excluding the consideration of item popularity and similarity in the item enhancement component, instead selecting candidate neighbors based solely on the time interval of the project. In addition to the overall dataset results, we evaluate performance on several subsets: frequent-item, less-frequent-item, uniform-sequence, and non-uniform-sequence. Using the ML-1M dataset as an example, Figure 5 shows the evaluation results of SASRec, UniRec, and UniRec without several components across various subsets.

First, UniRec demonstrates significant performance improvements over SASRec across all strategies, particularly in the less-frequent-item and non-uniform-sequence subsets. According to the experimental data, UniRec shows a 9.2% improvement in MRR@10 over SASRec in the frequent-item subset and an 18.0% improvement in the less-frequent-item subset. Additionally, in uniform and non-uniform subsets, UniRec achieves a 2.1% and 4.3% improvement in HR@10 over SASRec, respectively. These findings indicate that UniRec excels in enhancing performance for less-frequent items and non-uniform sequences.

Secondly, removing each component of the model results in varying degrees of performance degradation, indicating the importance of each component to the overall model performance. Particularly, w/o B leads to the most significant performance drop, particularly reflected in the HR metric, highlighting the effectiveness of the sequence enhancement module. This module not only improves the uniformity of non-uniform sequences but also increases the frequency of less-frequent items, significantly contributing to the accuracy of user interest modeling.

Furthermore, the performance on the frequent-item subset and uniform-sequence subset is consistent with the overall data. However, there are some differences between the less-frequent-item subset and the non-uniform-sequence subset. In the less-frequent-item subset, w/o A shows a significant drop in NDCG@10 and MRR@10, indicating that temporal information has a substantial impact on less-frequent items, as certain less-frequent items are more likely to be interacted with during specific periods. The declines in NDCG@10 and MRR@10 for w/o C and w/o D also demonstrate the effectiveness of these components in modeling less-frequent items. In particular, w/o D underscores the importance of considering item popularity, similarity, and relevance in selecting candidate neighbors to enhance less-frequent items' representations. In the non-uniform-sequence subset, the significant performance drop in w/o B indicates that sequence enhancement indeed improves the model's capability to handle sequences with rich interest drifts.

In summary, Figure 5 clearly illustrates the contributions of each component to the performance of UniRec, validating the necessity and effectiveness of multidimensional time modeling, sequence enhancement, item enhancement, and item popularity & similarity in improving the model's recommendation performance.




Figure 6: Performance comparison using different partition thresholds for item frequency and sequence uniformity on the Beauty dataset.




Figure 7: Time sensitivity comparison of uniform and non-uniform sequences on Amazon Sports and Amazon Industrial datasets.

4.5 Hyperparameter Experiment

In this subsection, we explore the relationship between the performance of UniRec and two hyperparameters: the item frequency partition threshold and the user uniformity partition threshold. As shown in Figure 6, we conduct experiments on the Amazon Beauty dataset, testing the impact of item frequency partition thresholds ranging from 40% to 90% (a), and sequence uniformity partition thresholds ranging from 30% to 80% (b). The results indicate that all tested partition thresholds yield good performance, but the most significant improvement occurs at specific values. For the Beauty dataset, the optimal split thresholds are 70% for high-frequency items and 30% for less-frequent items, while the ratio of uniform to non-uniform sequences is 60% to 40%. In summary, UniRec exhibits robust performance across different threshold settings, yet carefully selecting division thresholds can enhance the performance the most.

4.6 Time Sensitivity Analysis

As mentioned in section 3, we hypothesize that uniform sequences and non-uniform sequences may exhibit different dependencies on temporal information. In this subsection, to validate this hypothesis, we compare the effects of coarse-grained time modeling and fine-grained time modeling on both uniform and non-uniform sequence subsets. As shown in Figure 7, a positive score indicates that coarse-grained modeling outperforms fine-grained modeling, while the negative indicates the opposite. In both Amazon datasets, we observe that coarse-grained modeling performs better on uniform-sequence subsets, whereas fine-grained modeling is more effective

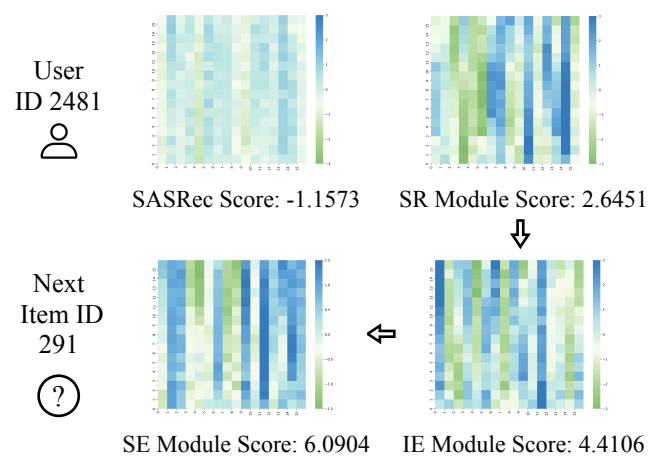


Figure 8: Prediction scores and corresponding sequence embedding heatmaps of a non-uniform sequence across different models and modules.

on non-uniform-sequence subsets. For uniform sequences, user behavior patterns are more consistent, capturing global patterns can yield satisfactory predictive outcomes. Conversely, non-uniform sequences exhibit greater diversity and dynamism in user behavior, necessitating a fine-grained temporal encoding strategy to accurately model shifts and changes in user interests.

4.7 Case Study

We conduct a case study to illustrate the progressive enhancement of a non-uniform sequence through various models and modules. As shown in Figure 8, we select a non-uniform sequence (user ID 2481) and demonstrate the changes in prediction scores for the next item (item ID 291) and the corresponding sequence embeddings after modeling with four different approaches: SASRec, SR module of UniRec, both the SR and IE modules, and the SR, IE, and SE modules. The progression of the model incorporating more modules is indicated by the arrows in the figure. SASRec shows a low prediction score, indicating its limited capability in handling sequences with significant interest drift. Adding the SR module significantly improves the model’s predictive ability. The inclusion of the IE module brings further improvement, and the model achieves its best performance with the addition of the SE module. In the heatmaps, blue indicates larger positive values and green indicates smaller negative values. The transition in heatmap colors from SASRec to the enhanced models, with increasing contrast, demonstrates the model’s growing ability to capture detailed information and features from various positions within the sequence.

5 RELATED WORKS

5.1 Sequential Recommendation

Sequential recommendation systems identify patterns in user behavior to predict future actions. Initially, Markov models [33, 20] are pivotal for analyzing transitions between states. The rise of deep learning leads to RNN models like GRU4Rec [19], which improves predictions by capturing long-term dependencies [14, 15, 16].

Convolutional Neural Network (CNN [24])-based models, such as Caser [37], improves recommendations by examining local behavior sequence patterns. Models like SHAN [35] and STAMP [26] effectively address shifts in user interests through memory strategies. Recently, attention mechanisms and Transformer-based models, like SASRec [21] and Bert4Rec [36], have gained prominence. They leverage self-attention to understand complex sequence dependencies, while LightSANs [7] introduces lightweight self-attention structures. The SSE-PT [41] integrates personalized embeddings with Stochastic Shared Embeddings (SSE) [42]. Research also extends to cross-domain [49, 1, 11], interpretable [48, 18], graph neural network [17, 46, 2, 29], and contrastive learning approaches [28, 47, 44, 5] for sequential recommendations.

5.2 Time-Aware Sequential Recommendation

Time-aware systems incorporate timing to capture the dynamic nature of user preferences, offering more accurate and timely recommendations. These models surpass traditional ones by adapting recommendations to both the shifts in user preferences over time and their current interests [45, 8]. The TiSASRec [25] model innovatively adjusts self-attention weights based on the timing between actions, significantly improving performance. MEANTIME [4] enriches time perception through diverse embedding techniques, whereas TASER [45] explores both absolute and relative time patterns. TGSRec [9] considers temporal dynamics in sequence patterns, and MOJITO [38] analyzes preferences from various temporal perspectives through a hybrid self-attention mechanism. FEARec [6] transitions sequence analysis from the time to the frequency domain, employing a hybrid attention mechanism and multitask learning for enhanced performance.

While these models ingeniously integrate temporal information, optimizing the use of such data remains a challenge. The diversity of data characteristics necessitates adaptable approaches for handling time intervals, timestamps, and cyclic patterns, given the varied and often irregular temporal behavior patterns among users. Recently, the TiCoSeRec [5] introduces an innovative approach by considering sequence uniformity during the data augmentation phase, marking a deeper understanding of sequential recommendation data. While this model treats sequence uniformity as a target of data enhancement, it does not delve into modeling and analyzing this characteristic of the data further. In contrast, in this paper, we incorporate sequence uniformity into model construction. Our method not only addresses the limitations encountered by existing models when dealing with data of varied temporal distributions but also proposes a novel perspective for feature enhancement.

6 CONCLUSION

In this paper, we demonstrate that sequential recommendation algorithms perform better on uniform sequences and frequent items compared to non-uniform sequences and less-frequent items. To address this, we present a novel bidirectional enhancement architecture that leverages sequence uniformity and item frequency for feature enhancement, optimizing the performance of sequential recommendations. Additionally, we introduce a multidimensional time modeling method to better capture temporal information. Experimental results show that our method significantly outperforms

twelve competitive models across four real-world datasets. To the best of our knowledge, this is the first work that utilizes the uniformity of sequences and frequency of items to enhance recommendation performance and it also indicates a promising direction and a new perspective for feature enhancement in future research.

ACKNOWLEDGMENTS

This paper is sponsored by anonymous with Grant No. XXXXXXXXX.

REFERENCES

- [1] Jiangxia Cao, Xin Cong, Jiawei Sheng, Tingwen Liu, and Bin Wang. 2022. Contrastive cross-domain sequential recommendation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 138–147.
- [2] Jianxin Chang, Chen Gao, Yu Zheng, Yiqun Hui, Yanan Niu, Yang Song, Depeng Jin, and Yong Li. 2021. Sequential recommendation with graph neural networks. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, 378–387.
- [3] Eunjoon Cho, Seth A. Myers, and Jure Leskovec. 2011. Friendship and mobility: user movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1082–1090.
- [4] Sung Min Cho, Eunhyeok Park, and Sungjoo Yoo. 2020. MEANTIME: Mixture of Attention Mechanisms with Multi-temporal Embeddings for Sequential Recommendation. In *Fourteenth ACM Conference on Recommender Systems*, 515–520.
- [5] Yizhou Dang, Enneng Yang, Guibing Guo, Linying Jiang, Xingwei Wang, Xiaoxiao Xu, Qinghui Sun, and Hong Liu. 2023. Uniform Sequence Better: Time Interval Aware Data Augmentation for Sequential Recommendation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37, 4, 4225–4232. Number: 4.
- [6] Xinyu Du, Huanhuan Yuan, Pengpeng Zhao, Jianfeng Qu, Fuzhen Zhuang, Guanfeng Liu, Yanchi Liu, and Victor S. Sheng. 2023. Frequency Enhanced Hybrid Attention Network for Sequential Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 78–88.
- [7] Xinyan Fan, Zheng Liu, Jianxun Lian, Wayne Xin Zhao, Xing Xie, and Ji-Rong Wen. 2021. Lighter and Better: Low-Rank Decomposed Self-Attention Networks for Next-Item Recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*, 1733–1737.
- [8] Ziwei Fan, Zhiwei Liu, Jiawei Zhang, Yun Xiong, Lei Zheng, and Philip S Yu. 2021. Continuous-time sequential recommendation with temporal graph collaborative transformer. In *Proceedings of the 30th ACM international conference on information & knowledge management*, 433–442.
- [9] Ziwei Fan, Zhiwei Liu, Jiawei Zhang, Yun Xiong, Lei Zheng, and Philip S. Yu. 2021. Continuous-Time Sequential Recommendation with Temporal Graph Collaborative Transformer. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 433–442.
- [10] Hui Fang, Danning Zhang, Yiheng Shu, and Guibing Guo. 2020. Deep learning for sequential recommendation: algorithms, influential factors, and evaluations. *ACM Transactions on Information Systems (TOIS)*, 39, 1, 1–42.
- [11] Lei Guo, Jinyu Zhang, Tong Chen, Xinhua Wang, and Hongzhi Yin. 2022. Reinforcement learning-enhanced shared-account cross-domain sequential recommendation. *IEEE Transactions on Knowledge and Data Engineering*.
- [12] F. Maxwell Harper and Joseph A. Konstan. 2015. The movielens datasets: history and context. arXiv:1503.0863. (2015). arXiv: 1503.0863 [cs. IR].
- [13] Ruining He and Julian McAuley. 2016. Ups and downs: modeling the visual evolution of fashion trends with one-class collaborative filtering. In *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 507–517.
- [14] Balázs Hidasi and Alexandros Karatzoglou. 2018. Recurrent Neural Networks with Top-k Gains for Session-based Recommendations. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, 843–852.
- [15] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939*.
- [16] Balázs Hidasi, Massimo Quadrana, Alexandros Karatzoglou, and Domonkos Tikk. 2016. Parallel Recurrent Neural Network Architectures for Feature-rich Session-based Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems*, 241–248.

- [17] Cheng Hsu and Cheng-Te Li. 2021. RetaGNN: Relational Temporal Attentive Graph Neural Networks for Holistic Sequential Recommendation. In *Proceedings of the Web Conference 2021*, 2968–2979.
- [18] Xiaowen Huang, Quan Fang, Shengsheng Qian, Jitao Sang, Yan Li, and Changsheng Xu. 2019. Explainable interaction-driven user modeling over knowledge graph for sequential recommendation. In *proceedings of the 27th ACM international conference on multimedia*, 548–556.
- [19] Dietmar Jannach and Malte Ludewig. 2017. When Recurrent Neural Networks meet the Neighborhood for Session-Based Recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*, 306–310.
- [20] Joseph Johnson and Yiu-Kai Ng. 2017. Enhancing long tail item recommendations using tripartite graphs and Markov process. In *Proceedings of the International Conference on Web Intelligence (WI '17)*, 761–768.
- [21] Wang-Cheng Kang and Julian McAuley. 2018. Self-Attentive Sequential Recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*, 197–206.
- [22] Yejin Kim, Kwangseob Kim, Chanyoung Park, and Hwanjo Yu. 2019. Sequential and diverse recommendation with long tail. In *IJCAI*. Vol. 19, 2740–2746.
- [23] Diederik P. Kingma and Jimmy Ba. 2014. Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [24] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature*, 521, 7553, 436–444.
- [25] Jiacheng Li, Yujie Wang, and Julian McAuley. 2020. Time Interval Aware Self-Attention for Sequential Recommendation. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, 322–330.
- [26] Qiao Liu, Yifu Zeng, Refuoe Mokhosi, and Haibin Zhang. 2018. STAMP: Short-Term Attention/Memory Priority Model for Session-based Recommendation. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1831–1839.
- [27] Siyi Liu and Yujia Zheng. 2020. Long-tail session-based recommendation. In *Proceedings of the 14th ACM Conference on Recommender Systems*, 509–514.
- [28] Zhiwei Liu, Yongjun Chen, Jia Li, Philip S Yu, Julian McAuley, and Caiming Xiong. 2021. Contrastive self-supervised sequential recommendation with robust augmentation. *arXiv preprint arXiv:2108.06479*.
- [29] Chen Ma, Liheng Ma, Yingxue Zhang, Jianing Sun, Xue Liu, and Mark Coates. 2020. Memory augmented graph neural networks for sequential recommendation. In *Proceedings of the AAAI conference on artificial intelligence* number 04. Vol. 34, 5045–5052.
- [30] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton van den Hengel. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 43–52.
- [31] Massimo Quadrona, Paolo Cremonesi, and Dietmar Jannach. 2018. Sequence-aware recommender systems. *ACM computing surveys (CSUR)*, 51, 4, 1–36.
- [32] Mostafa Rahmani, James Caverlee, and Fei Wang. 2023. Incorporating Time in Sequential Recommendation Models. In *Proceedings of the 17th ACM Conference on Recommender Systems*, 784–790.
- [33] 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 (WWW '10)*, 811–820.
- [34] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. *Proceedings of the 10th International Conference on World Wide Web*, 285–295.
- [35] 2018. Sequential recommender system based on hierarchical attention network. In *Sequential recommender system based on hierarchical attention network*. ISSN: 1045-0823, 3926–3932.
- [36] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 1441–1450.
- [37] Jiayi Tang and Ke Wang. 2018. Personalized Top-N Sequential Recommendation via Convolutional Sequence Embedding. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining (WSDM '18)*, 565–573.
- [38] Viet Anh Tran, Guillaume Salha-Galvan, Bruno Sguerra, and Romain Hennequin. 2023. Attention Mixtures for Time-Aware Sequential Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1821–1826.
- [39] Shoujin Wang, Longbing Cao, Yan Wang, Quan Z. Sheng, Mehmet A. Orgun, and Defu Lian. 2022. A Survey on Session-based Recommender Systems. *ACM Computing Surveys*, 54, 7, 1–38.
- [40] Shoujin Wang, Liang Hu, Yan Wang, Longbing Cao, Quan Z Sheng, and Mehmet Orgun. 2019. Sequential recommender systems: challenges, progress and prospects. *arXiv preprint arXiv:2001.04830*.
- [41] Liwei Wu, Shuqing Li, Cho-Jui Hsieh, and James Sharpnack. 2020. SSE-PT: Sequential Recommendation Via Personalized Transformer. In *Fourteenth ACM Conference on Recommender Systems*, 328–337.
- [42] Liwei Wu, Shuqing Li, Cho-Jui Hsieh, and James L Sharpnack. 2019. Stochastic Shared Embeddings: Data-driven Regularization of Embedding Layers. In *Advances in Neural Information Processing Systems*. H. Wallach, H. Larochelle, A. Beygelzimer, F. d' Alché-Buc, E. Fox, and R. Garnett, (Eds.) Vol. 32.
- [43] Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kamran Achan. 2019. Self-attention with functional time representation learning. *Advances in neural information processing systems*, 32.
- [44] Yuhao Yang, Chao Huang, Lianghao Xia, and Chenliang Li. 2022. Knowledge graph contrastive learning for recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, 1434–1443.
- [45] Wenwen Ye, Shuaiqiang Wang, Xu Chen, Xuepeng Wang, Zheng Qin, and Dawei Yin. 2020. Time Matters: Sequential Recommendation with Complex Temporal Information. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1459–1468.
- [46] Yaowen Ye, Lianghao Xia, and Chao Huang. 2023. Graph Masked Autoencoder for Sequential Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 321–330.
- [47] Junliang Yu, Hongzhi Yin, Xin Xia, Tong Chen, Jundong Li, and Zi Huang. 2023. Self-supervised learning for recommender systems: a survey. *IEEE Transactions on Knowledge and Data Engineering*.
- [48] Yongfeng Zhang, Xu Chen, et al. 2020. Explainable recommendation: a survey and new perspectives. *Foundations and Trends® in Information Retrieval*, 14, 1, 1–101.
- [49] Feng Zhu, Yan Wang, Chaochao Chen, Jun Zhou, Longfei Li, and Guanfeng Liu. 2021. Cross-domain recommendation: challenges, progress, and prospects. *arXiv preprint arXiv:2103.01696*.