



A survey on sequential recommendation

Li-Wei PAN¹, Wei-Ke PAN¹✉, Mei-Yan WEI¹, Hong-Zhi YIN², Zhong MING³

1. College of Computer Science and Software Engineering, Shenzhen University, Shenzhen 518060, China

2. School of Electrical Engineering and Computer Science, The University of Queensland, Queensland 4072, Australia

3. College of Big Data and Internet, Shenzhen Technology University, Shenzhen 518118, China

Received July 12, 2024; accepted March 11, 2025

E-mail: panweike@szu.edu.cn

© The Author(s) 2025. This article is published with open access at link.springer.com and journal.hep.com.cn

Abstract

Different from most conventional recommendation problems, sequential recommendation (SR) focuses on learning users' preferences by exploiting the internal order and dependency among the interacted items, which has received significant attention from both researchers and practitioners. In recent years, we have witnessed great progress and achievements in this field, necessitating a new survey. In this survey, we study the SR problem from a new perspective (i.e., the construction of an item's properties), and summarize the most recent techniques used in sequential recommendation such as multi-modal SR, generative SR, LLM-powered SR, ultra-long SR, and data-augmented SR. Moreover, we introduce some frontier research topics in SR, e.g., open-domain SR, data-centric SR, cloud-edge collaborative SR, continuous SR, SR for good, and explainable SR. We believe that our survey could be served as a valuable roadmap for readers in this field.

Keywords

sequential recommendation; ID-based; side information; recent advancements; new problems

1 Introduction

Recommender systems are often designed to deliver items that users are interested in from a large collection, which are expected to address the information overload problem and save time for users. Meanwhile, recommender systems are usually deemed effective in improving business profits by motivating people to purchase their items of interest. So far, they have been deployed in various real-world applications, e.g., e-commerce (e.g., Amazon and Alibaba), streaming services (e.g., YouTube and TikTok), social media (e.g., WeChat and Twitter), online advertising.

Various sequential recommendation (SR) models have been proposed, which have achieved significant recommendation improvements in performance recently [1]. The main idea of an SR model is to exploit the position and sequential information of the interacted items, so as to capture the dependencies among them. Given a user's interaction sequence, a typical SR model aims to predict the next likely-to-be-preferred items for this user. In recent years, there has been a significant increase in the number of SR models. By searching the keywords (i.e., sequential, recommend) in the DBLP website¹⁾, we have the specific publication numbers of papers about sequential recommendation, which are shown in Fig. 1. We can see that the number of papers about sequential

recommendation increases with time going by.

Some early SR models [1,2] use a unique ID to denote an item. These SR models are simple yet effective. For each item, an SR model usually needs to learn a unique item embedding, which is often efficient. Meanwhile, the item embeddings can capture the correlations among the items and extract their latent characteristics. However, these SR models have some weaknesses. For instance, only using item IDs cannot solve cold-start and sparsity problems. Therefore, some SR models [3,4] combine item IDs and item features. The item features mainly include categorical features (e.g., categories), numerical features (e.g., prices) and graph structure features (e.g., social networks and knowledge graph). Due to privacy issues, we often know little about the users' features. By leveraging these features, we can effectively learn item representations even when interactions with some items are sparse. Therefore, these SR models can alleviate the cold-start and data sparsity problems. However, these SR models also have some limitations. For example, they usually heavily rely on item IDs during training. If data are different, the item IDs are different. Therefore, these SR models trained on one data cannot be easily transferred to another data through fine-tuning or other methods. With the fast development of natural language processing and computer vision, models like BERT

¹⁾ Website at dblp.org/ (Last visited on May 12, 2025).




Fig. 1 The number of articles on sequential recommendation published between 2015 and 2024

[5] and ResNet [6] can extract rich features from text descriptions and images, respectively. To solve the issue of data dependency and high-difficulty in cross-data transfer, some SR models [7,8] leverage BERT and ResNet to extract semantic representations of items from text descriptions and images, respectively, and then directly use them to represent the items. Due to the widespread availability of multi-modal features, these SR models trained on one data from a certain platform can be deployed on another platform through fine-tuning. However, the semantic representations extracted from multi-modal features cannot fully capture users' fine-grained preferences and collaborative signals. Therefore, some SR models turn to combine the multi-modal features with item IDs to learn users' preferences more comprehensively.

Large language models (LLMs) are trained using a large amount of data and demonstrate promising performance. More importantly, LLMs have a strong reasoning ability. Therefore, LLMs have been applied to many applications, including recommender systems. For example, some SR models leverage LLMs to recommend items to users directly [9]. Some SR models leverage the output representations of LLMs for recommending personalized items more accurately since these representations contain rich semantic knowledge [10]. To alleviate data sparsity and cold-start problems, some SR models leverage LLMs to generate some data. These SR models then combine the generated data with the original data to achieve better recommendation performance. However, LLMs-based SR models also have some shortcomings. For example, training LLMs-based SR models requires significant computational resources. This survey summarizes how to leverage LLMs in SR models.

In sequential recommendation, modeling ultra-long interaction sequences becomes increasingly complex. Over time, users interact with more and more items. With the length of interaction sequences increasing, the time of both training and inference will increase. Meanwhile, as users interact with an increasing number of items, noise also becomes more serious in the interaction sequences. Therefore, modeling ultra-long interaction sequences becomes increasingly difficult. Common approaches retrieve a few interacted items from ultra-long interaction sequences [11]. Our survey will summarize some common methods applied to modeling ultra-long interaction sequences modeling.

In our survey, we also introduce data augmentation methods in SR

models. This is because data augmentation methods can be leveraged by different SR models. They mainly focus on generating more interaction data to alleviate data sparsity and cold-start problems. Common augmentation methods include some operations, such as resorting, masking, cropping, and others on interaction sequences [12].

The contributions of our survey are summarized as follows:

- Our survey summarizes the early and recent works about SR models from a more comprehensive manner.
- Our survey categorizes the existing SR models into four categories according to the construction of item properties.
- Our survey summarizes the latest techniques applied to sequential recommendation.
- Our survey introduces empirical studies and future research directions in SR models.

The remainder of our survey is organized as follows. In Section 2, we discuss some surveys on sequential recommendation and session-based recommendation. In Section 3, we provide a comprehensive view about sequential recommendation, including problem definition, properties, four different kinds of SR models, and challenges. In Section 4, we discuss some pure ID-based SR models. In Section 5, we present some SR models combining IDs and side information. In Section 6, we elaborate on recent advancements in sequential recommendation, including multi-modal SR, generative SR, LLM-powered SR, ultra-long sequence modeling in SR, and data augmentation in SR. In Section 7, we introduce datasets, evaluation protocols and metrics, as well as experimental results in sequential recommendation. Finally, we provide some insights into prospects and future directions about sequential recommendation in Section 8 and draw a conclusion in Section 9.

■ 2 Related surveys

In recent years, an increasing number of surveys on sequential and session-based recommender systems have been published [13–18]. They introduce some common algorithms and models for these systems. However, with the development of LLMs and other techniques, these surveys are out-of-date to some extent as LLMs have strong deduction abilities and can fully extract semantic information for recommendation. Meanwhile, some SR models introduced in these surveys still have weaknesses. For example, only using an item ID to denote an item might influence the transferability of SR models. Some surveys [19,20] focus on cross-domain sequential recommender systems, or multi-behavior sequential recommender systems. They only introduce a small fraction of models in sequential recommendation. Compared with these surveys, our survey introduces the development of sequential recommender systems more comprehensively. There are also some surveys about recommender systems. Jing et al. [21] introduced contrastive learning in recommender systems comprehensively. Lai et al. [22] discussed the importance of data and some methods of data augmentation. Zhang et al. [23] explored the usage of different kinds of side information in recommender systems. However, these surveys only

focus on one specific technique in recommender systems. Compared with these surveys, we not only classify existing research on SR but also introduce a new taxonomy to comprehend, structure, and analyze these works.

3 Sequential recommendation

In this section, we first define the problem of sequential recommendation. Then, we introduce properties about users, items, behaviors, sequences, as well as different kinds of features. Next, we describe four different types of SR models. Finally, we discuss some challenges faced by SR models.

3.1 Problem definition

In sequential recommendation, we often have one or more sequences of interacted items w.r.t. each user, as well as some auxiliary information to help learn user preferences. Our goal is then to generate a ranked list of items accurately for each user.

3.2 Properties

In sequential recommendation systems, there are two important entities, an item and a user. Therefore, item properties and user properties are vital for SR models. A user interacts with items, establishing connections between the user and the items via different types of behaviors. The user's interaction sequence is constructed according to interaction order. Therefore, the properties of both the behaviors and sequences are also important in sequential recommendation systems. The specific details about these properties in sequential recommendation systems are shown in Fig. 2.

For item properties in SR, each item is denoted by a unique number, i.e., an item ID. Apart from the item ID, each item also has side information. The side information mainly has two types, including multi-modal features and general features. The general features have three types: categorical features, numerical features, and knowledge graph. The categorical features can be denoted by one-hot vectors. The numerical features are values represented by numbers, which have real meaning. Each numerical feature usually lies in a continuous or discrete range. User properties consist of two parts: a unique user ID and some side information specific to each user. To avoid the leakage of users' privacy, the side information about users is usually unseen. Behavior properties are also comprised of general features and multi-modal features. Some features such as ratings and reviews, which might not be present in certain interactions, because not every user can provide a review or rating for an item they interact with. However, if a user interacts with an item, the behavior type of this interaction is bound to exist. Context information refers to an environment in which a behavior happens, such as a specific time and location. Sequential properties primarily include the length of interaction sequences and the internal order of items within these sequences. Meanwhile, there exist two important graph-structure features, knowledge graph and social networks.

Entity IDs (i.e., item IDs and user IDs), general features and multi-modal features are different in sequential recommender systems. The entity IDs are specific to a particular data. In different data or different platforms, the encoding methods for the entity IDs are quite different. For example, a movie might be encoded as item 1 in

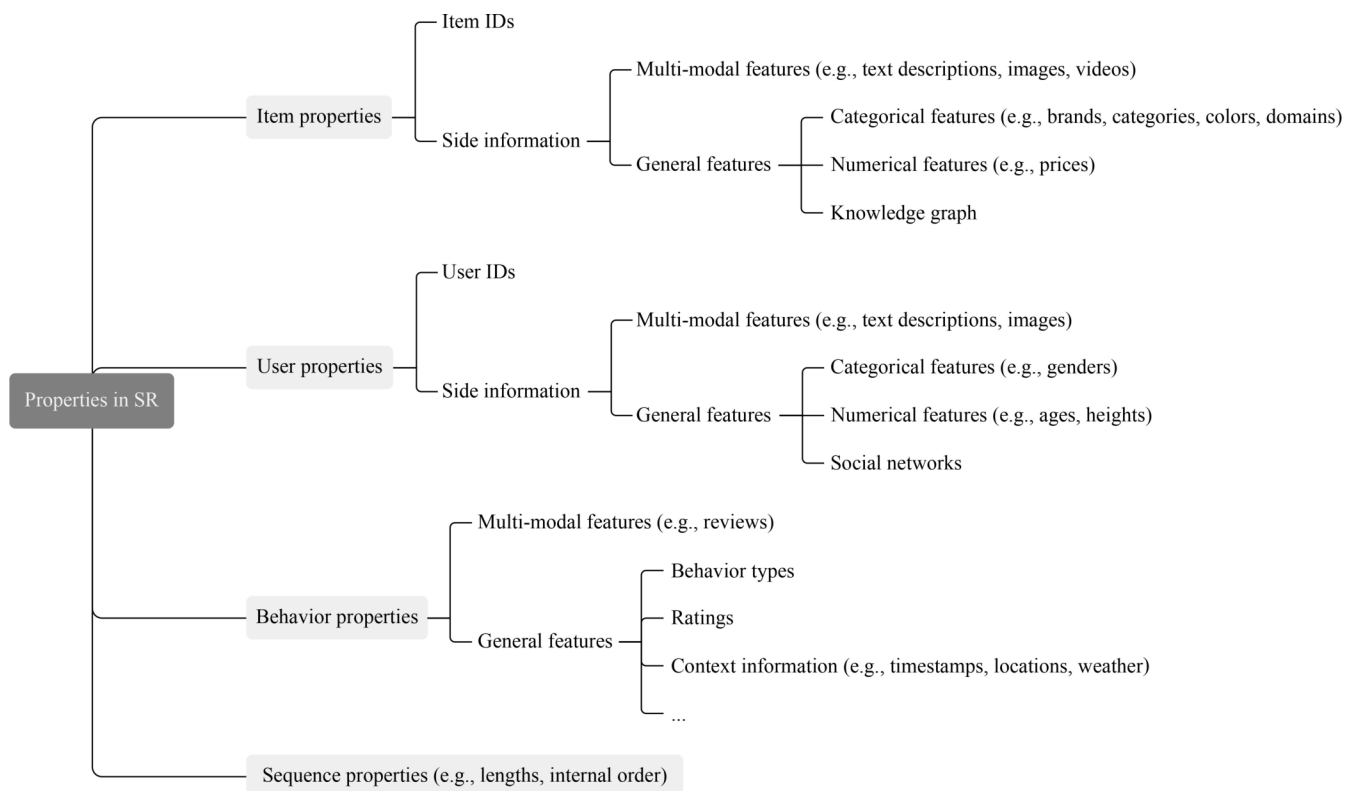


Fig. 2 Properties in sequential recommendation

MovieLens²). However, a book might also be encoded as item 1 in the book domain of Amazon³). Therefore, item embeddings based on item IDs cannot be transferred from one platform to another platform. The general features share some common characteristics across different datasets or platforms. For example, basic behavior types in different datasets might be the same. All e-commerce platforms contain basic behavior types (i.e., click and purchase). However, the categories of items might differ across different domains or platforms, leading to different distributions of general features. If we apply the embeddings of general features in one domain to another domain, it might result in negative transfer. Techniques like a cross-attention mechanism can address this issue to some extent. Multi-modal features, such as text, images, and videos, are universal across all domains and platforms. We use language to record information as text, cameras to capture beautiful images, videos to document moments. The text, images, and videos have separate embedding spaces. Meanwhile, they all have shared space. Therefore, the multi-features can be leveraged in all domains and all platforms. If we leverage the text descriptions of an item to represent the item in a model and train the model on one platform, it can be directly applied to another platform or fine-tuned in another platform as needed. Therefore, leveraging multi-modal features to denote an item can fully realize the transfer of SR models from one platform to another.

In the meantime, the general features and multi-modal features can be converted into each other. As shown in Fig. 3, we can convert an item's brand, price, and color into text descriptions. Then, SR models can leverage these text descriptions of this item to learn the universal representations of this item. However, this type of conversion has shortcomings. For example, if we convert timestamps to text, fine-grained time information might be damaged. Meanwhile, we can also extract general features from the text descriptions of an item, such as its brand, color, and price. However, the extracted general features might contain some noise.

There are several important features in sequential recommender systems. Different features have various meanings and play different roles. Some specific details are shown as follows:

- Timestamp refers to the precise time when a user interacts with an item. Leveraging timestamp information can capture the user's evolving interests and periodic behaviors precisely.

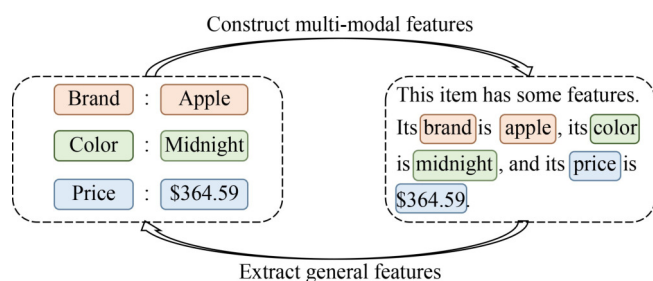


Fig. 3 The conversion between multi-modal features and general features

- Category refers to the hierarchical classification of an item. Leveraging the category of the item can capture users' preferences more precisely and narrow the range of candidate items.
- Price refers to the amount of money spent to purchase an item. It plays a dominant role in deciding whether to purchase the item or not.
- Text description refers to the description of an item and the title of the item. It can describe the item from different perspectives.
- Image refers to the visual representations of an item. It can provide a clear and intuitive understanding of the item. By looking at images, users can know what the item looks like.
- Rating refers to the degree of a user's satisfaction with an item. It can reflect the user's preferences for the item and the quality of the item.
- Review refers to comments that a user provides for an item. It can offer more detailed information about the user's preferences for the item, and it can reflect the overall quality of this item.
- Behavior type refers to a specific behavior type in which a user interacts with an item, such as click, favorite and purchase. Leveraging different types of behaviors can capture the user's different preferences from different perspectives.

3.3 Categories

In our surveys, we classify SR models into four categories based on the properties of items. As depicted in Fig. 4, for user u , there are three different interaction sequences containing item IDs, multi-modal features, and categorical and numerical features, respectively.

As shown in Fig. 4, some early works leverage a unique item ID to denote an item for recommendation [1,24]. These SR models leverage only item ID interaction sequences, which can be formalized as: $P(i|S_u)$. All notations used in our survey are shown in Table 1. In these works, given a candidate item i and a user u 's historical interaction sequence S_u , these models need to calculate the probability $P(i|S_u)$ in which user u interacts with the candidate item i . We will introduce this type of SR models in Section 4.

However, SR models based on pure item IDs cannot tackle the commonly encountered cold-start and data-sparsity problems well. General features are important for recommending personalized items, which can alleviate the cold-start and data-sparsity problems and fully extract users' preferences. As depicted in Figs. 4(a) and (c), some works leverage both item ID sequences and interaction sequences containing general features (i.e., categorical features and numerical features) to recommend personalized items to users [4,25]. These SR models can be formalized as: $P(i|S_u, N_u^{(1)}, N_u^{(2)}, \dots, N_u^{(K_g)})$. Given a user u 's interaction sequence S_u , which contains item IDs and general features $N_u^{(1)}, \dots, N_u^{(K_g)}$ embedded in the sequence, these models aim to predict the probability in which the user u interacted with a candidate item i . To account for the

²) Website at grouplens.org/datasets/movielens/ (Last visited on May 12, 2025).

³) Website at huggingface.co/datasets/McAuley-Lab/Amazon-Reviews-2023 (Last visited on May 12, 2025).

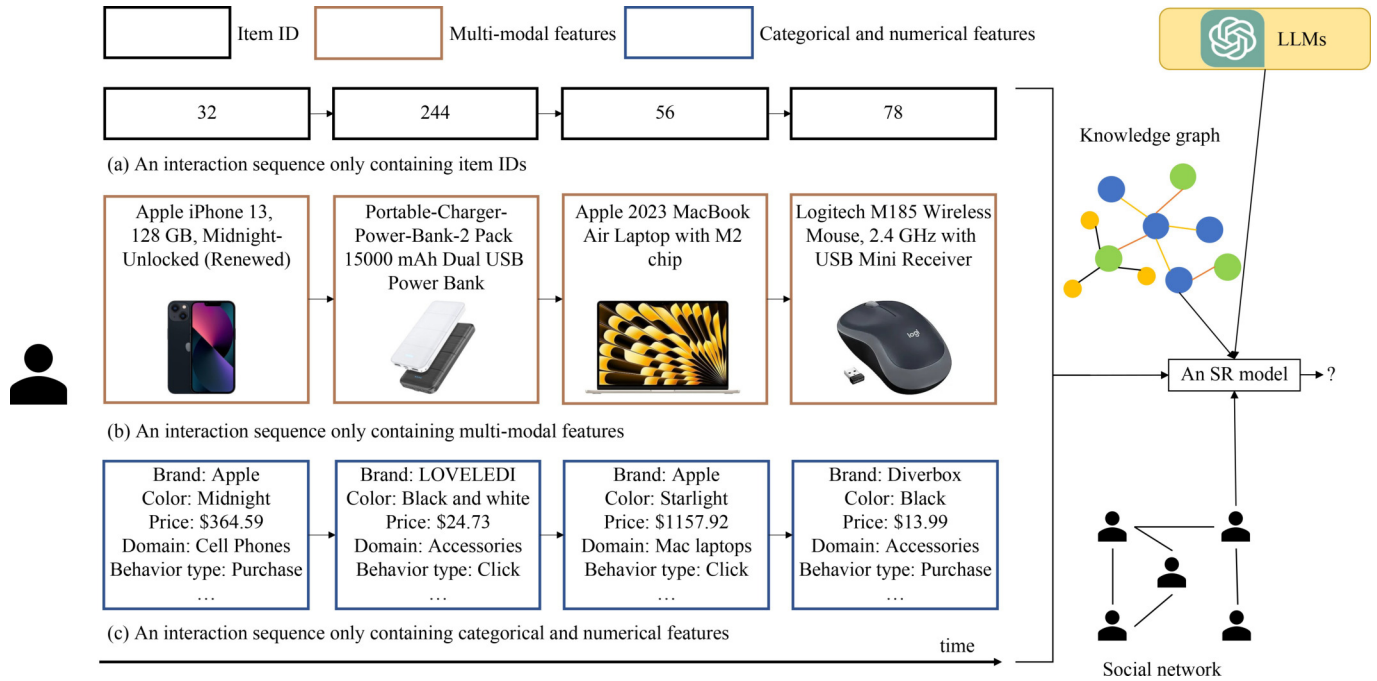


Fig. 4 An illustration of sequential recommendation

Table 1 Some notations used in the paper

Notation	Description
$\mathcal{U} = \{u_1, u_2, \dots, u_{ \mathcal{U} }\}$	A set of users
$\mathcal{I} = \{i_1, i_2, \dots, i_{ \mathcal{I} }\}$	A set of items
L	The length of an interaction sequence
$S_u = \{s_{u,1}, s_{u,2}, \dots, s_{u,L}\}$	The interaction sequence of user u
$F_u^{(k)} = \{f_{u,1}^{(k)}, f_{u,2}^{(k)}, \dots, f_{u,L}^{(k)}\}$	The k th multi-modal feature w.r.t. user u 's interaction sequence.
$N_u^{(k)} = \{n_{u,1}^{(k)}, n_{u,2}^{(k)}, \dots, n_{u,L}^{(k)}\}$	The k th categorical and numerical feature w.r.t. user u 's interaction sequence.
K_m	The number of multi-modal features
K_g	The number of categorical and numerical features
t_u	The item that user u interacts with at the most recent time step
$R_{u,i}$	The rank generated by an SR model of candidate item i for user u
$P(i S_u)$	The probability in which user u interacted item i given historical sequence S_u
GNNs(-), CNNs(-), Transformer(-)	The GNNs, CNNs, and Transformer models
Embed(-)	The item embedding layer
$\mathbf{p}_u^{\text{CNNs}}, \mathbf{p}_u^{\text{GNNs}}$	The user u 's preferences extracted from CNNs and GNNs
\mathbf{p}_u^{Tr}	The user u 's preferences extracted from Transformer
$\mathbf{p}_u^{\text{final}}$	The user u 's final preferences
\mathcal{G}	The graph constructed by all users' interaction sequences
Fuse(-)	The fusion module
$\text{BM}_k(\cdot)$	The k th backbone model
E_u	The item embedding matrix for user u
$A_u^{(k)}$	The k th feature embedding matrix for user u

Table 1 (Continued)

Notation	Description
$S_{u,1}^{(k)}, S_{u,2}^{(k)}, \dots, S_{u,m}^{(k)}$	The m sub-sequences divided by the k th feature
$\mathbf{p}_u^{\text{feat}}$	the u 's preferences extracted from the features of interacted items
$\text{Text}_i, \text{Img}_i, \text{Vid}_i$	The text descriptions, images, and videos of item i
$\text{TextEncoder}(), \text{ImageEncoder}(), \text{VideoEncoder}(\cdot)$	The text encoder, image encoder, and video encoder
\mathbf{e}_i	The representations about item i
$\mathbf{Y}_i = \{y_{i,1}, y_{i,2}, \dots, y_{i,l}\}$	The l tokens representing item i
x_{prompt}	The prompt which is fed into LLMs
$\text{LLMs}(\cdot)$	The large language models
\hat{x}_{prompt}	The outputs of LLMs while inputting x_{prompt}
$\mathbf{p}_u^{\text{short}}, \mathbf{p}_u^{\text{long}}$	The user u 's short-term and long-term interests
$\text{Att}(\cdot), \text{Ret}(\cdot)$	The attention and retrieval modules
$S_{u,\text{recent}}$	The user u 's recently behavior sequence

influence of users' friends, social networks might be leveraged in these works [26]. The general features of interacted items can be represented as knowledge graph. Therefore, these works might leverage knowledge graph to capture relationships among general features. We will introduce this type of SR models in Section 5.

Though the above-mentioned models achieve significant performance, they still have certain shortcomings. Firstly, these models heavily rely on item IDs, which hinders their transferability. Secondly, these general features might not be fully utilized by some of latest techniques effectively, such as LLMs. Therefore, some recently proposed models leverage the multi-modal features of an item to denote the item directly [8,27]. These models can realize the transferability across different platforms. By leveraging latest techniques like LLMs, they can fully extract some semantic knowledge. These models can be formalized as: $P(i|F_u^{(1)}, F_u^{(2)}, \dots, F_u^{(K_m)})$. Given the multi-modal features derived from the items within a user u 's interaction sequence, these models can calculate the probability in which user u interacts with a candidate item i . The core ideas of these models are illustrated in part (b) of Fig. 4. When both sequence information and multi-modal features (i.e., text descriptions and images) of each interaction are available, an SR model predicts the next interacted item interacted with by user u . By replacing item IDs with text descriptions and images, SR models can acquire rich semantic knowledge and enhance their transferability.

However, relying solely on items' multi-modal features cannot capture users' fine-grained interests. In a warm-start situation, only leveraging item IDs can achieve better performance than only leveraging multi-modal features [7]. Therefore, some SR models leverage both item IDs and multi-modal features of these items. The simple diagram of these models is shown by combining part (a) and part (b) in Fig. 4. In these SR models, each item property contains both a unique item ID and multi-modal features. To extract semantic information from multi-modal features, LLMs can be leveraged.

These SR models can be formalized as: $P(i|S_u, F_u^{(1)}, F_u^{(2)}, \dots, F_u^{(K_m)})$. We will introduce the aforementioned two types of SR models in Subsection 6.1.

3.4 Taxonomy

Our survey introduces SR models through a comprehensive taxonomy. As shown in Fig. 5, our survey will introduce SR models from three different aspects, i.e., pure ID-based SR, SR with side information, and recent SR advancements. Our survey will discuss specific details about these SR models in the following sections.

3.5 Challenges

In sequential recommender systems, there are some challenges as follows:

- Cold-start and data sparsity problems: A new user and a new item usually have sparse interactions. It is a challenge to improve recommendation performance in a new user and a new item.
- Ultra-long interaction sequences modeling: Over time, some users might generate ultra-long interaction sequences. It is a challenge to model these ultra-long interaction sequences effectively and efficiently.
- Dynamic interests: Users' interests evolve over time. It is a challenge to capture users' current and long-term interests.
- Diverse interests: Users exhibit diverse interests across different items. It is a challenge to capture users' diverse interests.
- Knowledge transfer across platforms: Items might be quiet different across different platforms. It is a challenge to achieve knowledge transfer across different platforms.

4 Pure ID-based SR

In SR models based on pure IDs, each item and user are represented by a unique item identity (i.e., item ID) and a unique user identity



Fig. 5 Taxonomy of existing works on sequential recommendation

(i.e., user ID), respectively. In the past few years, the number of SR models based on pure IDs has increased significantly. We categorize them into different classes according to their techniques. In this section, we first introduce some SR methods using traditional machine learning. Then, we introduce those integrating various deep learning techniques. Next, we introduce those using reinforcement learning. Finally, we introduce those combining different kinds of deep learning techniques. In this section, all SR models are designed based on pure IDs.

4.1 Traditional models

Traditional SR models do not use deep learning techniques. Though

these models are simple and effective, they cannot capture long-term sequential dependencies. The traditional SR models can mainly be classified into three classes: frequent pattern mining, Markov models, and latent factor models.

FPM (frequent pattern mining) extracts different kinds of patterns from interaction sequences [28]. These models are of good explainability, but they cannot tackle complex data. If a data is too large, the number of extracted sequential patterns might become excessively large.

Markov models have an assumption that a predicted action relies on a few recent actions. For example, FPMC [29] predicts the next

interacted item for each user by considering both the last interacted item and personalized interests. However, these models cannot exploit long-term sequential dependencies. If the order dependencies are weak in data, Markov models might perform poorly.

Latent factor models mainly learn latent representations to estimate unobserved transitions. TransRec [2] proposes a transition operator: a previously interacted item + a user \approx the user next interacted item. However, the transition operator might vary across different data. Meanwhile, these SR models cannot capture complex transition relations.

4.2 Deep learning models

4.2.1 Recurrent neural networks

Compared with traditional SR models, RNNs-based SR models can not only capture short-term sequential dependencies, but also capture long-term sequential dependencies and non-linear dynamics. In recent years, some RNNs-based SR models have been proposed. GRU4Rec [24] applies RNNs to capture long-term dependencies in session-based data. Based on GRU4Rec, GRU4Rec+ [30] proposes a novel negative item sampling and a novel ranking loss function. Some RNNs-based SR models are designed to capture each sequence independently, thus ignoring the global structure of all sequences. To solve this problem, KrNN-P [31] incorporates neighbor sequences into RNNs to capture global relationships while refining local ones. To consider the dependencies among users' sequences, MrRNN [32] introduces manifold regularization into RNNs based on the multi-facets of collaborative filtering. However, these RNNs-based models have some shortcomings. For example, they cannot avoid vanishing gradients while handling ultra-long interaction sequences.

4.2.2 Convolutional neural networks

Convolutional neural networks (CNNs) have been widely used in image processing. They mainly consist of two operations, i.e., a convolution operation and a pooling operation. Compared with multi-layer perceptron (MLPs), CNNs can extract features from images more effectively. By regarding user u 's interaction sequence as an image, CNNs can also be applied to sequential recommendation. CNNs usually have multiple convolutional filters. Therefore, CNNs-based SR models can extract multi-faceted interests from user u 's interaction sequence.

In recent years, several CNNs-based SR models have been proposed. DACNN [33] leverages CNNs to capture local sequential features while considering item sequences. Caser [34] leverages convolutional filters to capture both general preferences and sequential patterns. CosRec [35] leverages 2D convolutional neural networks to model complex interaction relationships. However, CNNs-based SR models have certain shortcomings. For example, they might have a large number of convolutional filters, which will result in excessive computational resource consumption.

4.2.3 Transformer

Transformer has accomplished great achievements in natural language processing, largely due to its self-attention mechanism. The self-attention mechanism can identify crucial information by calculating attention scores.

Transformer architecture has two types: single-direction Transformer (i.e., uni-Transformer) and double-direction Transformer (i.e., bi-Transformer). SASRec [1] is currently the most popular SR model, which leverages the uni-Transformer. SASRec predicts the next interacted items based on historical interacted items. BERT4Rec [36] leverages bi-Transformer, which predicts the next interacted items according to historical interacted items and future interacted items. ALBERT4Rec and DeBERTa4Rec [37] demonstrate superior performance and training efficiency compared to BERT4Rec [36]. ADT [38] learns disentangled diverse interests using a Transformer-based encoder and decoder. However, the large weights of attention mechanism might not be accurate. Additionally, position encoding and noise inputs might negatively impact the recommendation performance. Therefore, AC-TSR [39] leverages contrastive learning to achieve noise reduction. DT4IER [40] leverages decision Transformer to solve user retention challenges. Due to the uncertainty of users' sequential behaviors, STOSA [41] leverages stochastic Gaussian distribution to embed items and introduces a novel Wasserstein self-attention module to capture item-item relationships. MFGAN [42] uses a Transformer-based generator to recommend the next possible items and multiple discriminators to evaluate generated sub-sequences. SSE-PT [43] introduces user embeddings to realize the personalization of Transformer. STRec [44] leverages sparse Transformer to focus on the most relevant interactions while predicting target items. Meanwhile, STRec [44] replaces the self-attention mechanism with a cross-attention mechanism. MoSeR [45] captures motifs hidden in behavior sequences to model micro-structure features. PHGR [46] leverages a novel hyperbolic inner product operator to enable global (all user-item interactions) and local (each user's interactions) graph representation learning in Poincaré ball. CBiT [47] leverages contrastive learning and bi-Transformer to learn more fine-grained preferences. DCRec [48] leverages cross-view contrastive learning to capture correlations among different users' interaction sequences. DFAR [49] leverages a dual-interest disentangling layer to disentangle positive and negative interests for learning users' transition patterns better. DEPS [50] leverages both users' and items' interaction sequences as the inputs to two Transformers for learning preferences better. MrTransformer [51] leverages a preference separation module to learn both users' common and unique preferences. CRR-Transformer [52] applies a pre-trained transformer model to an online RL algorithm. Due to the uncertainty of interactions, DT4SR [53] leverages an elliptical Gaussian distribution to describe items and adopts Wasserstein distance to measure the similarity among distributions. Transformers4Rec [54] applies Transformer-based architectures for recommendation, such as GPT-2, XLNET, and Transformer-XL. PAUP [55] captures both short- and long-term patterns through a progressive attention distribution mechanism. In this model, long behavior sequences are segmented into a series of sub-sequences using a down-sampling convolution module. SAMRec [56] enhances the recommendation performance of Transformer from the perspective of loss geometry.

Though Transformer-based models have achieved significant progress, these models still have some shortcomings. For example,

Transformer-based SR models suffer from a quadratic computational complexity.

4.2.4 Diffusion model

Diffusion model has been extensively applied to various generative tasks, such as image generation.

In recent years, several SR models based on diffusion model have been proposed. DiffuASR [57] leverages a diffusion-based framework to generate high-quality interactions guaranteeing the similarity between generated and original sequences. DiffuRec [58] is the first work that applies the diffusion model to sequential recommendation. By corrupting and reconstructing target item embeddings, DiffuRec [58] injects uncertainty during the recommendation process. Diff4Rec [59] leverages the diffusion model to augment data by adopting a curriculum scheduling strategy, including interaction augmentation and objective augmentation. However, Diffusion-based SR models have some weaknesses. Compared with Transformer-based SR models, diffusion-based SR models require more computational resources and training time.

4.2.5 Graph neural networks

GNNs-based SR models can learn complex high-order interactions and capture collaborative information. Therefore, GNNs-based SR models have received widespread attention. AutoSeqRec [60] leverages an encoder and three decoders to reconstruct user-item interaction and item transition matrices. CD-GES-SASRec [61] leverages a causal graph to distinguish causal and non-causal transitions. DGSR [62] constructs a user-item interaction graph to capture personal preferences and an item-item transition graph to model transitional relationships between adjacent items. Meanwhile, DGSR [62] injects high-order connections into graphs. MSSG [63] transforms items within a sequence into a star graph. Meanwhile, MSSG introduces an additional internal node to capture global information. MSSG overcomes the over-smoothing issue and realizes a linear time complexity. However, GNNs-based SR models might not fully capture sequential patterns and could be computationally expensive to train.

4.2.6 Multi-layer perceptions

Compared with Transformer, leveraging MLPs to calculate is more efficient (i.e., a linear time complexity). In recent years, some MLPs-based SR models have been proposed. FMLP-Rec [64] incorporates learnable filters into MLPs architecture for reducing noise. TRIMLP [65] disables triangle neurons in MLPs to account for chronological order in users' interaction sequences. MLM4Rec [66] learns users' global and local preferences by incorporating convolution operators into MLPs. However, these MLPs-based SR models struggle to capture complex sequential patterns.

4.2.7 Other deep learning models

Variational autoencoder (VAE) can learn compressed features from input embeddings. SVAE [67] combines both VAE and RNNs to capture sequential dependencies, with VAE extracting features at each time step.

Temporal convolutional network (TCN) is a specialized type of

convolutional neural network. Compared with RNNs, TCN is more computationally efficient. HierTCN [68] combines both TCNs and RNNs to learn short-term interests and long-term interests in sequences.

Capsule networks can recognize objects from different perspectives. GAC [69] leverages capsule networks to model personalized item-level and factor-level sequential dependencies, and combines the two kinds of sequential dependencies to recommend items to users. TAMIC [70] utilizes a time-aware dynamic routing algorithm and two kinds of time-aware voting gates to inject temporal information into sequential modeling.

4.3 Reinforcement learning

Reinforcement learning can efficiently model interactions between users and a recommender system. LIRD [71] regards sequential interactions between the users and the recommender system as a Markov decision process (MDP), and leverages reinforcement learning to learn optimal strategies for recommending items automatically. SAR [72] leverages reinforcement learning to learn dependencies and the length of the interaction sequence for each user.

4.4 Hybrid models

Some SR models demonstrate significant improvements in recommendation performance by combining different techniques. For example, combining Transformer and GNNs can achieve better performance. As shown in Table 2, each technique offers distinct advantages. Therefore, the hybrid models can combine the strengths of different techniques. The hybrid models can be roughly categorized into three classes (i.e., CNNs + Transformer, GNNs + Transformer, and other combined types). Integrating various neural network architectures enables SR models to achieve enhanced recommendation performance. Four kinds of commonly used combined models are depicted in Fig. 6. Two different techniques can be leveraged in parallel. For example, CNNs and Transformer can extract short-term dependencies and long-term dependencies from user interaction sequences, respectively. And then the two kinds of dependencies are fused together. Alternatively, the two techniques can be used in a cascaded manner. For example, GNNs can be denoted as a feature extractor for extracting features, and then the extracted features can be fed into Transformer.

4.4.1 Combination of CNNs and transformer

Transformer and CNNs exhibit unique abilities to capture different types of dependencies. Transformer has a good ability to capture long-term dependencies. CNNs are effective at modeling short-term dependencies and multiple-faced preferences. As shown in Fig. 6(a), CNNs and Transformer can be combined in a parallel manner, which is formalized as follows:

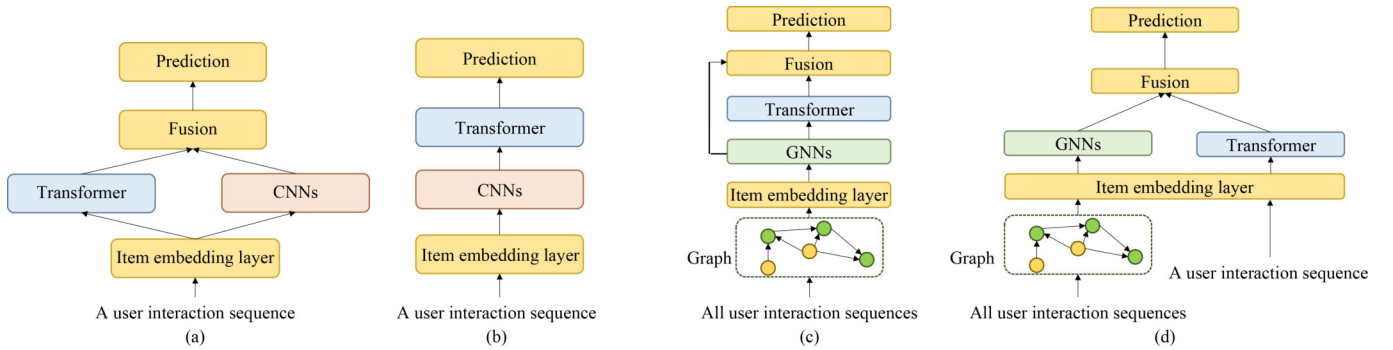
$$\mathbf{P}_u^{\text{CNNs}} = \text{CNNs}(\text{Embed}(S_u)),$$

$$\mathbf{P}_u^{\text{Tr}} = \text{Transformer}(\text{Embed}(S_u)),$$

$$\mathbf{P}_u^{\text{final}} = \text{Fuse}(\mathbf{P}_u^{\text{CNNs}}, \mathbf{P}_u^{\text{Tr}}).$$

Table 2 The advantages and disadvantages of different SR models

SR model	Advantage	Disadvantage
Traditional	Simple, can capture short-term sequential dependencies, efficient, good explainability	cannot capture long-term sequential dependencies, cannot handle complex data
RNNs-based	Can capture long-term sequential dependencies	Are sensitive to the order of items in an interaction sequence
CNNs-based	Can model union-level sequential patterns, are not sensitive to the order of interacted items	High computational costs due to lots of convolutional filters
Transformer-based	Can identify the importance of interacted items, can capture short-term and long-term dependencies	Time complexity is related to the length of interaction sequences
Diffusion-based	Can alleviate data sparsity, can generate diversity interaction data	Complex training process, high computational costs
GNNs-based	Can model complex transitions among interacted items, can capture structure characteristics and collaborative signals from a global view	High computational costs, over-smoothing problem
MLPs-based	Computational efficiency, a linear time and space complexity, simple model architecture	Limited performance on ultra-long sequences, are weak in capturing sequential dependencies
VAE-based	Can capture complex latent relationships, can handle the uncertainty of users' preferences	High computational costs, are influenced easily by potential noise
TCNs-based	Can preserve sequence order during convolution process and conduct efficient sequence modeling	Are difficult to capture long-term sequential dependencies, high complexity for hyper-parameter searching
Capsule-networks-based	Can model multi-facet interests, can model complex interaction sequences	Are difficult to capture long-term sequential dependencies
RL-based	Can adapt to changing sequence lengths, can model interaction process	cannot solve the cold-start problem, rely heavily on real-time feedback


Fig. 6 Illustration of the common combination strategies in sequential recommendation. (a) CNNs and Transformer in parallel; (b) CNNs and Transformer in cascade; (c) GNNs and Transformer in cascade; (d) GNNs and Transformer in parallel

The final preferences $\mathbf{P}_u^{\text{final}}$ can be fed into a prediction module for recommending personalized items for user u . Meanwhile, as shown in Fig. 6(b), CNNs and Transformer can be integrated in a cascaded manner, which is formalized as follows:

$$\mathbf{P}_u^{\text{final}} = \text{Transformer}(\text{CNNs}(\text{Embed}(S_u))).$$

AdaMCT [73] captures long-term and short-term dependencies by incorporating locality inductive bias into the Transformer with a local convolutional filter. CASFE [74] first leverages CNNs to capture periodic information, and then leverages Transformer to evaluate the importance of different behaviors for target behaviors.

4.4.2 Combination of GNNs and transformer

GNNs and Transformer excel at modeling graph structure and sequences, respectively. Transformer can capture temporal patterns and long-term dependencies. Meanwhile, GNNs can capture

structural information and users' long-term preferences. Therefore, combining GNNs and Transformer enables SR models to achieve better performance. As shown in Fig. 6(d), GNNs and Transformer can be combined in a parallel manner, which is formalized as follows:

$$\mathbf{P}_u^{\text{GNNs}} = \text{GNNs}(\text{Embed}(\mathcal{G}), u),$$

$$\mathbf{P}_u^{\text{Tr}} = \text{Transformer}(\text{Embed}(S_u)),$$

$$\mathbf{P}_u^{\text{final}} = \text{Fuse}(\mathbf{P}_u^{\text{GNNs}}, \mathbf{P}_u^{\text{Tr}}).$$

In the meantime, as shown in Fig. 6(c), GNNs and Transformer can be integrated in a cascade manner, which is formalized as follows:

$$\mathbf{P}_u^{\text{GNNs}} = \text{GNNs}(\text{Embed}(\mathcal{G}), u),$$

$$\mathbf{P}_u^{\text{Tr}} = \text{Transformer}(\text{GNNs}(\text{Embed}(\mathcal{G}), S_u)),$$

$$\mathbf{P}_u^{\text{final}} = \text{Fuse}(\mathbf{P}_u^{\text{GNNs}}, \mathbf{P}_u^{\text{Tr}}).$$

Compared with the former manner, it is noted that the latter leverages GNNs to extract item embeddings as the input of Transformer.

metaCSR [75] leverages Meta Learner to extract and propagate transferable knowledge from prior users and effectively learn presentations for cold-start users. GL-GraphFormers [76] leverages a global user-item bipartite graph to learn both first- and second-order graph information, and then injects the information into Transformer in the form of input and attention weights. MAERec [77] adaptively and dynamically distills sequential transition patterns for constructing a contrastive learning framework. SCG-SPRe [78] leverages a complement graph and a substitute graph to encode the complementary and substitutable relations, respectively. To explore global item-to-item interactions (collaborative information), APGL4SR [79] constructs a global item graph. Meanwhile, to capture personalized patterns, user embeddings are leveraged. MRGSRec [80] leverages a sequential encoder and a graph encoder to learn local behavioral representations and global behavioral representations, respectively. Subsequently, the two preferences are fused together for final recommendation. However, these methods have some drawbacks. For example, learning collaborative signals through graph neural networks is computationally expensive.

4.4.3 Other combinations

There are also some SR models combining various other techniques. RANN [81] leverages the strengths of RNNs and the self-attention mechanism in Transformer to capture user preferences from different views. RCNN [82] combines the advantages of RNNs and CNNs, in which RNNs can capture long-term dependencies and CNNs can capture short-term sequential patterns. NISRec [83] captures and aggregates the long-term and short-term intentions of neighbor users. Then NISRec [83] fuses these intentions for recommendation. Mixture networks can leverage multiple experts to learn diverse interests. For example, M3 [84] leverages a mixture of experts, where each expert specializes in capturing a specific temporal range. These experts are dynamically combined through a learned gating mechanism. Therefore, M3 [84] can capture diverse behavior patterns effectively.

By combining different techniques, SR models can achieve better performance. However, these models have several drawbacks. Firstly, they are more complex and computationally expensive. Secondly, different techniques might capture conflicting preferences. For example, Transformer and GNNs capture distinct preferences, which may conflict. Thirdly, different techniques might converge at different rates.

■ 5 SR with side information

In this section, we first discuss SR models combining IDs and one of the general features, categorical features and numerical features. Then, we show SR models combining IDs and one of the graph-structure features, knowledge graph and social networks. Next, we introduce SR models combining IDs and reviews, followed by SR models combining IDs and multiple feature types.

5.1 One categorical or numerical feature

Many SR models combine either a categorical or a numerical feature with item IDs to recommend items for users. Some specific methods are shown in Fig. 7. From Fig. 7, we can know that these methods can be grouped into four main types. As shown in Fig. 7(a), the first method leverages two backbones to learn two kinds of preferences from one sequence containing item IDs and the other sequence containing feature k , respectively. These preferences are then fused using methods such as contrastive learning, concatenation, addition, gate addition, or a cross-attention mechanism. The fusion process is formalized as follows:

$$\mathbf{P}_u^{\text{final}} = \text{Fuse}(\text{BM}_1(\mathbf{E}_u), \text{BM}_2(\mathbf{A}_u^{(k)})).$$

As shown in Fig. 7(b), the second method leverages the value of feature k to split the original sequence containing item IDs into sub-sequences. For example, some SR models split hybrid interaction sequences into separate interaction sequences in each domain [85,86]. Different preferences are then extracted from these sub-sequences and fused together to predict personalized items. The fusion process is formalized below:

$$\mathbf{P}_u^{\text{final}} = \text{Fuse}(\text{BM}_1(\text{Embed}(S_{u,1}^{(k)})), \dots,$$

$$\text{BM}_m(\text{Embed}(S_{u,m}^{(k)}))).$$

It is noted that the fusion usually exhibits a certain directionality, where user preferences are transferred from a source domain to a target domain [87]. As shown in Fig. 7(c), the third method first leverages simple techniques to combine item embeddings and feature embeddings. Then, the fused embeddings are fed to a backbone model for recommendation. For example, adding them together is to obtain the input of a backbone model. The fusion process is formalized as follows:

$$\mathbf{P}_u^{\text{final}} = \text{BM}_1(\text{Fuse}(\mathbf{E}_u, \mathbf{A}_u^{(k)})).$$

As shown in Fig. 7(d), the last method integrates item embeddings with feature embeddings in a backbone model. For example, leveraging feature embeddings changes the weights in attention matrices. The fusion process is formalized as follows:

$$\mathbf{P}_u^{\text{final}} = \text{BM}_1(\mathbf{E}_u, \mathbf{A}_u^{(k)}).$$

Though all four kinds of methods can obtain final user u 's preferences, they vary in terms of complexity and efficiency [87,88]. Meanwhile, as shown in Fig. 7(e), the predicted goal might contain the feature of the next interacted item [25,89].

5.1.1 Behaviors

Different behaviors show users' different preferences. Click behaviors usually reflect users' short-term preferences. Favorite behaviors usually reflect users' middle-term preferences. Purchase behaviors usually reflect users' long-term preferences. Meanwhile, dependencies exist among these different behaviors. MB-STR [90] injects users' behaviors into Transformer to capture multi-behavior heterogeneous dependencies. PBAT [91] models users' personalized behavior patterns and behavioral collaborations by a personalized behavior pattern generator and a behavior-aware collaboration

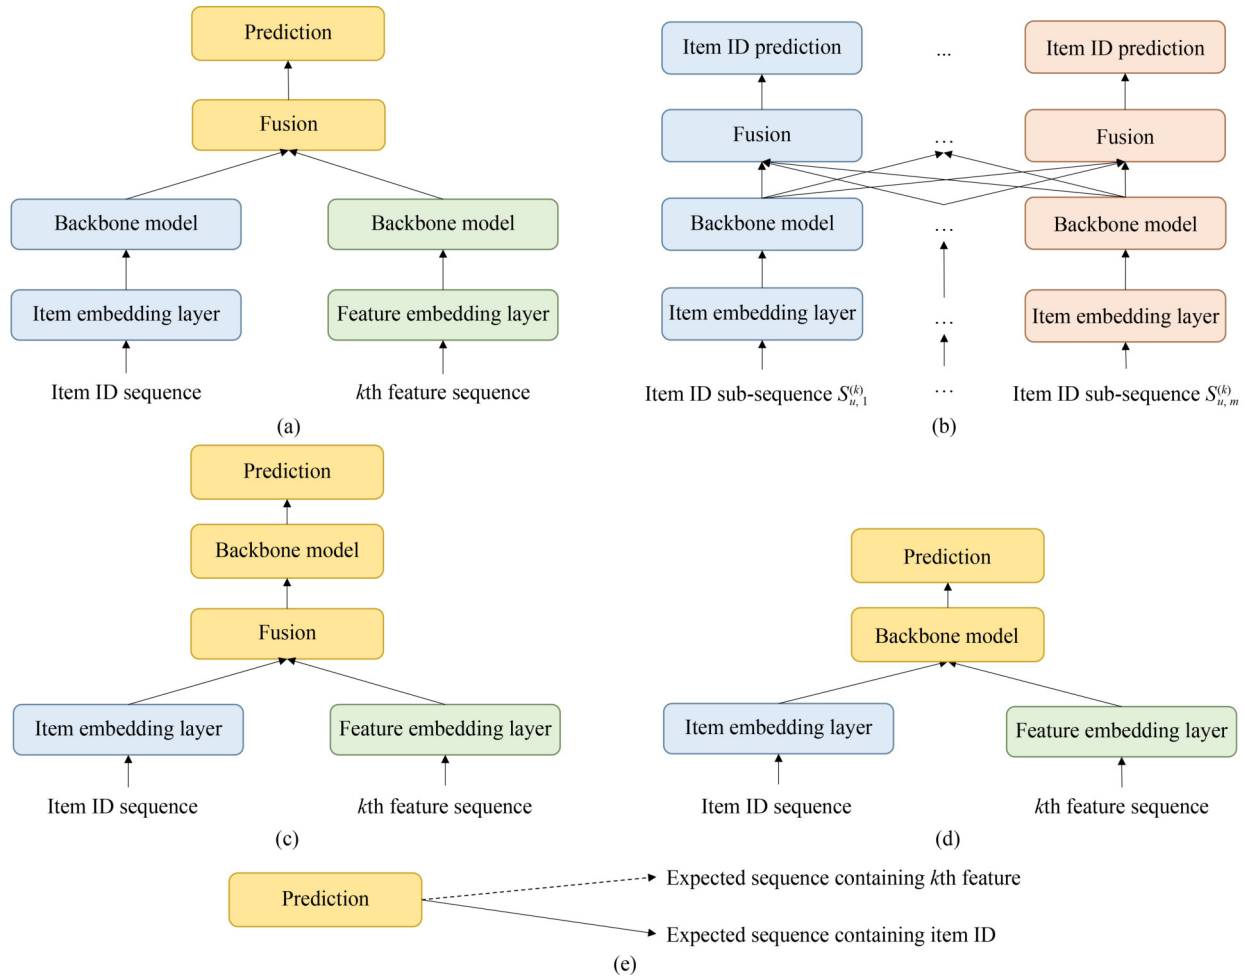


Fig. 7 Illustration of the usage of one categorical or one numerical feature in sequential recommendation. (a) Late fusion; (b) Sequence splitting; (c) Early fusion; (b) Hybrid fusion; (e) Prediction

extractor, respectively. GHTID [92] leverages a global heterogeneous graph (constructed from all interaction sequences) and some local heterogeneous graphs (constructed from each interaction sequence) to explicitly learn heterogeneous item transitions.

5.1.2 Timestamp

Users usually exhibit certain behaviors at specific times. For example, users tend to purchase T-shirts in summer. By considering timestamps, SR models can capture users' evolving interests more precisely. DRL-SRe [93] exploits time-sliced graph neural networks to learn complex user-item interactions from a global perspective, and leverages GRUs to learn complex item-to-item correlations. To account for temporal context in recommendations, MOJITO [94] leverages a Gaussian mixture method to combine attention-based temporal context with item representations. TGSRec [95] leverages a temporal collaborative Transformer layer to capture collaborative signals and temporal dynamics. TIGSA [96] leverages a time interval-aware graph to consider the impact of time intervals for recommendation, and leverages Transformer to extract users' sequential preferences. TMGN [97] first identifies the k most similar users to each candidate user. And then, TMGN [97] injects multiple users' sequences and temporal information into multi-head attention.

By leveraging timestamps, HyperRec [98] splits a whole user-item interaction hypergraph into sub-hypergraphs. HyperRec [98] adopts these sub-hypergraphs and multiple convolutional layers to capture multi-order connections and short-term user intents. Then the dynamic item embeddings and short-term user intents are incorporated by a fusion layer, and the incorporated representations are fed into a self-attention layer. PTGCN [99] leverages a position-enhanced and time-aware graph convolution operation to capture sequential patterns and temporal dynamics.

5.1.3 Category

The category of interacted items can reflect users' preferences more precisely. HierTrans [100] leverages a novel hierarchical temporal graph to extend traditional item-level relations to category-level relations. HPM [101] leverages Transformer to learn both low-level preferences (i.e., item IDs) and high-level preferences (i.e., item categories). While considering both item IDs and the categories of items, AutoMLP [102] captures users' long-term and short-term interests, with short-term interests learned by an adaptive search algorithm.

5.1.4 Domain

Domain information plays a crucial role in sequential recommender

systems. SR models usually first handle separate interaction sequences in each domain. Then SR models transfer the obtained knowledge from a source domain to a target domain [86,103]. It is because data distribution is different across domains. C²DSR [87] constructs a global item-to-item graph across two domains to learn transitions between two domains. DDGHM [104] constructs local graphs and a global graph to capture intra-domain sequential transitions and inter-domain sequential transitions, respectively. And, DDGHM leverages hybrid metric learning to alleviate data sparsity and improve recommendation performance.

5.1.5 Rating

Users usually rate interacted items, indicating their preferences. By considering ratings, users' preferences can be more accurately captured. TRA-SR [105] integrates rating information into the weight calculation of self-attention to enhance the learning of users' preferences.

5.1.6 Price

In real life, users often prefer purchasing cheaper items. Even if users have a strong liking for an item, they may not purchase it if it is too expensive. Both price and interest preferences might influence users' purchase choices simultaneously. Therefore, CoHHN [106] designs a co-guided heterogeneous hypergraph network to simultaneously extract users' price preferences and interest preferences. Then, CoHHN leverages a co-guided learning schema to model the complex relations between the two kinds of preferences. Finally, CoHHN predicts users' behaviors based on the two kinds of preferences and item features.

5.2 Graph structure features

5.2.1 Knowledge graph

Knowledge graph is a structure representation of knowledge, where entities (nodes) are connected by relationships (edges). The entities can represent users, items, and other related features. The relationships represent meaningful connections (e.g., "a user purchases a book" and "The color of the iPhone is black"). Therefore, knowledge graph can describe the relations among items, users, and features, as well as interactions between users and items. Knowledge graph provides rich extra knowledge for recommending personalized items. Therefore, by leveraging knowledge graph, SR models can alleviate cold-start and data sparsity problems. Meanwhile, SR models based on knowledge graph are also capable of providing explanations for recommendation.

There are several SR models based on knowledge graph. EIUM [107] leverages knowledge graph to conduct explainable recommendation. KV-MN [108] integrates knowledge graph into RNNs to enhance semantic representations. KP4SR [109] leverages knowledge graph to capture users' fine-grained preferences. Meanwhile, KP4SR [109] converts knowledge graph to knowledge instructions for noise reduction. KERL [110] incorporates knowledge graph into reinforcement learning. In order to learn sequential dependencies while considering semantic information, KERL [110] leverages a composite reward function to compute both sequence-level and knowledge-level rewards. TKG-SRec [111] leverages

temporal knowledge graph to capture temporal information, and leverages GRUs to conduct temporal knowledge evolution training. However, SR models based on knowledge graph have disadvantages as follows: 1) Complexity: Knowledge graph contains many entities and relationships, which may increase the complexity of SR models and cost more computational resources. 2) Noise: Constructing a high-quality knowledge graph is challenging. Incorrect relationships within the knowledge graph might degrade the recommendation performance.

5.2.2 Social network

Social networks provide insight into the friend relationships among users. Since users tend to have similar preferences to their friends, leveraging social networks allows SR models to capture users' preferences more accurately. Satori [112] leverages a graph attention network and a self-attention mechanism to extract auxiliary features and user intentions, respectively. Then, hybrid user and item representations are sent to a prediction module to calculate predicted scores. SERec [113] first constructs a heterogeneous graph by combining a social network and users' interaction sequences. Then by leveraging GNNs, SERec improves user and item representations by integrating knowledge from the social network.

5.3 Review

Compared with rating information, reviews can reflect users' preferences about items comprehensively in the form of text and images. Meanwhile, reviews can also provide more detailed information about items. By reading reviews given by other people, users' action decisions might be influenced. RNS [114] learns users' long-term preferences by an aspect-aware convolutional network over users' document. Meanwhile, RNS learns users' short-term preferences by hierarchical attention. Finally, RNS combines the two kinds of preferences for recommendation. To consider both user-item interactions and reviews simultaneously, PIRSP [115] learns item sequential patterns and review sequential patterns together. Then PIRSP introduces a fusion gating mechanism to learn short-term preferences from the two patterns. Finally, PIRSP combines short-term and long-term preferences for recommendation. CCA [116] designs a cascaded cross-attention mechanism to capture the complex relationships among an item sequence, a review sequence, and candidate items.

5.4 Fusion of different features

In SR models, multiple features are usually used to improve the recommendation performance. Different features can be fused by different fusion methods. These fusion methods can mainly be summarized into three types. Some specific details are shown in Fig. 8. The first type leverages simple methods to fuse item embeddings and feature embeddings together, for example, adding them together is to obtain the input of a backbone model. The fusion is shown as follows:

$$\mathbf{P}_u^{final} = \text{BM}_1(\text{Fuse}(\mathbf{E}_u, \mathbf{A}_u^{(1)}, \dots, \mathbf{A}_u^{(K_g)})).$$

As shown in part (b) of Fig. 8, the second type fuses different features in a backbone model [25], for example, leveraging a

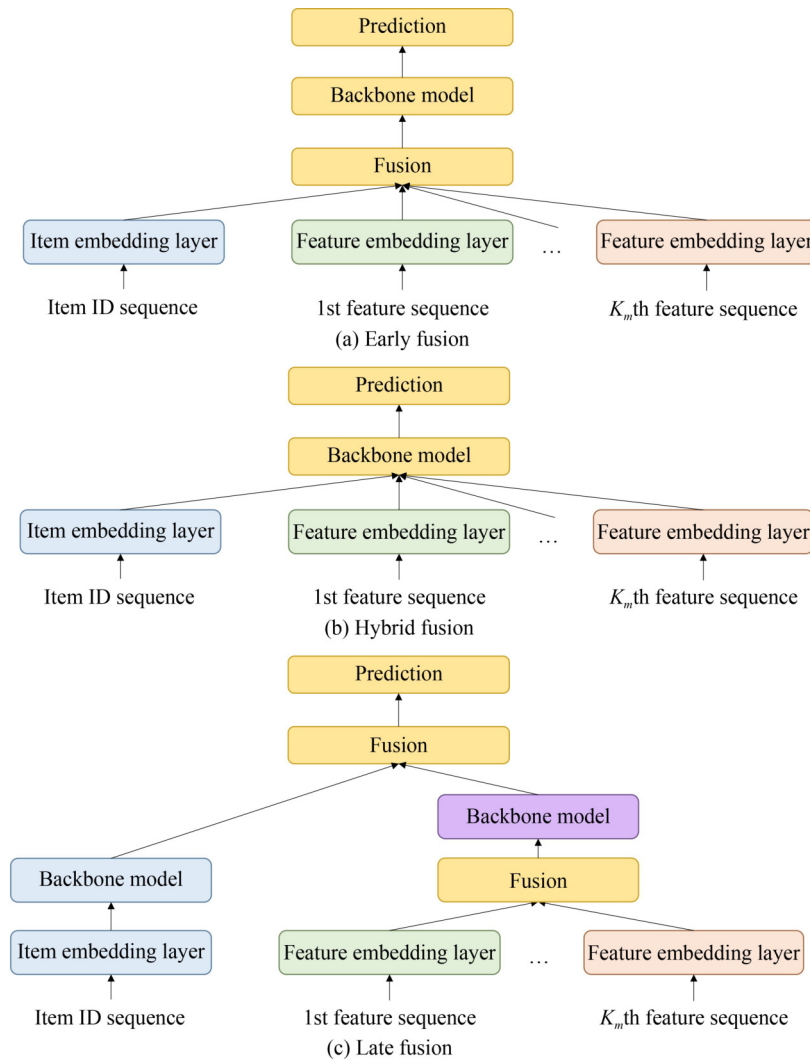


Fig. 8 Illustration of different fusion strategies for multiple general features

decoupled features fusion mechanism in a self-attention mechanism. The fusion is shown as follows:

$$\mathbf{P}_u^{final} = \text{BM}_1(\mathbf{E}_u, \mathbf{A}_u^{(1)}, \dots, \mathbf{A}_u^{(K_g)}).$$

As shown in part (c) of Fig. 8, the third type first fuses different feature embeddings together, and then passes the fused embeddings through a backbone model to learn feature representations. Meanwhile, item representations are learned by the other backbone model. Finally, the item representations and feature representations are fused together for prediction. The fusion is described as follows:

$$\mathbf{P}_u^{\text{feat}} = \text{BM}_1(\text{Fuse}(\mathbf{A}_u^{(1)}, \dots, \mathbf{A}_u^{(K_g)})),$$

$$\mathbf{P}_u^{\text{final}} = \text{Fuse}(\text{BM}_2(\mathbf{E}_u), \mathbf{P}_u^{\text{feat}}).$$

We will introduce common features fusion, context information fusion and special features fusion one by one.

5.4.1 Common features fusion

In this section, we mainly introduce some SR models that fuse some common features (e.g., position, category, brand) and item IDs. The structure of these SR models is feature-agnostic to some extent.

ICAI-SR [3] leverages a heterogeneous graph to capture complex item-to-attribute relationships by inner attribute aggregation and attribute-to-item aggregation, and then passes aggregation embeddings through an entity sequential model. MLP4Rec [117] designs a fusion mechanism to capture sequential, cross-channel and cross-feature correlations. SMLP4Rec [118] leverages layer normalization for a sequence mixing module, a feature mixing module and a channel mixing module simultaneously. Meanwhile, to enhance efficiency and effectiveness, SMLP4Rec leverages a parallel mode. RetaGNN [119] leverages a relational attentive GNN in which learnable weight matrices focus on various relations among users, items and attributes to have inductive and transferable properties. Then, RetaGNN [119] leverages Transformer to capture short-term and long-term temporal patterns. S³-Rec [120] leverages four self-supervised learning tasks, i.e., associated attribute prediction, masked item prediction, masked attribute prediction, and segment prediction, to learn the associations between item IDs and item features.

In Transformer-based SR models, several advancements have been made in combining common features with item IDs. FDSA [121]

first fuses different features by a vanilla attention mechanism. Then, FDSA leverages separate self-attention blocks to independently learn feature-level and item-level representations. Finally, FDSA fuses the two kinds of preferences for recommendation. SASRec_F [25] is a sample model to fuse common feature embeddings and item ID embeddings through a concatenation operation. This fusion method is considered invasive [4], similar to an addition operation and a gating addition operation. The specific process is shown in Fig. 8(a). This fusion method might compromise the original information in item embeddings, potentially leading to negative effects. To address this issue, NOVA [4] leverages a non-invasive method to enhance the learning of attention matrices. In NOVA, key and query matrices are derived by integrating different feature embeddings and item embeddings while value matrices are directly obtained from item embeddings. However, NOVA [4] still leverages integrated embeddings to learn key and value matrices. The compounded embedding space might introduce random disturbances and share identical gradients, which might degrade recommendation performance. DIF-SR [25] leverages decoupled attention to adaptively learn different features by flexible gradients. MSSR [122] captures the associations among an item ID sequence and each feature sequence. ASIF [123] maintains semantic consistency between item IDs and features. Meanwhile, ASIF [123] makes full use of different kinds of features without compromising item embeddings, thanks to contrastive learning.

5.4.2 Context information fusion

Context information can reflect the specific environment where behaviors happen. SR models can provide more accurate recommendations by considering the context information. Bi-CASER [124] uses a context graph to represent context information, and then leverages bi-Transformer to learn users' sequential behaviors. CRNNs [125] incorporates context information (e.g., interaction timestamps) by combining item embeddings with context embeddings to recommend personalized items. STAR [126] leverages stacked RNNs to model the sequential and context-aware information simultaneously. STEM [127] considers item behaviors and leverages a Transformer-based architecture for exploiting spatial and temporal information.

5.4.3 Special features fusion

In this section, we mainly introduce SR models designed to fuse special features. These SR models are typically applied to situations where certain special features (e.g., social networks and timestamps) are available. STEN [26] introduces temporal information and a social network to capture users' fine-grained dynamic interests in an event-level direct paradigm. DUVRec [128] constructs a timespan-aware sequence graph and an attribute-augmented graph to learn item-view and factor-view user representations, respectively. Recommendation performance is then improved by fusing the two representations. Based on features and labels, STARec [129] retrieves a target user's historical behaviors through a search-based retriever. Then, STARec [129] leverages a graph structure to enhance representation learning. SGCross [130] transfers knowledge from an auxiliary domain to a target domain from a personal view, a temporal view and a collaborative view. The knowledge encompasses personal preferences, dynamic preferences, and collaborative preferences.

6 Recent SR advancements

In this section, we introduce recent research directions in SR models. Firstly, we discuss SR models using multi-modal features. Secondly, we explore a novel paradigm, generative recommendation, in sequential recommendation. Thirdly, we introduce SR models using LLMs. Fourthly, we introduce SR models designed to address ultra-long interaction sequences. Finally, we present augmentation techniques aimed at enhancing SR performance.

6.1 Multi-modal SR

A user u can be represented by the user's profile. The profile may consist of user u 's demographic information or user u 's interacted items. An item can be denoted by an item ID embedding or a pre-trained modality embedding. Modality content features consist of text, images, videos, audio, text-image multi-modal pairs, and so on. ID-agnostic representations can be learned using pre-trained modality models. A common text encoder is BERT, which is usually used to extract semantic features from text. Swin Transformer is an image encoder, which can be leveraged to extract features from images. As shown in Fig. 9, unlike ID-based SR models, one type of modal-based SR models leverages a single modality feature (e.g., text or

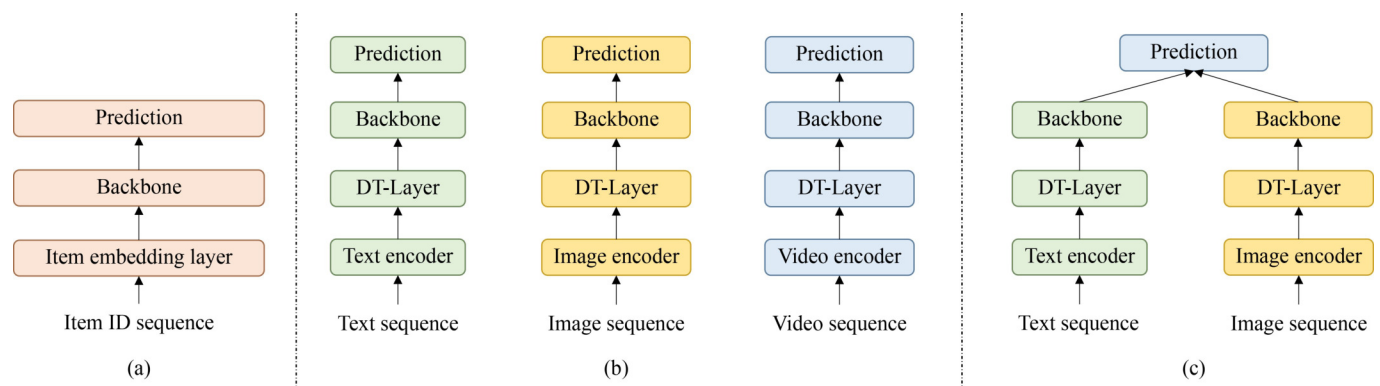


Fig. 9 Illustration of different modal-based SR models. Note that DT denotes the dimension Transformer. (a) ID-based SR models; (b) single-modality-based SR models; (c) multi-modality-based SR models

images) to denote items. The other type of modal-based SR models leverages some modal features (e.g., text and images) to represent items [131]. In other words, the representations of item i are obtained from the multi-modal features of item i in modal-based SR models, which is formalized as follows:

$$\mathbf{e}_i = \begin{cases} \text{TextEncoder}(\text{Text}_i), \\ \text{ImageEncoder}(\text{Image}_i), \\ \text{VideoEncoder}(\text{Video}_i). \end{cases}$$

SR problems can be solved by NLP models. Different from natural language, interacted items do not typically follow a specific order in user interactions. If we swap two items in a user's interaction sequence, the user's preferences might not be changed. If we swap two words in a sentence, the meaning of this sentence will be changed. Text contains rich semantic information in the real world. By using natural language like text, semantic gaps across different domains and platforms can be bridged. Therefore, SR models can be trained in a universal semantic space.

There are some single-modality-based SR models. UniSRec [8] leverages an MoE-enhanced adaptor and parametric whitening for domain adaptation. To learn universal sequence representations better, UniSRec proposes a seq-item contrastive task and a seq-to-seq contrastive task. The model is pre-trained in one domain and subsequently fine-tuned in another domain using either an inductive or a transductive approach. Meanwhile, UniSRec adopts two augmentation strategies (i.e., item drop and word drop). VQ-Rec [27] is a novel approach to learn vector-quantized item representations for solving the problem about over-emphasizing the effect of text features, which might lead to poor recommendation performance. VQ-Rec converts item text to a series of codes. Then with a lookup over a code embedding table, VQ-Rec can obtain item embeddings. Based on representation schema, VQ-Rec leverages contrastive learning among the next ground-truth items and negative items (i.e., semi-synthetic negative items and mixed-domain negative items). Finally, based on the code-embedding method, VQ-Rec implements a novel cross-domain fine-tuning method. TransRec [132] is pre-trained in one domain, and then is fine-tuned in another domain. TransRec [132] achieves excellent recommendation performance. It is because that TransRec [132] first obtains knowledge from a source domain, and then transfers the obtained knowledge to a target domain. The three models mentioned above share a similar approach: they are initially pre-trained in one domain and then fine-tuned in another domain. IDA-SR [133] adopts a pre-trained text encoder to learn item representations by encoding item text descriptions. The pre-training tasks consist of next-item prediction, masked item prediction and permuted item prediction. These item representations can then be used for personalized recommendations through semantic transfer. SASRec_v [134] leverages videos to denote an item. Compared with SASRec using an item ID to denote an item, SASRec_v [134] can achieve comparable performance using an end-to-end framework. RECFORMER [135] mainly focuses on embedding initialization. By adopting behavior-tuned pre-trained language models for the initialization of ID-based SR models, RECFORMER [135] can achieve better recommendation

performance without introducing additional computational costs. SRT-1K [136] calculates item embeddings using a trainable feature extractor, which can make the number of parameters be independent of the number of items. Additionally, SRT-1K [136] leverages contrastive learning to enhance catalog diversity representation. RECFORMER [137] flattens item attributes into sentences to represent items. The SR models mentioned above achieve improved recommendation performance by combining item IDs and multi-modal features.

There are some multi-modality-based SR models. ODMT [131] fuses different modal information (i.e., text and images) together. Meanwhile, ODMT uses knowledge distillation to model weights for online recommendations. MISSRec [138] leverages a multi-modal pre-training and transfer learning framework to integrate multi-modal information into SR models. MMSR [122] integrates multi-modal information into nodes. MMMLP [139] incorporates multi-modal data into MLPs to enhance recommendation performance. p-RNN [140] leverages four different frameworks to combine some item features (i.e., images and text) and item IDs for modeling interaction sequences. Based on the aforementioned modal-based SR models, the following conclusions can be drawn:

- Modal-based SR models with an end-to-end training framework achieve comparable recommendation performance with pure ID-based SR models. Because the pre-trained features might not be suitable for recommendation.
- Modal-based SR models with larger parameters tend to have better recommendation performance than the counterpart containing smaller parameters in recommendation tasks.
- In a cold-start setting, modal-based SR models can obtain better recommendation performance than pure ID-based SR models.
- If we leverage a pre-trained modality model to extract features, we can obtain better performance than training a model from scratch.

In the meantime, the modal-based SR models have some advantages as follows:

- Leveraging modality information to denote an item has the advantages of interpretability and visualization.
- Modal-based SR models can alleviate the cold-start problem to some extent.
- Modal-based SR models can realize the transferability of the models across different platforms.
- Multi-modal features act as a bridge to better integrate recommendation and search algorithms.

The modal-based SR models also have several disadvantages as follows:

- Modal-based SR models require substantial computational resources and extended training time.
- Hyper-parameters need to be searched carefully, as improper hyper-parameters easily lead to model collapse.

- Semantic information from different modalities might not be suited for recommendation.

6.2 Generative recommendation

Generative recommendation is proposed on the basis of a widely used technique in document retrieval (i.e., generative retrieval) [141]. Different from some above-mentioned SR models, generative recommendation models [141,142] leverage semantic IDs to denote an item. As depicted in Fig. 10, a semantic ID consists of a series of tokens. Compared with only using an item ID to denote an item, this method has several advantages as follows:

- Compared with item IDs, semantic IDs encode semantic information extracted from side information (e.g., text descriptions) more efficiently.
- Because the semantic IDs of different items might contain some of the same tokens, knowledge can be shared among similar items.
- Using semantic IDs prevents the embedding table size from increasing linearly with the number of items in a data. It is because an item is denoted by a series of tokens instead of a unique item ID. The total number of tokens is far smaller than the number of all item IDs.

Some above-mentioned methods leverage the multi-modal features of an item to denote the item. However, these methods might damage fine-grained semantic information [143]. Semantic IDs can solve this problem. Due to their hierarchical nature, semantic IDs contain semantic information that ranges from a coarse-grained level to a fine-grained level. Generative recommendation generally involves: semantic ID generation and generative recommendation with semantic IDs. As depicted in Fig. 10, semantic IDs are obtained by encoding item IDs and side information about items. RQ-VAE is a widely-used semantic ID generator, which leverages a residual quantizer to learn quantized representations [141]. Because target items are denoted by semantic IDs, a generative recommender model need to generate the target items by an auto-regressive manner. A widely-used generative recommender model is based on a Transformer-based encoder-decoder architecture [144,145]. The process is shown in Fig. 11. In generative recommender models, the probability of user u interacts with item i is calculated by chain rule, which is formalized as follows:

$$P(i|S_u) = \prod_{j=1}^l p(y_{i,j}|S_u, y_{i,1}, y_{i,2}, \dots, y_{i,j-1}).$$

In recent years, some generative sequential recommender models have been proposed. TIGER [141] is the first model that leverages




Fig. 10 Illustration of semantic ID generation

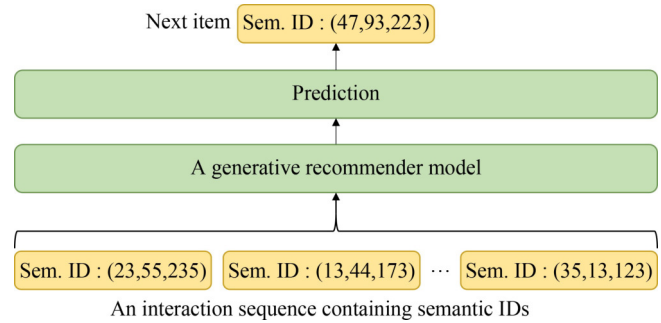


Fig. 11 Illustration of generative recommendation with semantic IDs

semantic IDs for generative recommendation. MBGen [142] considers the role of behavioral information in generative recommendation. To consider the complementary nature between behavioral information and semantic information, EAGER [144] leverages a shared encoder and two separate decoders to decode the two kinds of information. CoST [145] leverages a contrastive quantization-based semantic tokenization approach to capture neighborhood relationships among items. LETTER [143] leverages semantic, collaborative and diversity regularization. The first two types of regularization can extract semantic information and collaborative signals, respectively. The diversity regularization can mitigate code assignment bias. To consider the multi-modal features of items, MMGRec [146] leverages a hierarchical quantization method, Graph RQ-VAE, by incorporating multi-modal features. SpecGR [147] leverages a drafter model to generate some candidate items, including some existing items and new items. Meanwhile, SpecGR [147] leverages a verifier to decide whether to accept or reject candidate items, enabling SpecGR to recommend new items. Compared with traditional SR models, generative recommender models have some advantages as follows:

- Generative recommender models have strong capabilities in recommending unpopular items. Therefore, they can alleviate the cold-start problem.
- By leveraging temperature-based sampling, generative recommender models can generate diverse recommendation lists.

However, generative recommender models also have some disadvantages. For example, they might generate invalid semantic IDs during the prediction phase.

6.3 LLM-powered SR

6.3.1 Large language models

Large language models (LLMs) can handle more complex tasks compared with traditional language models based on deep learning. The main characteristics of LLMs are their billions of parameters as well as the vast amounts of data and computational resources used during training. Therefore, LLMs have a deeper understanding of natural language and extensive knowledge extracted from amounts of data. There are many kinds of LLMs. In our survey, we introduce them briefly as follows:

- LLaMA⁴⁾ (Large Language Model Meta AI), developed by Meta. It is a family of large language models, including LLaMa1, LLaMa2 and LLaMa3.
- ChatGLM (Generative Language Model), developed by Tsinghua University, which is a bilingual model for understanding and generating both Chinese and English.
- T5 (Text-To-Text Transfer Transformer), developed by Google, which can solve many kinds of tasks by modeling each task as a text-to-text task.
- ChatGPT⁵⁾, the most popular large language model, which is developed by OpenAI. It includes models such as GPT-3.5, GPT-4, GPT-4o, and other variants.
- PaLM (Pathways Language Model), developed by Google, which is one of the largest language models. It shows strong capabilities in multilingual tasks and source code generation.
- Vicuna⁶⁾, designed by fine-tuning LLaMA on high-quality conversational data.
- Qwen, developed by Alibaba and excels at processing some tasks described in both Chinese and English.
- Baichuan, contains a series of large language models developed by Baichuan Intelligent Technology, including Baichuan 2 and Baichuan-M1.
- DeepSeek⁷⁾, has recently achieved remarkably impressive performance on various metrics in comparison with the most powerful LLMs.

LLMs have a strong ability to understand text descriptions and integrate external knowledge. Consequently, incorporating LLMs into SR models can enhance performance and alleviate challenges such as cold-start and data sparsity problems. As shown in Table 3, there are some SR models leveraging LLMs for recommendation. The specific details are shown as follows.

6.3.2 LLMs for recommendation

There are several models leveraging LLMs for recommendation. Some specific details are shown in Fig. 12. The first type of models leverages LLMs to recommend items to users directly, which is formalized below:

Table 3 The SR models based on LLMs

Large language model	Representative works
T5	[148], [149], [150]
LLaMA	[151], [152], [153], [154], [155], [156], [157], [158], [159], [9], [160], [161], [162]
ChatGLM	[163]
ChatGPT	[164], [165], [166]
Vicuna	[167], [168]
Qwen	[169]
Baichuan	[170]

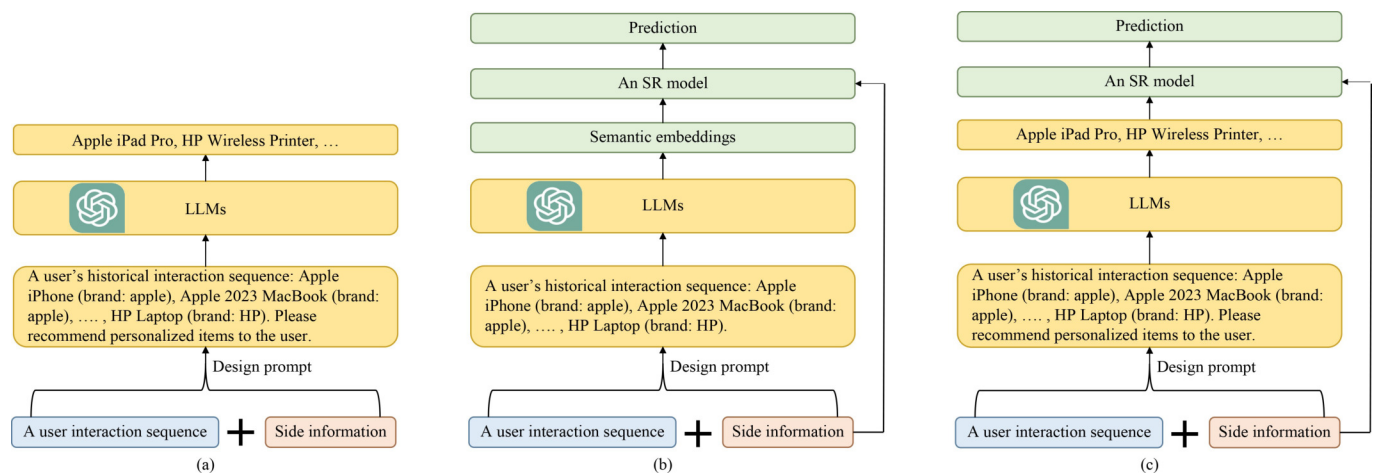


Fig. 12 The common SR models based on LLMs. (a) LLMs for recommendation; (b) LLMs for learning semantic embeddings; (c) LLMs for data generation

⁴⁾ Website at huggingface.co/meta-llama/Llama-2-13b (Last visited on May 12, 2025).
⁵⁾ Website at platform.openai.com/ (Last visited on May 12, 2025).
⁶⁾ Website at github.com/lm-sys/FastChat (Last visited on May 12, 2025).
⁷⁾ Website at github.com/deepseek-ai (Last visited on May 12, 2025).

$$\hat{x}_{\text{prompt}} = \text{LLMs}(x_{\text{prompt}}).$$

By reading the outputs of LLMs \hat{x}_{prompt} , we can identify the names of recommended items.

The second type utilizes semantic embeddings extracted from LLMs to recommend personalized items. As shown in Fig. 12(b), the prompt x_{prompt} might contain the text descriptions about a user's whole historical interaction sequence. Similar to the modal-based SR models, the prompt x_{prompt} might be the text descriptions of item i , and LLMs are leveraged to extract semantic information about item i , which is formalized as follows:

$$\mathbf{e}_i = \text{LLMs}(x_{\text{prompt}}).$$

Compared with conventional text feature extractors such as BERT, LLMs can extract richer semantic information.

As shown in Fig. 12(c), the third type of models leverages LLMs to generate high-quality data which is then combined with the original data to train SR models. The data generation process is formalized as follows:

$$\hat{x}_{\text{prompt}} = \text{LLMs}(x_{\text{prompt}}).$$

Compared with the above-mentioned first type, the outputs of LLMs are regarded as augmented data for training. It is noted that the x_{prompt} might vary with the usage of LLMs changing.

The most straightforward approach is to directly leverage LLMs for recommending personalized items. Given user u 's historical sequence and a target item, LLMs predict u 's preferences for the target item. Meanwhile, given the interacted item sequences, LLMs can generate corresponding recommendation lists. For example, P5 [148] can generate personalized recommend items by inputting historical interaction items into LLMs. There are different tasks, such as generating an item or a list of items, rating items, and ranking the order of the interacted items. POD [150] converts discrete prompt into a set of continuous prompt vectors. This method can reduce the inference time by bridging the gap between IDs and words. GenRec [9] leverages an adaptation method (LoRA) to fine-tune and perform inference tasks on LLaMA. However, these models have some weaknesses, such as the potential to generate non-existent items.

6.3.3 LLMs for learning semantic embeddings

Some SR models use LLMs to initialize item embeddings. Firstly, RecInterpreter [153] leverages a linear layer to map item embeddings from an item ID embedding space to a text embedding space. By providing task-specific textual prompts and projected hidden representations, LLaMA is encouraged to generate textual descriptions of items in interaction sequences. The second task aims to discover residual items by task-specific textual prompts and hidden representations. To solve the problem of long-tail items, LLM-ESR [164] leverages LLMs to generate semantic embeddings for enhancing the semantic information of items. Common methods are to combine item ID embeddings and generated semantic embeddings for achieving better recommendation performance. To combine the advantages of traditional SR models and LLMs, E4SRec [156] uses item IDs as input. This method ensures that all the generated and recommended items fall within candidate lists.

Therefore, E4SRec can handle item IDs more effectively. LLMEmb [162] is a novel method that leverages LLMs to generate item embeddings. To integrate collaborative signals into LLM-generated embeddings, LLMEmb leverages a recommendation adaptation training strategy. Though knowledge can be transferred from some source domains to a target domain, the cold-start problem cannot be completely solved. LLMs cannot capture collaborative information effectively. Traditional cross-domain sequential recommendation models can capture collaborative signals while LLMs-based models excel at capturing semantic information. URLLM [157] incorporates both advantages. Firstly, URLLM leverages a dual-graph neural network to capture collaborative signals and transfer knowledge from a source domain to a target domain. Meanwhile, URLLM leverages a retrieve-generation model to incorporate structure information into an LLMs-based recommender model. Finally, URLLM combines collaborative information and semantic information to recommend items. LLMCDSR [170] leverages LLMs to predict unobserved cross-domain interactions, which improves the recommendation performance for non-overlapped users. LLMCDSR [170] leverages a collaborative-textual contrastive pre-training approach to integrate collaborative information into textual features. From these models, we can draw the following conclusions: 1) Initializing item embeddings with a large language model can obtain recommendation performance gains. 2) Fine-tuning LLMs for recommendation tasks enables LLMs to learn domain-specific and domain-shared knowledge.

To reduce computational overload and enhance the efficiency and effectiveness of inference, Lite-LLM4Rec [149] leverages a beam search decoding method and a hierarchical LLM structure for recommendation. We can leverage prompts to conduct fine-tuning of LLMs. TALLRec [154] employs two tuning methods, such as Alpaca tuning and rec-tuning, to optimize the tuning process of LLMs more efficiently. SeRALM [158] leverages item IDs and knowledge generated by LLMs as input. Meanwhile, SeRALM leverages the feedback from recommender systems to refine the knowledge of items generated by LLMs. To extract semantic information generated by LLMs and enhance computational efficiency, SAID [159] leverages a projector to map item IDs into item embeddings and feeds them into LLMs for refining the item embeddings. Finally, SAID leverages the trained projector to recommend items more effectively.

Knowledge distillation and contrastive learning are effective techniques for extracting knowledge from LLMs. Knowledge distillation can enhance the efficiency of LLMs. SLIM [151] can obtain semantic knowledge from LLMs through knowledge distillation. DLLM2Rec [155] distills the knowledge from LLMs-based SR models to conventional SR models by leveraging two kinds of distillation (i.e., collaborative embedding distillation and importance-aware ranking distillation). LSVCR [163] is trained in two stages (i.e., personalized preferences alignment and recommendation-oriented fine-tuning). It uses contrastive learning between preferences generated by LLMs-based recommendation models and preferences generated by conventional SR models. Therefore, these conventional SR models can obtain richer semantic

information. SLMRec [161] leverages a layer-wise knowledge distillation approach to enhance the efficiency of SR models based on LLMs.

LLMs often struggle with processing long text descriptions. LLM-TRSR [152] can leverage too-long text descriptions as the input of LLMs by segmenting original users' interaction sequences. The LLMs-based SR models will consume lots of computational resources when processing long interacted sequences. To solve this problem, BAHE [169] leverages a behavior-aggregated hierarchical encoding method to enhance computational efficiency. By encoding a user's separate interacted sequences in parallel, BAHE can obtain the user's final representations by concatenating them together. Because LoRA's shared parameters across all users limit its ability to recommend personalized items. Meanwhile, ultra-long interaction sequences might affect the efficiency of LLMs. Therefore, RecLoRA [168] proposes a personalized LoRA module to overcome this limitation and designs a long-short modality retriever to retrieve interaction sequences of different lengths based on different modalities.

6.3.4 LLMs for data generation

LLMs can be leveraged to generate data for alleviating the data sparsity and cold-start problems. For example, KDA_{LRD} [166] leverages LLMs to describe relationships among items. Most existing models rely on predefined relations among items, which might result in sparse issues. To generate more item-to-item relations, KDA_{LRD} leverages LLMs to learn item representations. An item reconstruction module leverages the estimated relation and another item to reconstruct the target item. Finally, KDA_{LRD} leverages the predefined and generated relations for recommendation. FineRec [165] initially leverages LLMs to extract attribute pairs from reviews. Then by constructing an attribute-aware graph and aggregating, FineRec can obtain improved user and item embeddings for recommendation. For few-shot recommendation, ReLLa [167] can be used as a data augmentation method to augment data. ReLLa [167] leverages LLMs to retrieve the top- k semantically relevant behaviors according to a target item. Leveraging the semantic knowledge of LLMs, ReLLa enhances the quality of data. Finally, ReLLa is trained on a mixed data containing the original data and the retrieval-enhanced data. LLM4DSR [160] leverages LLMs to identify noisy items and replace them with suggested ones.

6.4 Ultra-long SR

An ultra-long behavior sequence typically refers to a scenario where the number of interacted items exceeds 1,000. Compared with conventional behavior sequence modeling in SR models, ultra-long behavior sequence modeling presents several challenges. Firstly, modeling ultra-long sequences requires substantial computational resources. Secondly, ultra-long sequences often contain noise due to irrelevant behaviors. Finally, during the inference phase, modeling ultra-long sequences can be computationally expensive. For example, the time complexity of Transformer is $O(L^2)$. Therefore, SR models while modeling ultra-long sequences might face high latency. The general ultra-long sequence modeling methods are shown in Fig. 13. Firstly, the most relevant k items are retrieved from a whole

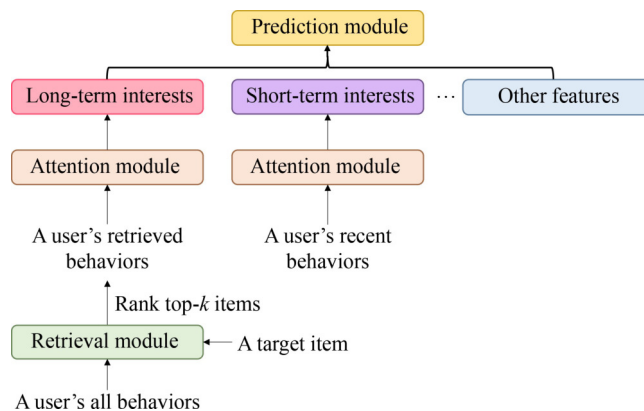


Fig. 13 Illustration of a typical ultra-long sequential recommendation method

interaction sequence. Then long-term interests are extracted from the retrieved k items by an attention module. Short-term interests are extracted from recently interacted items by an attention module. The process of extracting the two interests is formalized as follows:

$$\mathbf{P}_u^{\text{short}} = \text{Att}(S_{u,\text{recent}}),$$

$$\mathbf{P}_u^{\text{long}} = \text{Att}(\text{Ret}(S_u)).$$

The sequences fed into the attention module are relatively short, so computational efficiency can be enhanced to some extent. Finally, by leveraging the long-term interests, the short-term interests and other features, SR models can recommend personalized items for users.

Traditional SR methods focus on leveraging the most recent behaviors, while UBR4CTR [171] emphasizes the most relevant behaviors. When only the recent interactions are leveraged, the long-term dependencies and periodicity might not be captured accurately. ETA-Net [172] aims to retrieve relevant items from long sequences in a more efficient manner. Instead of using a traditional dot-product operation, ETA-Net leverages a hashing-based target attention algorithm to calculate the Hamming distance, and then selects the top- k relevant behavior items for a target item. Finally, ETA-Net leverages multi-head attention to learn long sequence embeddings for capturing diverse interests from the selected top- k items. During prediction, ETA-Net leverages five different embeddings (i.e., long sequence embeddings, target item embeddings, short-sequence embeddings, user profile embeddings, and context embeddings). For users with long interaction sequences, these sequences can be divided into sessions based on time intervals between interactions. As a result, HNUM [173] is designed with a two-layer architecture. The first layer is a session-level GRU model designed to capture users' short-term preferences within the current session. The second layer is a user-level memory network, which can capture users' long-term preferences across their entire interaction history. SIM [11] captures users' interests using both general search unit (GSU) and exact search unit (ESU). GSU leverages both hard search and soft search methods. The hard search aims to retrieve interacted items belonging to the same category of a candidate item from a user's long sequence. The soft search leverages inner product to identify top- k relevant items. These top- k items are used to construct users' sub-sequences.

By employing multi-head attention, ESU effectively models the relationship between candidate items and the sub-sequences to extract users' long-term preferences. To solve the inconsistency between CP-GSU and ESU for target items, TWIN [174] leverages the same target-behavior relevance metric in both modules. Furthermore, to enhance efficiency, TWIN compresses user-item cross features into bias terms for target attention. To tackle ultra-long user behavior sequences, TWIN-V2 [175] leverages a hierarchical method to compress life-cycle behaviors. Similar items are classified into the same cluster. Meanwhile, TWIN-V2 leverages cluster-aware target attention to learn users' diverse interests. R-RNN [176] learns a user's global interests and recent interests based on global user behaviors and recent user behaviors, respectively. To improve efficiency while modeling long behavior sequences, SAM [177] leverages dual-query attention by regarding target item embedding and memory vector as queries. User interest center (UIC) is designed to decouple user interest modeling from the whole model. Modeling user interests needs to consume lots of computational time. UIC learns a user's latest interest representations according to real-time user behavior events. MIMN [178] leverages memory utilization regularization and memory induction unit to learn multi-channel user interests. UBCS [179] proposes a behavior sampling module to sample short sequences based on candidate items, relevance and temporal information. To speed up the process of sampling, UBCS leverages an item clustering module to cluster candidate items for selecting some representative candidate items. CoFARS [180] takes into account the influence of contexts on user preferences. CoFARS leverages a prototype-based approach to identify contexts that reflect similar user preferences. Then, CoFARS constructs a temporal graph containing context nodes and prototype nodes for integrating temporal information into users' interests. Finally, CoFARS identifies the prototypes that are consistent with the target context to generate a sub-sequence. Different from single-domain long sequence modeling, LCN [181] can transfer knowledge from a source domain to a target domain. In a cross representation production module, LCN leverages contrastive learning to bridge the connections among similar items in different domains. In a lifelong attention pyramid module, LCN leverages three kinds of cascading attentions to extract preferences from user lifelong behavior sequences considering candidate items. LRURec [182] leverages linear recurrent units and a recursive parallelization framework to enhance computational efficiency and realize low-cost inference. LinRec [183] reduces the quadratic time complexity of self-attention to a linear time complexity while preserving the advantages of attention mechanisms. To address the issue of sparse search behaviors failing to capture users' interests, QIN [184] adopts a two-stage search method. Firstly, the top- k items from a long sequence for user u are retrieved using a pre-trained retrieval model. Secondly, QIN predicts whether user u will purchase items based on the retrieved top- k items. For QIN, it first searches sub-sequences based on the relevance between behaviors and queries. Secondly, QIN searches sub-sequences according to the relevance between behaviors and target items. By considering item ID field and attribute field separately, QIN leverages a fused attention unit to learn users'

representations better. To efficiently retrieve sub-sequences, SDIM [185] leverages an effective sampling method to generate sub-sequences. SDIM generates hash signatures for both the candidate item and each item in a user behavior sequence by multiple hash functions. SDIM can obtain the sub-sequences where items have the hash signatures the same as the candidate item. Item IDs might not exist in a data and different modal (i.e., text, images, and attributes) embeddings are not aligned. Meanwhile, users' interacted item sequence and users' search query sequence reflect users' interests from different perspectives. To address these problems, SEMINAR [186] leverages a pre-training search unit to learn query-item pairs. Then, SEMINAR leverages pre-trained embeddings to initialize the network. Additionally, SEMINAR leverages a multi-modal product quantization strategy to reduce time complexity. The basic framework of SEMINAR includes GSU and ESU. To model users' long sequences and capture evolving users' interests, HPMN [187] leverages memory slots to realize sequential patterns' personalized memorization for each user. To capture multi-scale and evolving sequence patterns, HPMN leverages a hierarchical incremental updating mechanism.

6.5 Data-augmented SR

There are two methods to conduct data augmentation. The first augmented methods only consider the information about item IDs. The second augmented methods that consider the role of side information. These approaches aim to alleviate the data sparsity problem and learn high-quality user representations. There are some SR models using data-augmented methods as follows.

Time intervals in users' interaction sequences might vary significantly, which might fail to model users' preferences due to preference drift. To solve this problem, TiCoSeRec [188] standardizes interaction sequences by transforming uneven time intervals into uniform distributions. TiCoSeRec consists of five operators (i.e., Ti-Crop, Ti-Reorder, Ti-Mask, Ti-Substitute, and Ti-Insert). Different from traditional operators, the five data augmentation methods explicitly account for time intervals. Finally, TiCoSeRec leverages contrastive learning to ensure a high similarity between the augmented and original sequences. CL4SRec [12] leverages three kinds of data augmentation approaches (i.e., item crop, item mask, and item reorder). Meanwhile, CL4SRec leverages contrastive learning between one augmented user interaction sequence and another augmented one. To alleviate the data sparsity problem and obtain stable users' interests, TGCL4SR [189] leverages two methods (i.e., neighbor sampling and time disturbance) to obtain augmented subgraphs. The neighbor sampling constructs augmented subgraphs by randomly sampling the neighbor nodes of each item. The time disturbance constructs the hypergraph to randomly add some noise that obeys Gaussian distribution to timestamps. Finally, by using disturbed graph contrastive loss and subgraph contrastive loss, high-quality representations can be obtained. Driven by the potential of data-centric AI, DR4SR+ [190] incorporates a model-aware data personalizer to tailor the regenerated data. In the process of model-agnostic data regeneration, a regenerator is pre-trained by learning transition patterns obtained through rule-based methods in

the original interaction sequences. DR4SR+ leverages a diversity-promoted regenerator to capture one-to-many mapping relationships between the original sequences and generated patterns. To generate a personalized data for target items, DR4SR+ uses a data personalizer to evaluate the quality of each generated sample for a target model. Traditional data augmentation methods often modify an original interaction sequence randomly. For example, some items are masked in sequences randomly. These methods might break sequential correlations. To solve this problem, MStein [191] leverages Wasserstein discrepancy measurement to measure mutual information among augmented sequences and ensures the maximum of the mutual information. MBASR [192] considers users' behaviors while data augmentation is conducted. DuoRec [193] leverages a contrastive learning regularizer to generate item embeddings with more uniform magnitude and frequency. END4Rec [194] improves the performance of Transformer by pruning noisy input. ASReP [195] leverages a pre-training Transformer to enhance the length of short users' interaction sequences, and then the Transformer is fine-tuned by augmented sequences.

7 Empirical studies

In this section, our survey will introduce some datasets, evaluation protocols and evaluation metrics used in sequential recommendation. Finally, we will present experimental results of some representative SR models.

7.1 Datasets

Training modality-based SR models with LLMs requires datasets that include multi-modal features. In our survey, we introduce some public datasets containing multi-modal features. Some specific details are shown in Table 4.

These datasets primarily contain three different modals (i.e., text, images, and videos). The strong sequential structure of these datasets is important for training SR models. Compared with other datasets in Table 4, Yelp and Steam might have weak sequential structure [196].

7.2 Evaluation protocols

Datasets can be divided into training, validation, and testing data using two primary methods. An early method is the leave-one-out strategy [1]. For this split method, all users' last interaction is split into the testing data, all users' penultimate interaction is split into the validation data, and the remaining interactions are split into the training data. This split method is simple and efficient, but it might

Table 4 The statistics about several widely used datasets

Dataset	Multi-modal features
MIND ⁸⁾	text
H-M ⁹⁾	text, images
Bili ¹⁰⁾	text, images
Art of the Mix ¹¹⁾	text
LastFM ¹²⁾	text
Amazon 2023 ¹³⁾	text, images, videos
MicroLens ¹⁴⁾	text, images, videos
MovieLens ¹⁵⁾	text
Steam ¹⁶⁾	text, images
Yelp ¹⁷⁾	text, images
Goodreads ¹⁸⁾	text, images
Douban ¹⁹⁾	text

lead to the problem of data leakage [197]. This is because interactions in the validation data may precede those in the training data based on their timestamps. The second method obtains the training, validation, and testing data by specific timestamps [116]. We can set two timestamps (i.e., T_1 and T_2). All interactions happening before timestamp T_1 are split into the training data. All interactions happening between timestamp T_1 and timestamp T_2 are split into the validation data. All interactions happening after timestamp T_2 are assigned to the testing data. This splitting method helps to prevent data leakage. However, this splitting method has a shortcoming that some users might not be evaluated. Because all interacted records for some users could fall within training data.

During the validation and testing phase, candidate items can be selected by three methods (i.e., random sampling, sampling based on popularity, and full sampling). Random sampling selects a small number of candidate items randomly [1]. This method is simple yet effective. Because the SR models are usually deployed in ranking or re-ranking phase, there are only a small number of candidate items. The number of candidate items is consistent with the actual situation. However, random sampling might lead to variability in the candidate items. The evaluation results might be unstable with the varying candidate items. Sampling based on popularity is that the candidate items are selected based on popularity [103]. If the popularity of

⁸⁾ Website at msnews.github.io (Last visited on May 12, 2025).

⁹⁾ Website at www.kaggle.com/c/h-and-m-personalized-fashion-recommendations/data (Last visited on May 12, 2025).

¹⁰⁾ Website at github.com/westlake-repl/NineRec?tab=readme-ov-file (Last visited on May 12, 2025).

¹¹⁾ Website at brianmcfee.net/data/aotm2011.html (Last visited on May 12, 2025).

¹²⁾ Website at www.kaggle.com/datasets/harshal19t/lastfm-dataset (Last visited on May 12, 2025).

¹³⁾ Website at huggingface.co/datasets/McAuley-Lab/Amazon-Reviews-2023 (Last visited on May 12, 2025).

¹⁴⁾ Website at github.com/westlake-repl/MicroLens/ (Last visited on May 12, 2025).

¹⁵⁾ Website at grouplens.org/datasets/movielens/ (Last visited on May 12, 2025).

¹⁶⁾ Website at huggingface.co/datasets/FronkonGames/steam-games-dataset (Last visited on May 12, 2025).

¹⁷⁾ Website at www.yelp.com/dataset (Last visited on May 12, 2025).

¹⁸⁾ Website at www.kaggle.com/datasets/jealousleopard/goodreadsbooks (Last visited on May 12, 2025).

¹⁹⁾ Website at www.kaggle.com/datasets/fengzhujoey/douban-datasetratingreviewside-information (Last visited on May 12, 2025).

items is larger, the probability of these items that are selected as candidate items is larger. This method can reduce the influence of randomness. Users interact with popular items easily, therefore these items appear in candidate items more frequently. Compared with random sampling, this split method is more stable and provides a more accurate evaluation of recommendation performance. The final method considers all items in the dataset as candidate items [190]. This method can evaluate the recommendation performance of different models from a comprehensive view. However, it might consume lots of computational resources especially when the dataset contains a large number of items.

7.3 Evaluation metrics

SR models aim to generate a top- N recommendation list. The evaluation metrics mainly aim to measure whether the user's most recently interacted item is in the top- N recommendation list and the position in the top- N recommendation list. Two commonly used evaluation metrics [1,18] are shown as follows.

Hit Ratio (HR) denotes the ratio of ground-truth items included in the top- N recommendation list. It can be formalized as:

$$HR@N = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \sigma(R_{u,t_u} \leq N),$$

where $\sigma(\cdot)$ is an indicator function. If the rank generated by an SR model is among top- N , the value of $\sigma(\cdot)$ is 1, otherwise is 0. Normalized Discounted Cumulative Gain ($NDCG$) evaluates the quality of the ranking by considering both the relevance and position of ground-truth items in the top- N recommendation list. It can be formalized as:

$$NDCG@N = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{\sigma(R_{u,t_u} \leq N)}{\log_2(1 + R_{u,t_u})}.$$

7.4 Experimental results

To show the relative recommendation performance of different SR models, we quote and copy the results of some recent works [138,147,149]. The compared SR models can be classified into five categories, including pure ID-based SR (i.e., SASRec [1] and BERT4Rec [36]), SR with side information (i.e., FDSA [121] and S³Rec [120]), multi-modal SR (i.e., UniSRec [8], MISSRec [138], and RECFORMER [137]), generative SR (i.e., TIGER [141] and SpecGR [147]), and LLM-powered SR (i.e., P5 [148], POD [150], and Lite-LLM4Rec [149]). “Office”, “Game”, and “Toy” are three domains in an Amazon dataset. Two widely-used metrics (i.e., HR@10 and NDCG@10) are adopted for performance evaluation. In each domain, the leave-one-out strategy is leveraged to obtain the training, validation and test data. Meanwhile, all items are taken as the candidate ones in recommendation. The data in each domain is processed according to the corresponding papers [138,147,149]. From the experimental results in Table 5, we can have the following observations:

- SR models with side information achieve better performance than pure ID-based SR ones, which demonstrates that leveraging item features can alleviate the cold-start and sparsity problem to some extent.
- By removing item IDs from multi-modal SR models, the recommendation performance will decrease. It is because that the item IDs can capture the collaborative signals and users'

Table 5 Experimental results of different SR models. The results are copied from papers [138,147,149] for direct comparison, where “–” means that the corresponding experimental results are not available in the original papers

Model	Office		Game		Toy	
	HR@10	NDCG@10	HR@10	NDCG@10	HR@10	NDCG@10
SASRec	0.1056	0.0710	0.0186	0.0093	0.0561	0.0312
BERT4Rec	0.0825	0.0634	–	–	0.0352	0.0179
FDSR	0.1118	0.0868	0.0190	0.0101	0.0568	0.0344
S ³ -Rec	0.1030	0.0653	0.0195	0.0094	0.0538	0.0276
UniSRec w/o IDs	0.1013	0.0619	–	–	–	–
UniSRec	0.1280	0.0831	0.0225	0.0115	–	–
MISSRec w/o IDs	0.1258	0.0795	–	–	–	–
MISSRec	0.1301	0.0842	–	–	–	–
RECFORMER	–	–	0.0243	0.0114	–	–
TIGER	–	–	0.0222	0.0114	–	–
SpecGR	–	–	0.0229	0.0115	–	–
P5	–	–	–	–	0.0543	0.0356
POD	–	–	–	–	0.0566	0.0421
Lite-LLM4Rec	–	–	–	–	0.0682	0.0488

fine-grained preferences.

- LLM-powered SR models show better recommendation performance than traditional SR models. We can see that the semantic knowledge extracted by LLMs are useful for enhancing the recommendation performance.
- Compared with pure ID-based and multi-modal SR models, generative SR models show comparable and even better performance. By leveraging semantic IDs, generative models can capture different grained semantic information.

■ 8 Prospects and future directions

SR models have been studied for several years, but there are certain areas that remain to be fully explored.

8.1 Open-domain SR

Traditional SR models are trained on a single and static data with predefined features. When new items emerge frequently and data resources are increasingly diverse, the traditional SR models are difficult to recommend personalized items. Leveraging the multi-modal features of items and LLMs, open-domain SR models can operate effectively in dynamic environments with a continuous influx of new items. Meanwhile, the open-domain SR models can be transferred across different platforms and domains. By transferring knowledge from external domains or platforms, the open-domain SR models can achieve better recommendation performance within the target domain or platform.

8.2 Data-centric SR

To improve the recommendation performance, data-centric SR models mainly focus on improving the quality of data rather than refining model structures. The data is important in sequential recommendation tasks. The data-centric SR models mainly leverage three techniques (i.e., data generation, data denoising, and data debiasing). The data generation has been researched for many years. However, challenges remain in generating high-quality data for certain existing datasets. The quality of the generated data cannot be measured by a specific method. Some SR models leverage data augmentation to generate additional data, aiming to improve the recommendation performance. However, SR models need to be trained in a larger number of data, which might consume more computational resources and training time. Therefore, we need to design a data augmentation method to improve the quality of generated data and reduce the training time. In the meantime, we can consider some multi-modal features to generate some high-quality data.

8.3 Cloud-edge collaborative SR

Cloud-edge collaborative SR models combine the strengths of cloud-based and edge-based SR models. Compared with conventional SR models, cloud-edge collaborative SR models not only can improve the recommendation performance but also protect user privacy and leverage resources more effectively. The conventional SR models often require lots of computational resources. Therefore, the conventional SR models are hard to deploy on their own computers. Moreover, transmitting interaction data to a server could risk

exposing private information. Therefore, the cloud-edge collaborative SR models are meaningful. Federated SR models are one such approach specifically designed to protect privacy.

8.4 Continuous SR

In many short-video scenarios, such as TikTok, the length of a user's interaction sequence tends to grow over time. Some ultra-long interaction sequences might influence the efficiency of SR models. Therefore, how to design an SR model that can handle ultra-long interaction sequences is a challenge. Meanwhile, a user's interests might evolve over time. Therefore, continuous SR models need to update the user's interests dynamically based on recent interactions.

8.5 SR for Good

Some existing SR models aim to improve economic growth. However, SR models for good aim to recommend personalized items benefit both individuals and society while also supporting economic development. For example, to reduce carbon emissions, it is important to recommend environmentally friendly items for users. In educational field, SR models should recommend learning materials to help users learn more effectively. In healthy field, SR models should recommend personalized exercises to help users improve their physical fitness and overall health. The SR models for good remain to be explored.

8.6 Explainable SR

Some SR models based on multi-modal features and large language models focus on improving the recommendation accuracy but often fail to provide explanations for why these related items are recommended. LLMs excel at processing and understanding multi-modal features (i.e., text, images, and so on), which provides an opportunity to generate some explanations. Therefore, in future work, we can leverage large language models and multi-modal features of interacted items to offer more interpretation in recommendations.

■ 9 Conclusions

In this survey, we have shown that SR models have achieved significant advancements in recommending personalized items for users. From their development, we can have some observations. Firstly, SR models combining IDs and general features can tackle the data sparsity and cold-start problems better than pure ID-based SR models. Secondly, modal-based SR models are often associated with better transferability than pure ID-based SR models and SR models combining IDs and general features. Thirdly, modal-based SR models are of high time complexity and need to consume more training time than pure ID-based SR models. Fourthly, multi-modal SR models using IDs can learn fine-grained preferences and achieve more accurate recommendation performance. Fifthly, LLMs-based SR models and generative SR models show further potential for improving recommendation effectiveness. Sixthly, SR models can effectively tackle ultra-long interaction sequences by retrieving a few items from these sequences. Finally, data augmentation methods can alleviate the data sparsity and cold-start problems to some extent.

As for future directions, we discuss several interesting topics worthy of exploration, which are about scenario, data, architecture, time, society, and explainability.

■ Acknowledgements

We thank the support for the National Natural Science Foundation of China (Grant Nos. 62461160311, 62172283 and 62272315), and Guangdong Basic and Applied Basic Research Foundation (No. 2024A1515010122).

■ Competing interests

The authors declare that they have no competing interests or financial conflicts to disclose.

■ Open Access

This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

■ References

- [1] Kang W C, McAuley J. Self-attentive sequential recommendation. In: Proceedings of 2018 IEEE International Conference on Data Mining. 2018, 197–206
- [2] He R, Kang W C, McAuley J. Translation-based recommendation. In: Proceedings of the 11th ACM Conference on Recommender Systems. 2017, 161–169
- [3] Yuan X, Duan D, Tong L, Shi L, Zhang C. ICAI-SR: item categorical attribute integrated sequential recommendation. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2021, 1687–1691
- [4] Liu C, Li X, Cai G, Dong Z, Zhu H, Shang L. Noninvasive self-attention for side information fusion in sequential recommendation. In: Proceedings of the 35th AAAI Conference on Artificial intelligence. 2021, 4249–4256
- [5] Devlin J, Chang M W, Lee K, Toutanova K. BERT: pre-training of deep bidirectional transformers for language understanding. In: Proceedings of 2019 Conference of the North American Chapter of the Association for Computational Linguistics. 2019, 4171–4186
- [6] He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition. 2016, 770–778
- [7] Yuan Z, Yuan F, Song Y, Li Y, Fu J, Yang F, Pan Y, Ni Y. Where to go next for recommender systems? ID-vs. Modality-based recommender models revisited. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2023, 2639–2649
- [8] Hou Y, Mu S, Zhao W X, Li Y, Ding B, Wen J R. Towards universal sequence representation learning for recommender systems. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 2022, 585–593
- [9] Ji J, Li Z, Xu S, Hua W, Ge Y, Tan J, Zhang Y. GenRec: large language model for generative recommendation. In: Proceedings of the 46th European Conference on Information Retrieval. 2024, 494–502
- [10] Boz A, Zorgdrager W, Kotti Z, Harte J, Louridas P, Karakoidas V, Jannach D, Fragkoulis M. Improving sequential recommendations with LLMs. *ACM Transactions on Recommender Systems*, 2024, doi: 10.1145/3711667
- [11] Pi Q, Zhou G, Zhang Y, Wang Z, Ren L, Fan Y, Zhu X, Gai K. Search-based user interest modeling with lifelong sequential behavior data for click-through rate prediction. In: Proceedings of the 29th ACM International Conference on Information & Knowledge Management. 2020, 2685–2692
- [12] Xie X, Sun F, Liu Z, Wu S, Gao J, Zhang J, Ding B, Cui B. Contrastive learning for sequential recommendation. In: Proceedings of the 38th IEEE International Conference on Data Engineering. 2022, 1259–1273
- [13] Wang S, Cao L, Wang Y, Sheng Q Z, Orgun M A, Lian D. A survey on session-based recommender systems. *ACM Computing Surveys (CSUR)*, 2021, 54(7): 154
- [14] Li Z, Yang C, Chen Y, Wang X, Chen H, Xu G, Yao L, Sheng M. Graph and sequential neural networks in session-based recommendation: a survey. *ACM Computing Surveys*, 2024, 57(2): 40
- [15] Quadrona M, Cremonesi P, Jannach D. Sequence-aware recommender systems. *ACM Computing Surveys (CSUR)*, 2018, 51(4): 66
- [16] Nasir M, Ezeife C I. A survey and taxonomy of sequential recommender systems for e-commerce product recommendation. *SN Computer Science*, 2023, 4(6): 708
- [17] Fang H, Zhang D, Shu Y, Guo G. Deep learning for sequential recommendation: algorithms, influential factors, and evaluations. *ACM Transactions on Information Systems (TOIS)*, 2020, 39(1): 10
- [18] Boka T F, Niu Z, Neupane R B. A survey of sequential recommendation systems: techniques, evaluation, and future directions. *Information Systems*, 2024, 125: 102427
- [19] Chen X, Li Z, Pan W, Ming Z. A survey on multi-behavior sequential recommendation, 2023. arXiv preprint arXiv: 2308.15701
- [20] Chen S, Xu Z, Pan W, Yang Q, Ming Z. A survey on cross-domain sequential recommendation. In: Proceedings of the 33rd International Joint Conference on Artificial Intelligence. 2024, 884
- [21] Jing M, Zhu Y, Zang T, Wang K. Contrastive self-supervised learning in recommender systems: a survey. *ACM Transactions on Information Systems*, 2023, 42(2): 59
- [22] Lai R, Chen L, Chen R, Zhang C. A survey on data-centric recommender systems, 2024. arXiv preprint arXiv: 2401.17878
- [23] Zhang X, Xu B, Li C, Zhou Y, Li L, Lin H. Side information-driven session-based recommendation: a survey, 2024. arXiv preprint arXiv: 2402.17129
- [24] Hidasi B, Karatzoglou A, Baltrunas L, Tikk D. Session-based recommendations with recurrent neural networks. In: Proceedings of the 4th International Conference on Learning Representations. 2016

- [25] Xie Y, Zhou P, Kim S. Decoupled side information fusion for sequential recommendation. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2022, 1611–1621
- [26] Li Y, Ding Y, Chen B, Xin X, Wang Y, Shi Y, Tang R, Wang D. Extracting attentive social temporal excitation for sequential recommendation. In: Proceedings of the 30th ACM International Conference on Information & Knowledge Management. 2021, 998–1007
- [27] Hou Y, He Z, McAuley J, Zhao W X. Learning vector-quantized item representation for transferable sequential recommenders. In: Proceedings of ACM Web Conference 2023. 2023, 1162–1171
- [28] Mobasher B, Dai H, Luo T, Nakagawa M. Using sequential and non-sequential patterns in predictive web usage mining tasks. In: Proceedings of 2002 IEEE International Conference on Data Mining. 2002, 669–672
- [29] Rendle S, Freudenthaler C, Schmidt-Thieme L. Factorizing personalized Markov chains for next-basket recommendation. In: Proceedings of the 19th International Conference on World Wide Web. 2010, 811–820
- [30] Hidasi B, Karatzoglou A. Recurrent neural networks with top-k gains for session-based recommendations. In: Proceedings of the 27th ACM International Conference on Information and Knowledge Management. 2018, 843–852
- [31] Lin X, Niu S, Wang Y, Li Y. K-plet recurrent neural networks for sequential recommendation. In: Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. 2018, 1057–1060
- [32] Niu S, Zhang R. Collaborative sequence prediction for sequential recommender. In: Proceedings of 2017 ACM on Conference on Information and Knowledge Management. 2017, 2239–2242
- [33] Chen Q, Li G, Zhou Q, Shi S, Zou D. Double attention convolutional neural network for sequential recommendation. *ACM Transactions on the Web*, 2022, 16(4): 19
- [34] Tang J, Wang K. Personalized top-N sequential recommendation via convolutional sequence embedding. In: Proceedings of the 11th ACM International Conference on Web Search and Data Mining. 2018, 565–573
- [35] Yan A, Cheng S, Kang W C, Wan M, McAuley J. CosRec: 2D convolutional neural networks for sequential recommendation. In: Proceedings of the 28th ACM International Conference on Information and Knowledge Management. 2019, 2173–2176
- [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. 2019, 1441–1450
- [37] Petrov A, Macdonald C. A systematic review and replicability study of BERT4Rec for sequential recommendation. In: Proceedings of the 16th ACM Conference on Recommender Systems. 2022, 436–447
- [38] Zhang Y, Wang X, Chen H, Zhu W. Adaptive disentangled transformer for sequential recommendation. In: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 2023, 3434–3445
- [39] Zhou P, Ye Q, Xie Y, Gao J, Wang S, Kim J B, You C, Kim S. Attention calibration for Transformer-based sequential recommendation. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. 2023, 3595–3605
- [40] Liu Z, Liu S, Zhang Z, Cai Q, Zhao X, Zhao K, Hu L, Jiang P, Gai K. Sequential recommendation for optimizing both immediate feedback and long-term retention. In: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2024, 1872–1882
- [41] Fan Z, Liu Z, Wang Y, Wang A, Nazari Z, Zheng L, Peng H, Yu P S. Sequential recommendation via stochastic self-attention. In: Proceedings of ACM Web Conference 2022. 2022, 2036–2047
- [42] Ren R, Liu Z, Li Y, Zhao W X, Wang H, Ding B, Wen J R. Sequential recommendation with self-attentive multi-adversarial network. In: Proceedings of the 43rd International ACM SIGIR conference on Research and Development in Information Retrieval. 2020, 89–98
- [43] Wu L, Li S, Hsieh C J, Sharpnack J. SSE-PT: sequential recommendation via personalized transformer. In: Proceedings of the 14th ACM Conference on Recommender Systems. 2020, 328–337
- [44] Li C, Wang Y, Liu Q, Zhao X, Wang W, Wang Y, Zou L, Fan W, Li Q. STRec: sparse transformer for sequential recommendations. In: Proceedings of the 17th ACM Conference on Recommender Systems. 2023, 101–111
- [45] Cui Z, Cai Y, Wu S, Ma X, Wang L. Motif-aware sequential recommendation. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2021, 1738–1742
- [46] Guo N, Liu X, Li S, Ma Q, Gao K, Han B, Zheng L, Guo S, Guo X. Poincaré heterogeneous graph neural networks for sequential recommendation. *ACM Transactions on Information Systems*, 2023, 41(3): 1–26
- [47] Du H, Shi H, Zhao P, Wang D, Sheng V S, Liu Y, Liu G, Zhao L. Contrastive learning with bidirectional transformers for sequential recommendation. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management. 2022, 396–405
- [48] Yang Y, Huang C, Xia L, Huang C, Luo D, Lin K. Debaised contrastive learning for sequential recommendation. In: Proceedings of ACM Web Conference 2023. 2023, 1063–1073
- [49] Lin G, Gao C, Zheng Y, Chang J, Niu Y, Song Y, Li Z, Jin D, Li Y. Dual-interest factorization-heads attention for sequential recommendation. In: Proceedings of ACM Web Conference 2023. 2023, 917–927
- [50] Xu C, Xu J, Chen X, Dong Z, Wen J R. Dually enhanced propensity score estimation in sequential recommendation. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management. 2022, 2260–2269
- [51] Ma M, Ren P, Chen Z, Ren Z, Liang H, Ma J, De Rijke M. Improving transformer-based sequential recommenders through preference editing. *ACM Transactions on Information Systems*, 2023, 41(3): 71
- [52] Xi X, Zhao Y, Liu Q, Ouyang L, Wu Y. Integrating offline reinforcement learning with transformers for sequential recommendation. In: Proceedings of the 17th ACM Conference on Recommender Systems. 2023, 1
- [53] Fan Z, Liu Z, Wang S, Zheng L, Yu P S. Modeling sequences as

- distributions with uncertainty for sequential recommendation. In: Proceedings of the 30th ACM International Conference on Information & Knowledge Management. 2021, 3019–3023
- [54] de Souza Pereira Moreira G, Rabhi S, Lee J M, Ak R, Oldridge E. Transformers4Rec: Bridging the gap between NLP and sequential/session-based recommendation. In: Proceedings of the 15th ACM Conference on Recommender Systems. 2021, 143–153
- [55] Zhu Y, Huang B, Jiang S, Yang M, Yang Y, Zhong W. Progressive self-attention network with unsymmetrical positional encoding for sequential recommendation. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2022, 2029–2033
- [56] Lai V, Chen H, Yeh C C M, Xu M, Cai Y, Yang H. Enhancing transformers without self-supervised learning: a loss landscape perspective in sequential recommendation. In: Proceedings of the 17th ACM Conference on Recommender Systems. 2023, 791–797
- [57] Liu Q, Yan F, Zhao X, Du Z, Guo H, Tang R, Tian F. Diffusion augmentation for sequential recommendation. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. 2023, 1576–1586
- [58] Li Z, Sun A, Li C. DiffuRec: a diffusion model for sequential recommendation. *ACM Transactions on Information Systems*, 2023, 42(3): 66
- [59] Wu Z, Wang X, Chen H, Li K, Han Y, Sun L, Zhu W. Diff4Rec: sequential recommendation with curriculum-scheduled diffusion augmentation. In: Proceedings of the 31st ACM International Conference on Multimedia. 2023, 9329–9335
- [60] Liu S, Liu J, Gu H, Li D, Lu T, Zhang P, Gu N. AutoSeqRec: autoencoder for efficient sequential recommendation. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. 2023, 1493–1502
- [61] Chen X, Li Q. Causality-driven user modeling for sequential recommendations over time. In: Companion Proceedings of ACM Web Conference 2024. 2024, 1400–1406
- [62] Ding Y, Ma Y, Wong W K, Chua T S. Leveraging two types of global graph for sequential fashion recommendation. In: Proceedings of 2021 International Conference on Multimedia Retrieval. 2021, 73–81
- [63] Peng B, Chen Z, Parthasarathy S, Ning X. Modeling sequences as star graphs to address over-smoothing in self-attentive sequential recommendation. *ACM Transactions on Knowledge Discovery from Data*, 2023, 18(8): 207
- [64] Zhou K, Yu H, Zhao W X, Wen J R. Filter-enhanced MLP is all you need for sequential recommendation. In: Proceedings of ACM Web Conference 2022. 2022, 2388–2399
- [65] Jiang Y, Xu Y, Yang Y, Yang F, Wang P, Li C, Zhuang F, Xiong H. TRIMLP: a foundational MLP-like architecture for sequential recommendation. *ACM Transactions on Information Systems*, 2024, 42(6): 157
- [66] Long C, Yuan H, Fang J, Xian X, Liu G, Sheng V S, Zhao P. Learning global and multi-granularity local representation with MLP for sequential recommendation. *ACM Transactions on Knowledge Discovery from Data*, 2024, 18(4): 87
- [67] Sachdeva N, Manco G, Ritacco E, Pudi V. Sequential variational autoencoders for collaborative filtering. In: Proceedings of the 12th ACM International Conference on Web Search and Data Mining. 2019, 600–608
- [68] You J, Wang Y, Pal A, Eksombatchai P, Rosenberg C, Leskovec J. Hierarchical temporal convolutional networks for dynamic recommender systems. In: Proceedings of World Wide Web Conference. 2019, 2236–2246
- [69] Zhang Q, Wu B, Sun Z, Ye Y. Gating augmented capsule network for sequential recommendation. *Knowledge-Based Systems*, 2022, 247: 108817
- [70] Wang Z, Shen Y. Time-aware multi-interest capsule network for sequential recommendation. In: Proceedings of 2022 SIAM International Conference on Data Mining. 2022, 558–566
- [71] Zhao X, Zhang L, Xia L, Ding Z, Yin D, Tang J. Deep reinforcement learning for list-wise recommendations, 2019. arXiv preprint arXiv: 1801.00209
- [72] Antaris S, Rafailidis D. Sequence adaptation via reinforcement learning in recommender systems. In: Proceedings of the 15th ACM Conference on Recommender Systems. 2021, 714–718
- [73] Jiang J, Zhang P, Luo Y, Li C, Kim J B, Zhang K, Wang S, Xie X, Kim S. AdaMCT: adaptive mixture of CNN-Transformer for sequential recommendation. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. 2023, 976–986
- [74] Wang J, Liu Q, Liu Z, Wu S. Towards accurate and interpretable sequential prediction: a CNN & attention-based feature extractor. In: Proceedings of the 28th ACM International Conference on Information and Knowledge Management. 2019, 1703–1712
- [75] Huang X, Sang J, Yu J, Xu C. Learning to learn a cold-start sequential recommender. *ACM Transactions on Information Systems (TOIS)*, 2022, 40(2): 30
- [76] Chen H, Huang B, Wang X, Zhou Y, Zhu W. Global-local graphformer: towards better understanding of user intentions in sequential recommendation. In: Proceedings of the 5th ACM International Conference on Multimedia in Asia. 2024, 8
- [77] Ye Y, Xia L, Huang C. Graph masked autoencoder for sequential recommendation. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2023, 321–330
- [78] Zhang W, Chen Z, Zha H, Wang J. Learning from substitutable and complementary relations for graph-based sequential product recommendation. *ACM Transactions on Information Systems (TOIS)*, 2021, 40(2): 26
- [79] Yin M, Wang H, Xu X, Wu L, Zhao S, Guo W, Liu Y, Tang R, Lian D, Chen E. APGL4SR: a generic framework with adaptive and personalized global collaborative information in sequential recommendation. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. 2023, 3009–3019
- [80] Baikolov V, Frolov E. End-to-end graph-sequential representation learning for accurate recommendations. In: Companion Proceedings of ACM Web Conference 2024. 2024, 501–504
- [81] Tan L, Xu J, Gong D, Liu F. Recurrent attentive neural networks for sequential recommendation. In: Proceedings of 2023 International

- Conference on Communication Network and Machine Learning. 2024, 356–360
- [82] Xu C, Zhao P, Liu Y, Xu J, Sheng V S S S, Cui Z, Zhou X, Xiong H. Recurrent convolutional neural network for sequential recommendation. In: Proceedings of World Wide Web Conference. 2019, 3398–3404
- [83] Yu M, Zhu K, Zhao M, Yu J, Xu T, Jin D, Li X, Yu R. Learning neighbor user intention on user–item interaction graphs for better sequential recommendation. *ACM Transactions on the Web*, 2024, 18(2): 21
- [84] Tang J, Belletti F, Jain S, Chen M, Beutel A, Xu C, Chi E H. Towards neural mixture recommender for long range dependent user sequences. In: Proceedings of World Wide Web Conference. 2019, 1782–1793
- [85] Alharbi N, Caragea D. Cross-domain self-attentive sequential recommendations. In: Saraswat M, Roy S, Chowdhury C, Gandomi A H, eds. Proceedings of International Conference on Data Science and Applications. Singapore: Springer, 2022, 601–614
- [86] Ma H, Xie R, Meng L, Chen X, Zhang X, Lin L, Zhou J. Triple sequence learning for cross-domain recommendation. *ACM Transactions on Information Systems*, 2024, 42(4): 91
- [87] Cao J, Cong X, Sheng J, Liu T, Wang B. Contrastive cross-domain sequential recommendation. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management. 2022, 138–147
- [88] Li J, Wang Y, McAuley J. Time interval aware self-attention for sequential recommendation. In: Proceedings of the 13th International Conference on Web Search and Data Mining. 2020, 322–330
- [89] Cho J, Hyun D, Kang S, Yu H. Learning heterogeneous temporal patterns of user preference for timely recommendation. In: Proceedings of Web Conference 2021. 2021, 1274–1283
- [90] Yuan E, Guo W, He Z, Guo H, Liu C, Tang R. Multi-behavior sequential transformer recommender. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2022, 1642–1652
- [91] Su J, Chen C, Lin Z, Li X, Liu W, Zheng X. Personalized behavior-aware transformer for multi-behavior sequential recommendation. In: Proceedings of the 31st ACM International Conference on Multimedia. 2023, 6321–6331
- [92] Li X, Chen H, Yu J, Zhao M, Xu T, Zhang W, Yu M. Global heterogeneous graph and target interest denoising for multi-behavior sequential recommendation. In: Proceedings of the 17th ACM International Conference on Web Search and Data Mining. 2024, 387–395
- [93] Chen Z, Zhang W, Yan J, Wang G, Wang J. Learning dual dynamic representations on time-sliced user-item interaction graphs for sequential recommendation. In: Proceedings of the 30th ACM International Conference on Information & Knowledge Management. 2021, 231–240
- [94] Tran V A, Salha-Galvan G, Sguerra B, Hennequin R. Attention mixtures for time-aware sequential recommendation. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2023, 1821–1826
- [95] 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. 2021, 433–442
- [96] Chen Z, Wang W. Time interval-aware graph with self-attention for sequential recommendation. In: Proceedings of 2022 5th International Conference on Algorithms, Computing and Artificial Intelligence. 2023, 71
- [97] Du X, Li X. Moving forward together: a multi-user pattern extraction neural model for sequential recommendations. In: Proceedings of 2023 6th International Conference on Artificial Intelligence and Pattern Recognition. 2024, 1355–1361
- [98] Wang J, Ding K, Hong L, Liu H, Caverlee J. Next-item recommendation with sequential hypergraphs. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval. 2020, 1101–1110
- [99] Huang L, Ma Y, Liu Y, Danny Du B, Wang S, Li D. Position-enhanced and time-aware graph convolutional network for sequential recommendations. *ACM Transactions on Information Systems*, 2023, 41(1): 6
- [100] Zhang Y, He Y, Wang J, Caverlee J. Adaptive hierarchical translation-based sequential recommendation. In: Proceedings of the Web Conference 2020. 2020, 2984–2990
- [101] Huang C, Wang S, Wang X, Yao L. Dual contrastive transformer for hierarchical preference modeling in sequential recommendation. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2023, 99–109
- [102] Li M, Zhang Z, Zhao X, Wang W, Zhao M, Wu R, Guo R. AutoMLP: automated MLP for sequential recommendations. In: Proceedings of ACM Web Conference 2023. 2023, 1190–1198
- [103] Xu Z, Pan W, Ming Z. A multi-view graph contrastive learning framework for cross-domain sequential recommendation. In: Proceedings of the 17th ACM Conference on Recommender Systems. 2023, 491–501
- [104] Zheng X, Su J, Liu W, Chen C. DDGHM: dual dynamic graph with hybrid metric training for cross-domain sequential recommendation. In: Proceedings of the 30th ACM International Conference on Multimedia. 2022, 471–481
- [105] Li Y, Li Q, Meng S, Hou J. Transformer-based rating-aware sequential recommendation. In: Proceedings of the 21st International Conference on Algorithms and Architectures for Parallel Processing. 2021, 759–774
- [106] Zhang X, Xu B, Yang L, Li C, Ma F, Liu H, Lin H. Price DOES matter!: modeling price and interest preferences in session-based recommendation. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2022, 1684–1693
- [107] Huang X, Fang Q, Qian S, Sang J, Li Y, Xu C. Explainable interaction-driven user modeling over knowledge graph for sequential recommendation. In: Proceedings of the 27th ACM International Conference on Multimedia. 2019, 548–556
- [108] Huang J, Zhao W X, Dou H, Wen J R, Chang E Y. Improving sequential recommendation with knowledge-enhanced memory networks. In: Proceedings of the 41st International ACM SIGIR Conference on

Research & Development in Information Retrieval. 2018, 505–514

- [109] Zhai J, Zheng X, Wang C D, Li H, Tian Y. Knowledge prompting for sequential recommendation. In: Proceedings of the 31st ACM International Conference on Multimedia. 2023, 6451–6461
- [110] Wang P, Fan Y, Xia L, Zhao W X, Niu S, Huang J. KERL: a knowledge-guided reinforcement learning model for sequential recommendation. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval. 2020, 209–218
- [111] Hu H, Guo W, Liu X, Liu Y, Tang R, Zhang R, Kan M Y. User behavior enriched temporal knowledge graphs for sequential recommendation. In: Proceedings of the 17th ACM International Conference on Web Search and Data Mining. 2024, 266–275
- [112] Chen J, Cao Y, Zhang F, Sun P, Wei K. Sequential intention-aware recommender based on user interaction graph. In: Proceedings of 2022 International Conference on Multimedia Retrieval. 2022, 118–126
- [113] Chen T, Wong R C W. An efficient and effective framework for session-based social recommendation. In: Proceedings of the 14th ACM International Conference on Web Search and Data Mining. 2021, 400–408
- [114] Li C, Niu X, Luo X, Chen Z, Quan C. A review-driven neural model for sequential recommendation. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence. 2019, 2866–2872
- [115] Zhang J, Mu X, Zhao P, Kang K, Ma C. Improving current interest with item and review sequential patterns for sequential recommendation. *Engineering Applications of Artificial Intelligence*, 2021, 104: 104348
- [116] Huang B, Luo J, Du W, Pan W, Ming Z. Cascaded cross attention for review-based sequential recommendation. In: Proceedings of 2023 IEEE International Conference on Data Mining. 2023, 170–179
- [117] Li M, Zhao X, Lyu C, Zhao M, Wu R, Guo R. MLP4Rec: a pure MLP architecture for sequential recommendations. In: Proceedings of the 31st International Joint Conference on Artificial Intelligence. 2022, 2138–2144
- [118] Gao J, Zhao X, Li M, Zhao M, Wu R, Guo R, Liu Y, Yin D. SMLP4Rec: an efficient all-MLP architecture for sequential recommendations. *ACM Transactions on Information Systems*, 2024, 42(3): 86
- [119] Hsu C, Li C T. RetaGNN: relational temporal attentive graph neural networks for holistic sequential recommendation. In: Proceedings of Web Conference 2021. 2021, 2968–2979
- [120] Zhou K, Wang H, Zhao W X, Zhu Y, Wang S, Zhang F, Wang Z, Wen J R. S3-Rec: Self-supervised learning for sequential recommendation with mutual information maximization. In: Proceedings of the 29th ACM international Conference on Information & Knowledge Management. 2020, 1893–1902
- [121] Zhang T, Zhao P, Liu Y, Sheng V S, Xu J, Wang D, Liu G, Zhou X. Feature-level deeper self-attention network for sequential recommendation. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence. 2019, 4320–4326
- [122] Hu H, Guo W, Liu Y, Kan M Y. Adaptive multi-modalities fusion in sequential recommendation systems. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. 2023, 843–853
- [123] Wang S, Shen B, Min X, He Y, Zhang X, Zhang L, Zhou J, Mo L. Aligned side information fusion method for sequential recommendation. In: Companion Proceedings of ACM Web Conference 2024. 2024, 112–120
- [124] Zhou X, Lumbantoruan R, Ren Y, Chen L, Yang X, Shao J. Dynamic bi-layer graph learning for context-aware sequential recommendation. *ACM Transactions on Recommender Systems*, 2024, 2(2): 14
- [125] Smirnova E, Vasile F. Contextual sequence modeling for recommendation with recurrent neural networks. In: Proceedings of the 2nd Workshop on Deep Learning for Recommender Systems. 2017, 2–9
- [126] Rakkappan L, Rajan V. Context-aware sequential recommendations with stacked recurrent neural networks. In: Proceedings of World Wide Web Conference. 2019, 3172–3178
- [127] Zhang J, Lin F, Yang C, Jiang W. A new sequential prediction framework with spatial-temporal embedding. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2022, 2282–2286
- [128] Xue L, Yang D, Zhai S, Li Y, Xiao Y. Learning dual-view user representations for enhanced sequential recommendation. *ACM Transactions on Information Systems*, 2023, 41(4): 86
- [129] Zheng L, Chai H, Chen X, Jin J, Zhang W, Yu Y, Guo X, Ge C, Feng Z. Search-based time-aware graph-enhanced recommendation with sequential behavior data. *ACM Transactions on Recommender Systems*, 2024, 2(4): 26
- [130] Li H, Yu L, Niu X, Leng Y, Du Q. Sequential and graphical cross-domain recommendations with a multi-view hierarchical transfer gate. *ACM Transactions on Knowledge Discovery from Data*, 2023, 18(1): 8
- [131] Ji W, Liu X, Zhang A, Wei Y, Ni Y, Wang X. Online distillation-enhanced multi-modal Transformer for sequential recommendation. In: Proceedings of the 31st ACM International Conference on Multimedia. 2023, 955–965
- [132] Zhang J, Cheng Y, Ni Y, Pan Y, Yuan Z, Fu J, Li Y, Wang J, Yuan F. NineRec: a benchmark dataset suite for evaluating transferable recommendation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024, doi: [10.1109/TPAMI.2024.3373868](https://doi.org/10.1109/TPAMI.2024.3373868)
- [133] Mu S, Hou Y, Zhao W X, Li Y, Ding B. ID-agnostic user behavior pre-training for sequential recommendation. In: Proceedings of the 28th China Conference on Information Retrieval. 2022, 16–27
- [134] Ni Y, Cheng Y, Liu X, Fu J, Li Y, He X, Zhang Y, Yuan F. A content-driven micro-video recommendation dataset at scale, 2023. arXiv preprint arXiv: 2309.15379
- [135] Qu Z, Xie R, Xiao C, Kang Z, Sun X. The elephant in the room: rethinking the usage of pre-trained language model in sequential recommendation. In: Proceedings of the 18th ACM Conference on Recommender Systems. 2024, 53–62
- [136] Zivic P, Vazquez H, Sánchez J. Scaling sequential recommendation models with transformers. In: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2024, 1567–1577
- [137] Li J, Wang M, Li J, Fu J, Shen X, Shang J, McAuley J. Text is all

- you need: learning language representations for sequential recommendation. In: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 2023, 1258–1267
- [138] Wang J, Zeng Z, Wang Y, Wang Y, Lu X, Li T, Yuan J, Zhang R, Zheng H T, Xia S T. MISSRec: Pre-training and transferring multi-modal interest-aware sequence representation for recommendation. In: Proceedings of the 31st ACM International Conference on Multimedia. 2023, 6548–6557
- [139] Liang J, Zhao X, Li M, Zhang Z, Wang W, Liu H, Liu Z. MMMLP: multi-modal multilayer perceptron for sequential recommendations. In: Proceedings of ACM Web Conference 2023. 2023, 1109–1117
- [140] Hidasi B, Quadrana M, Karatzoglou A, Tikk D. Parallel recurrent neural network architectures for feature-rich session-based recommendations. In: Proceedings of the 10th ACM Conference on Recommender Systems. 2016, 241–248
- [141] Rajput S, Mehta N, Singh A, Keshavan R, Vu T, Heidt L, Hong L, Tay Y, Tran V Q, Samost J, Kula M, Chi E H. Recommender systems with generative retrieval. In: Proceedings of the 37th International Conference on Neural Information Processing Systems. 2023, 452
- [142] Liu Z, Hou Y, McAuley J. Multi-behavior generative recommendation. In: Proceedings of the 33rd ACM International Conference on Information and Knowledge Management. 2024, 1575–1585
- [143] Wang W, Bao H, Lin X, Zhang J, Li Y, Feng F, Ng S K, Chua T S. Learnable item tokenization for generative recommendation. In: Proceedings of the 33rd ACM International Conference on Information and Knowledge Management. 2024, 2400–2409
- [144] Wang Y, Xun J, Hong M, Zhu J, Jin T, Lin W, Li H, Li L, Xia Y, Zhao Z, Dong Z. EAGER: two-stream generative recommender with behavior-semantic collaboration. In: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 2024, 3245–3254
- [145] Zhu J, Jin M, Liu Q, Qiu Z, Dong Z, Li X. CoST: contrastive quantization based semantic tokenization for generative recommendation. In: Proceedings of the 18th ACM Conference on Recommender Systems. 2024, 969–974
- [146] Liu H, Wei Y, Song X, Guan W, Li Y F, Nie L. MMGRec: multimodal generative recommendation with transformer model, 2024. arXiv preprint arXiv: 2404.16555
- [147] Ding Y, Hou Y, Li J, McAuley J. Inductive generative recommendation via retrieval-based speculation, 2024. arXiv preprint arXiv: 2410.02939
- [148] Geng S, Liu S, Fu Z, Ge Y, Zhang Y. Recommendation as language processing (RLP): a unified pretrain, personalized prompt & predict paradigm (P5). In: Proceedings of the 16th ACM Conference on Recommender Systems. 2022, 299–315
- [149] Wang H, Liu X, Fan W, Zhao X, Kini V, Yadav D, Wang F, Wen Z, Tang J, Liu H. Rethinking large language model architectures for sequential recommendations, 2024. arXiv preprint arXiv: 2402.09543
- [150] Li L, Zhang Y, Chen L. Prompt distillation for efficient LLM-based recommendation. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. 2023, 1348–1357
- [151] Wang Y, Tian C, Hu B, Yu Y, Liu Z, Zhang Z, Zhou J, Pang L, Wang X. Can small language models be good reasoners for sequential recommendation? In: Proceedings of ACM Web Conference 2024. 2024, 3876–3887
- [152] Zheng Z, Chao W, Qiu Z, Zhu H, Xiong H. Harnessing large language models for text-rich sequential recommendation. In: Proceedings of ACM Web Conference 2024. 2024, 3207–3216
- [153] Yang Z, Wu J, Luo Y, Zhang J, Yuan Y, Zhang A, Wang X, He X. Large language model can interpret latent space of sequential recommender, 2023. arXiv preprint arXiv: 2310.20487
- [154] Bao K, Zhang J, Zhang Y, Wang W, Feng F, He X. TALLRec: an effective and efficient tuning framework to align large language model with recommendation. In: Proceedings of the 17th ACM Conference on Recommender Systems. 2023, 1007–1014
- [155] Cui Y, Liu F, Wang P, Wang B, Tang H, Wan Y, Wang J, Chen J. Distillation matters: empowering sequential recommenders to match the performance of large language models. In: Proceedings of the 18th ACM Conference on Recommender Systems. 2024, 507–517
- [156] Li X, Chen C, Zhao X, Zhang Y, Xing C. E4SRec: an elegant effective efficient extensible solution of large language models for sequential recommendation, 2023. arXiv preprint arXiv: 2312.02443
- [157] Shen T, Wang H, Zhang J, Zhao S, Li L, Chen Z, Lian D, Chen E. Exploring user retrieval integration towards large language models for cross-domain sequential recommendation, 2024. arXiv preprint arXiv: 2406.03085
- [158] Ren Y, Chen Z, Yang X, Li L, Jiang C, Cheng L, Zhang B, Mo L, Zhou J. Enhancing sequential recommenders with augmented knowledge from aligned large language models. In: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2024, 345–354
- [159] Hu J, Xia W, Zhang X, Fu C, Wu W, Huan Z, Li A, Tang Z, Zhou J. Enhancing sequential recommendation via LLM-based semantic embedding learning. In: Companion Proceedings of ACM Web Conference 2024. 2024, 103–111
- [160] Wang B, Liu F, Zhang C, Chen J, Wu Y, Zhou S, Lou X, Wang J, Feng Y, Chen C, Wang C. LLM4DSR: leveraging large language model for denoising sequential recommendation, 2024. arXiv preprint arXiv: 2408.08208
- [161] Xu W, Wu Q, Liang Z, Han J, Ning X, Shi Y, Lin W, Zhang Y. SLMRec: distilling large language models into small for sequential recommendation, 2024. arXiv preprint arXiv: 2405.17890
- [162] Liu Q, Wu X, Wang W, Wang Y, Zhu Y, Zhao X, Tian F, Zheng Y. LLMEmb: large language model can be a good embedding generator for sequential recommendation, 2024. arXiv preprint arXiv: 2409.19925
- [163] Zheng B, Lin Z, Liu E, Yang C, Bai E, Ling C, Zhao W X, Wen J R. A large language model enhanced sequential recommender for joint video and comment recommendation, 2024. arXiv preprint arXiv: 2403.13574
- [164] Liu Q, Wu X, Zhao X, Wang Y, Zhang Z, Tian F, Zheng Y. Large language models enhanced sequential recommendation for long-tail user and item, 2024. arXiv preprint arXiv: 2405.20646v1
- [165] Zhang X, Xu B, Wu Y, Zhong Y, Lin H, Ma F. FineRec: exploring fine-grained sequential recommendation. In: Proceedings of the 47th International ACM SIGIR Conference on Research and

Development in Information Retrieval. 2024, 1599–1608

- [166] Yang S, Ma W, Sun P, Ai Q, Liu Y, Cai M, Zhang M. Sequential recommendation with latent relations based on large language model. In: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2024, 335–344
- [167] Lin J, Shan R, Zhu C, Du K, Chen B, Quan S, Tang R, Yu Y, Zhang W. ReLLa: retrieval-enhanced large language models for lifelong sequential behavior comprehension in recommendation. In: Proceedings of ACM Web Conference 2024. 2024, 3497–3508
- [168] Zhu J, Lin J, Dai X, Chen B, Shan R, Zhu J, Tang R, Yu Y, Zhang W. Lifelong personalized low-rank adaptation of large language models for recommendation, 2024. arXiv preprint arXiv: 2408.03533
- [169] Geng B, Huan Z, Zhang X, He Y, Zhang L, Yuan F, Zhou J, Mo L. Breaking the length barrier: LLM-enhanced CTR prediction in long textual user behaviors. In: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2024, 2311–2315
- [170] Xin H, Sun Y, Wang C, Xiong H. LLMCDJR: enhancing cross-domain sequential recommendation with large language models. ACM Transactions on Information Systems, 2025, doi: [10.1145/3715099](https://doi.org/10.1145/3715099)
- [171] Qin J, Zhang W, Wu X, Jin J, Fang Y, Yu Y. User behavior retrieval for click-through rate prediction. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval. 2020, 2347–2356
- [172] Chen Q, Xu Y, Pei C, Lv S, Zhuang T, Ge J. Efficient long sequential user data modeling for click-through rate prediction, 2022. arXiv preprint arXiv: 2209.12212
- [173] Dong J, Sun F, Wu T, Wu X, Zhang W, Wang S. A hierarchical network with user memory matrix for long sequence recommendation. Wireless Communications and Mobile Computing, 2022, 2022(1): 5457044
- [174] Chang J, Zhang C, Fu Z, Zang X, Guan L, Lu J, Hui Y, Leng D, Niu Y, Song Y, Gai K. TWIN: two-stage interest network for lifelong user behavior modeling in CTR prediction at KuaiShou. In: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 2023, 3785–3794
- [175] Si Z, Guan L, Sun Z, Zang X, Lu J, Hui Y, Cao X, Yang Z, Zheng Y, Leng D, Zheng K, Zhang C, Niu Y, Song Y, Gai K. TWIN V2: scaling ultra-long user behavior sequence modeling for enhanced CTR prediction at KuaiShou. In: Proceedings of the 33rd ACM International Conference on Information and Knowledge Management. 2024, 4890–4897
- [176] Gan M, Xiao K. R-RNN: extracting user recent behavior sequence for click-through rate prediction. IEEE Access, 2019, 7: 111767–111777
- [177] Lin Q, Zhou W J, Wang Y, Da Q, Chen Q G, Wang B. Sparse attentive memory network for click-through rate prediction with long sequences. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management. 2022, 3312–3321
- [178] Pi Q, Bian W, Zhou G, Zhu X, Gai K. Practice on long sequential user behavior modeling for click-through rate prediction. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. 2019, 2671–2679
- [179] Zhang Y, Chen E, Jin B, Wang H, Hou M, Huang W, Yu R. Clustering based behavior sampling with long sequential data for CTR prediction. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2022, 2195–2200
- [180] Feng Z, Xie J, Li K, Qin Y, Wang P, Li Q, Yin B, Li X, Lin W, Wang S. Context-based fast recommendation strategy for long user behavior sequence in Meituan Waimai. In: Companion Proceedings of ACM Web Conference 2024. 2024, 355–363
- [181] Hou R, Yang Z, Ming Y, Lu H, Zheng Z, Chen Y, Zeng Q, Chen M. Cross-domain lifelong sequential modeling for online click-through rate prediction. In: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 2024, 5116–5125
- [182] Yue Z, Wang Y, He Z, Zeng H, McAuley J, Wang D. Linear recurrent units for sequential recommendation. In: Proceedings of the 17th ACM International Conference on Web Search and Data Mining. 2024, 930–938
- [183] Liu L, Cai L, Zhang C, Zhao X, Gao J, Wang W, Lv Y, Fan W, Wang Y, He M, Liu Z, Li Q. LinRec: linear attention mechanism for long-term sequential recommender systems. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2023, 289–299
- [184] Guo T, Li X, Yang H, Liang X, Yuan Y, Hou J, Ke B, Zhang C, He J, Zhang S, Yu E, Ou W. Query-dominant user interest network for large-scale search ranking. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. 2023, 629–638
- [185] Cao Y, Zhou X, Feng J, Huang P, Xiao Y, Chen D, Chen S. Sampling is all you need on modeling long-term user behaviors for CTR prediction. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management. 2022, 2974–2983
- [186] Shen K, Ding X, Zheng Z, Gong Y, Li Q, Liu Z, Zhang G. SEMINAR: search enhanced multi-modal interest network and approximate retrieval for lifelong sequential recommendation, 2024. arXiv preprint arXiv: 2407.10714
- [187] Ren K, Qin J, Fang Y, Zhang W, Zheng L, Bian W, Zhou G, Xu J, Yu Y, Zhu X, Gai K. Lifelong sequential modeling with personalized memorization for user response prediction. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval. 2019, 565–574
- [188] Dang Y, Yang E, Guo G, Jiang L, Wang X, Xu X, Sun Q, Liu H. Uniform sequence better: time interval aware data augmentation for sequential recommendation. In: Proceedings of the 37th AAAI Conference on Artificial Intelligence. 2023, 4225–4232
- [189] Zhang S, Chen L, Wang C, Li S, Xiong H. Temporal graph contrastive learning for sequential recommendation. In: Proceedings of the 38th AAAI Conference on Artificial Intelligence. 2024, 9359–9367
- [190] Yin M, Wang H, Guo W, Liu Y, Zhang S, Zhao S, Lian D, Chen E. Dataset regeneration for sequential recommendation. In: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 2024, 3954–3965
- [191] Fan Z, Liu Z, Peng H, Yu P S. Mutual Wasserstein discrepancy minimization for sequential recommendation. In: Proceedings of ACM Web Conference 2023. 2023, 1375–1385
- [192] Xiao J, Pan W, Ming Z. A generic behavior-aware data augmentation framework for sequential recommendation. In: Proceedings

of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2024, 1578–1588

[193] Qiu R, Huang Z, Yin H, Wang Z. Contrastive learning for representation degeneration problem in sequential recommendation. In: Proceedings of the 15th ACM International Conference on Web Search and Data Mining. 2022, 813–823

[194] Han Y, Wang H, Wang K, Wu L, Li Z, Guo W, Liu Y, Lian D, Chen E. Efficient noise-decoupling for multi-behavior sequential recommendation. In: Proceedings of ACM Web Conference 2024. 2024, 3297–3306

[195] Liu Z, Fan Z, Wang Y, Yu P S. Augmenting sequential recommendation with pseudo-prior items via reversely pre-training transformer. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2021, 1608–1612

[196] Klenitskiy A, Volodkevich A, Pembek A, Vasilev A. Does it look sequential? An analysis of datasets for evaluation of sequential recommendations. In: Proceedings of the 18th ACM Conference on Recommender Systems. 2024, 1067–1072

[197] Sun A. On challenges of evaluating recommender systems in an offline setting. In: Proceedings of the 17th ACM Conference on Recommender Systems. 2023, 1284–1285



Li-Wei PAN received his BS degree from the Hubei University of Technology, China in 2023. He is currently working toward his PhD degree with the College of Computer Science and Software Engineering, Shenzhen University, China. His research interests include recommender systems, deep learning, and transfer learning.



Wei-Ke PAN received his PhD degree in computer science and engineering from the Hong Kong University of Science and Technology, China in 2012. He is currently a professor with the College of Computer Science and Software Engineering, Shenzhen University, China. His research interests include recommender systems, deep learning, transfer learning, and federated learning.



Mei-Yan WEI is currently pursuing a bachelor's degree at the Qingdao Institute of Software and the College of Computer Science and Technology, China University of Petroleum (East China). Her research interests include recommender systems and deep learning.



Hong-Zhi YIN received his PhD degree in computer science from Peking University, China in 2014. He is a full professor and director of the Responsible Big Data Intelligence Lab (RBDI), University of Queensland, Australia. He is an ARC future fellow and an ARC DECRA fellow. His research interests include recommendation system, user profiling, topic models, deep learning, social media mining, and location-based services.



Zhong MING received his PhD degree in computer science and technology from the Sun Yat-sen University, China in 2003. He is currently a professor with the College of Big Data and Internet, Shenzhen Technology University, China. His research interests include software engineering and artificial intelligence.