

A Survey on User Behavior Modeling in Recommender Systems

Zhicheng He¹, Weiwen Liu¹, Wei Guo², Jiarui Qin³, Yingxue Zhang⁴,
Yaochen Hu⁴ and Ruiming Tang¹

¹Huawei Noah's Ark Lab, Shenzhen, China

²Huawei Noah's Ark Lab, Singapore

³Shanghai Jiao Tong University, Shanghai, China

⁴Huawei Noah's Ark Lab, Montreal, Canada

{hezicheng9, liuweiwen8, guowei67, yingxue.zhang, yaochen.hu, tangruiming}@huawei.com,
qinjr96@sjtu.edu.com

Abstract

User Behavior Modeling (UBM) plays a critical role in user interest learning, which has been extensively used in recommender systems. Crucial interactive patterns between users and items have been exploited, which brings compelling improvements in many recommendation tasks. In this paper, we attempt to provide a thorough survey of this research topic. We start by reviewing the research background of UBM. Then, we provide a systematic taxonomy of existing UBM research works, which can be categorized into four different directions including *Conventional UBM*, *Long-Sequence UBM*, *Multi-Type UBM*, and *UBM with Side Information*. Within each direction, representative models and their strengths and weaknesses are comprehensively discussed. Besides, we elaborate on the industrial practices of UBM methods with the hope of providing insights into the application value of existing UBM solutions. Finally, we summarize the survey and discuss the future prospects of this field.

1 Background

With the rapid development of various Internet applications, Recommender Systems (RS) have become increasingly indispensable to both personalized services and the alleviation of information overloading [Zhang *et al.*, 2021]. One major bottleneck of RS research is the scarcity of explicit feedback regarding users' preferences [Hu *et al.*, 2008]. Instead, user preferences are implicitly recorded in the coarse and noisy behavior logs. To ameliorate this bottleneck, many researchers have been devoted to the study of User Behavior Modeling (UBM), aiming to explore and exploit user interest representations from the behavior history [Zhang *et al.*, 2021].

Despite the different forms of application tasks (e.g., ranking or next-item prediction), the common learning objective of recommendation models is to predict given users' interests in particular items. Specifically, for a UBM based recommendation model $F_{\Theta}^{UBM}(\cdot)$ parameterized by Θ , the interest

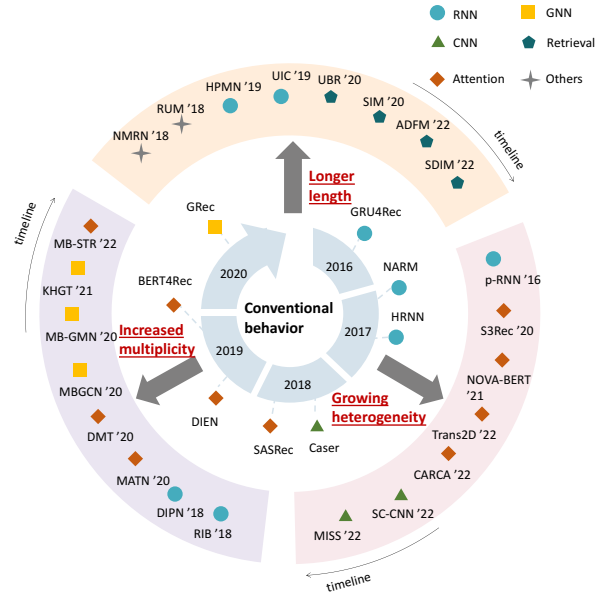


Figure 1: Taxonomy for user behavior modeling. The primary technique of each method is marked with a dyed shape.

$P(u, i)$ of a given user u for a target item i is predicted as

$$P(u, i) = F_{\Theta}^{UBM}(u, i, \mathcal{H}_u, f_u, f_i, f_c), \forall u \in \mathcal{U}, i \in \mathcal{I} \quad (1)$$

where \mathcal{U} and \mathcal{I} are the universal sets of users and items respectively, and \mathcal{H}_u is the user behavior history. f_u , f_i , and f_c are the non-behavior user features (e.g. age), item features (e.g. price), and context features (e.g., weekday) respectively. We do not discuss f_u , f_i , and f_c as they are not the key concerning issue of UBM, interested readers can refer to [Zhang *et al.*, 2021]. The core research issue of UBM is how to extract and model user interests from the behavior history \mathcal{H}_u . Within \mathcal{H}_u , each behavior record $h_{u,k} = \{v_k, t_k, f_k\}$ consists of the interacted item $v_k \in \mathcal{I}$, the time stamp t_k , and some related features f_k (e.g., behavior type and item description). As presented in Figure 1, various UBM methods have been proposed to learn effective user interest representations from \mathcal{H}_u . Despite there being already some

Branch	Representative	Research Motivation	Key Techniques
Conventional	GRU4Rec HRNN Caser SASRec	Learn evolving behavior patterns within session Integrate in- & cross-session behavior patterns Learn behavior patterns from different scopes and levels Learn adaptive weights for behavior dependencies	RNN RNN CNN Attention, Transformer
Long-sequence	UIC UBR	Memorize user interests for long histories (> 1000) Retrieve relevant behaviors from long histories (> 1000)	Memory Network Retrieval, Ranking
Multi-type	DMT MBGCN NMTR	Efficiently combine user interests from different behaviors Fuse fine-grained intra- & cross-type behavior semantics Jointly predict different types of user behaviors	Transformer, GCN Transformer, GCN Cascaded prediction
Side information	TiSASRec p-RNN S3Rec	Consider the time interval between two behaviors Consider the multi-modal text and image information Extract supervised signals from side information	Attention RNN SSL, Pre-training

GRU4Rec [Hidasi *et al.*, 2016a], HRNN [Quadrona *et al.*, 2017], Caser [Tang and Wang, 2018], UIC [Pi *et al.*, 2019], SASRec [Kang and McAuley, 2018], UBR [Qin *et al.*, 2020], DMT [Gu *et al.*, 2020], MBGCN [Jin *et al.*, 2020], NMTR [Gao *et al.*, 2019], TiSASRec [Li *et al.*, 2020], p-RNN [Hidasi *et al.*, 2016b], S3Rec [Zhou *et al.*, 2020].

Table 1: A taxonomy for existing UBM works according to the research issues they are proposed to solve.

survey papers on conventional UBM [Wang *et al.*, 2019; Fang *et al.*, 2020], however, they fail to cover the most recent development in user behavior modeling. In this survey, we will comprehensively analyze the latest UBM research progresses and discuss their advantages and weaknesses from both academic and industrial perspectives. Furthermore, we lay out three new research trends, i.e., the *longer length*, *increased multiplicity*, and *growing heterogeneity* of user behavior sequences, as illustrated in Figure 1.

The rest of this paper is organized as follows. In Section 2, we pick out some representative UBM works, categorize the proposed ideas, and analyze their differences and similarities. On top of that, we propose a new taxonomy of UBM which consists of four research directions as shown in Figure 1. The detailed discussion about the achievements, advantages and disadvantages of the four directions are presented in Section 3, 4, 5, and 6 respectively. In Section 7, we further discuss the application value of UBM in large-scale industrial recommender systems. Finally, Section 8 concludes this survey and sheds some light on the future research directions of UBM.

2 Taxonomy

Various UBM methods have been proposed to advance the exploration of implicit user interest modeling. To better understand the developing venation of UBM, we identify some representative and influential research works, analyze their research motivations, and summarize their key technical contributions. As presented in Table 1, this survey establishes a novel taxonomy for UBM that divides the existing research works into four major categories. Originating from the conventional UBM research on simple user behaviors, researchers have extended to further consider the *longer length*, *increased multiplicity*, and *growing heterogeneity* of user behavior histories, as illustrated in Figure 1. We briefly explain the core ideas of the three recent research trends as follows.

- **Conventional UBM** focuses on learning user interest representations from simple historical behavior sequences. The user behavior history \mathcal{H}_u is first processed into a

chronologically arranged item sequence with restricted length. Subsequently, researchers attempt to explore the complicated behavior patterns from various angles, such as the session structure [Hidasi *et al.*, 2016a], behavior granularity and influence scope [Tang and Wang, 2018], and pairwise dependency [Kang and McAuley, 2018].

- **Long-Sequence UBM** extends conventional UBM by learning from longer sequences of user behavior records (at least in the scale of thousands). As more behavior records in \mathcal{H}_u are kept, the model will have the potential to extract richer and long-term user interests. However, learning very long sequences with deep models is challenging and have potential computation bottlenecks. Thus researchers either adopt memory-augmented methods to store long-range behavior dependencies [Pi *et al.*, 2019] or take retrieval-based methods to pick out target-relevant behavior histories [Qin *et al.*, 2020].
- **Multi-Type UBM** further extends to study the multiplicity of user behaviors. In a unified recommendation model, the explicit modeling of different behavior types (e.g. click and purchase) provides a new perspective to understand behavior patterns and behavior correlations at a fine-grained granularity. Researchers have devoted themselves to a series of new challenges raised by the behavior multiplicity, such as the multi-behavior definition [Guo *et al.*, 2019], multi-behavior fusion [Jin *et al.*, 2020], and multi-behavior prediction [Gao *et al.*, 2019].
- **UBM with Side Information** further takes into account the heterogeneous features associated with behavior records. The side information helps to recover the contextual interactive logic when behavior records happen, which provides important supplements for understanding complicated user interests. Advanced NLP and CV models have been adopted to transform the rich multi-modal information into the user interest representation space [Liu *et al.*, 2021; Singer *et al.*, 2022].

The four research directions are closely related rather than mutually exclusive, and they enable the fast progression of the

UBM research field jointly. In the following sections, we further elaborate on the four UBM research directions in terms of the key research challenges, representative solutions, and developing trends.

3 Conventional UBM

Conventional user behavior modeling focuses on extracting item dependencies from relatively short-term behavior sequences with the same type [Kang and McAuley, 2018; Zhou *et al.*, 2018a]. Typical models can be formalized as:

$$P(u, i) = F_{\Theta}^{UBM}(u, i, \mathcal{H}_u^S), \forall u \in \mathcal{U}, i \in \mathcal{I}, \quad (2)$$

where \mathcal{H}_u^S is the behavior sequence extracted from \mathcal{H}_u . To both ensure the relevancy of extracted behavior records and to reduce model complexity, an upper bound value is used to restrict the length of the user behavior sequence \mathcal{H}_u^S . Various deep network structures have been adopted to learn from \mathcal{H}_u^S , including Recurrent Neural Networks (RNNs), Convolutional Neural Networks (CNNs), and Attention mechanisms.

3.1 RNN-based Methods

RNNs naturally capture the long- and short-term dependencies of a sequence. As one of the earliest attempts at user behavior modeling, GRU4Rec [Hidasi *et al.*, 2016a] uses Gated Recurrent Units (GRUs) to learn the evolving patterns for user behaviors within a session. Follow-up studies extend GRU4Rec with data augmentation techniques [Tan *et al.*, 2016] or improved ranking loss functions [Hidasi and Karatzoglou, 2018]. NARM [Li and et al, 2017] applies an RNN-based encoder-decoder model that combines both global and local interests in a session. Quadrona *et al.* [2017] further propose a hierarchy of two GRUs to model the evolving of user behaviors within and across sessions, respectively.

3.2 CNN-based Methods

Tang and Wang [2018] point out that RNN-based methods may not be able to capture the *skip behaviors* well, where the next step is influenced by the behaviors a few steps earlier rather than the adjacent behavior. Therefore, they propose a Case model that views recent behaviors as an “image” among time and latent dimensions, and learns the sequential patterns and skip behaviors with horizontal and vertical convolutional filters. NextItNet [Yuan *et al.*, 2019] introduces a generative CNN model with the residual block structure to capture both short- and long-range item dependencies.

3.3 Attention-based Methods

Recently, attention mechanisms are largely adopted in UBM because of their advantage in modeling interactions between any pair of behaviors, without degradation over the encoding distance [Vaswani *et al.*, 2017]. To exploit more supervised information, SASRec [Kang and McAuley, 2018] adopts self-attention to identify the importance of past behaviors in an auto-regressive prediction framework. DIN [Zhou *et al.*, 2018a] applies the attention mechanisms to adaptively learn the representation of user interests from historical behaviors with respect to a certain item. However, DIN ignores the evolving character of user interests, which is solved later

in DIEN [Zhou *et al.*, 2019] with the combination of attention and GRU. To further consider the intrinsic session structure of user behaviors, DSIN [Feng *et al.*, 2019] uses self-attention and Bi-LSTM to extract in-session interest representations and cross-session evolution patterns respectively. Despite the above achievements, researchers also notice the defects of conventional attention mechanisms in exploring complicated behavior patterns, thus resorting to more sophisticated attention architectures such as the Transformer encoder in BST [Chen *et al.*, 2019] and the stacked bidirectional Transformer encoder in BERT4Rec [Sun *et al.*, 2019].

3.4 Discussion

Research on conventional UBM explores different network structures for extracting dependency patterns from the simple behavior sequence. Despite the network structures mentioned above, other deep network structures (e.g., MLP [Wang *et al.*, 2014] and GNN [Chang *et al.*, 2021]) or a combination of multiple network structures [Zhou *et al.*, 2019; Feng *et al.*, 2019] have also been investigated in conventional UBM. For example, SURGE [Chang *et al.*, 2021] constructs item-item interest graphs from behavior sequences using metric learning. Different deep techniques are good at capturing different behavior dependencies. We observe that the captured dependencies in UBM are becoming more complex and practical. The dependency modeling progresses from simple unidirectional dependencies (by RNNs) to skip behavior dependencies (by CNNs), and finally to arbitrary pairwise dependency modeling (Attention) or multiple relationships modeling (GNNs).

4 Long-Sequence UBM

With the ever-growing volume of data from user-platform interactions, users accumulate a large amount of behavior data [Ren *et al.*, 2019]. Extending from the conventional UBM, long-sequence UBM further attempts to reserve a longer range of user behavior histories, which can be defined as

$$P(u, i) = F_{\Theta}^{UBM}(u, i, \mathcal{H}_u^L), \forall u \in \mathcal{U}, i \in \mathcal{I}, \quad (3)$$

where \mathcal{H}_u^L reserves as many behavior records in \mathcal{H}_u as possible. Long-sequence UBM enables recommender systems to further take advantage of long-term behavior dependencies and the periodicity of user behaviors. However, the user behavior sequences could become very long (at least in thousands), which might be prohibitively expensive to model all the historical user behavior. Longer sequences usually incline to contain more noise [Pi *et al.*, 2020; Qin *et al.*, 2020], and the strict online latency requirements further discourage incorporating long behavior sequences in large-scale industrial production. To ameliorate the above problems, many research works are proposed from both academia and industry, which can be divided into two major directions: memory-augmented methods and user behavior retrieval-based methods.

4.1 Memory-Augmented Methods

Memory-augmented networks are originally proposed in Natural Language Processing (NLP) tasks which explicitly store the extracted knowledge into external memory [Graves *et al.*,

2014; Kumar *et al.*, 2016]. In recommendation scenarios, the external memory stores the user interest representations, which are read and updated by a tailored neural network according to the user behavior records.

NMRN [Wang *et al.*, 2018] maintains an external memory bank of several interest representations for each user, which are updated according to the newly generated user behaviors. The interest representations are associated with personalized weights to reflect users’ different tastes. RUM [Chen *et al.*, 2018] stores the users’ historical behaviors explicitly. It uses a first-in-first-out mechanism to maintain the latest interacted items in the user memory. KSR [Huang *et al.*, 2018] takes the item attributes as keys to index the extracted knowledge from the external memory, which enhances the item representations for sequential recommendations.

To handle the long-sequence user behaviors more efficiently and effectively, specifically designed memory-augmented networks are proposed to summarize long behavior sequences into dense vectors. Ren *et al.* [2019] propose the lifelong sequential modeling framework with the Hierarchical Periodic Memory Network (HPMN). To capture the multi-scale sequential patterns, a multi-layer GRU architecture is designed to incrementally update external user interest representations with different update frequencies at each layer. With similar motivations, Pi *et al.* [2019] design a User Interest Center (UIC) module to store the multi-faceted user interests captured by the Multi-channel user Interest Memory Network (MIMN). At the inference time, user behavior representations are directly obtained from UIC without online calculation, which decouples the time-consuming modeling procedure from the real-time prediction process. UIC provides a systematic solution for long-sequence UBM from the industrial perspective and it has successfully deployed in a real-world recommender system.

4.2 User Behavior Retrieval Methods

Apart from storing user interests in an external memory, retrieving the most relevant and important behaviors from a very long sequence is an alternative solution. The retrieval process could be conducted efficiently and the retrieval step can reduce the potential noise inside the sequence. UBR [Qin *et al.*, 2020] uses search engine techniques to retrieve the most relevant behaviors with the target item. Instead of using a long consecutive behavior sequence, only a small retrieved set will be used in the following computation. Qin *et al.* [2023] further study how to optimize the retrieval function and propose a learning-to-rank-based optimization method. SIM [Pi *et al.*, 2020] proposes hard search and soft search approaches. For the hard search, it uses the user ID and item category to build the two-layer index. As for the soft search, SIM utilizes local sensitive hashing (LSH) to quickly fetch the relevant behaviors based on their embeddings. The above two methods are both two-phase models which means the retrieval function is trained separately from the prediction module. ETA [Chen *et al.*, 2021b], on the contrary, proposes to train the retrieval function in an end-to-end manner. It uses the SimHash algorithm to map the behavior embeddings and the target item to binary signatures. And based on the binary signatures, the behaviors with the smallest Hamming

distance to the target item are retrieved. Item embeddings are the only parameters used in the retrieval process, thus the entire process could be implemented end-to-end. SDIM [Cao *et al.*, 2022] further improves ETA by proposing a simple hash sampling-based approach that directly gathers behavior items that share the same hash signature with the target item. ADFM [Li *et al.*, 2022] proposes an adversarial filtering mechanism, which compresses the retrieved behavior sequences and removes the potential noisy behavior.

4.3 Discussion

The design philosophy of memory-augmented methods is to memorize more information by introducing more user-/item-specific parameters (i.e., external memory). Though effective in modeling longer sequences, the model is usually complicated and difficult to be deployed in real-world systems [Pi *et al.*, 2020]. By comparison, retrieval-based frameworks are more efficient and easy to scale for large systems. These advantages make them more suitable for handling long behavior sequences in industrial applications. However, retrieving behaviors inevitably causes information loss as many behaviors are dropped in the retrieval process. Therefore, it still remains an open problem regarding how to make a trade-off and exploit the advantages from both lines of research.

5 Multi-Type UBM

Building on top of conventional UBM, multi-type UBM aims to explicitly consider the different behavior types as they convey subtle differences in user interest modeling [Zhou *et al.*, 2018b; Jin *et al.*, 2020]. For example, in the e-commerce scenario, the *purchase* behavior usually stands for stronger interests than the *click* behavior. While in product reviewing scenarios, the *positive* and *negative* ratings are opposite. Specifically, multi-type UBM frameworks estimate the type-specific interest $P(u, i, b)$ of user u in item i as

$$P(u, i, b) = F_{\Theta}^{UBM}(u, i, \mathcal{H}_u^{MB}), \forall u \in \mathcal{U}, i \in \mathcal{I}, \quad (4)$$

where b is a specific behavior type in the whole type set \mathcal{B} . The multi-type behavior history \mathcal{H}_u^{MB} extends \mathcal{H}_u by explicitly considering the behavior type in each behavior record,

$$h_{u,k}^{MB} = \{v_k, t_k, b_k, f_k\}, b_k \in \mathcal{B}, \quad (5)$$

where b_k is particularly picked out from the behavior feature f_k to deal with the behavior multiplicity. The explicit consideration of behavior types not only enables to interpret user interests from a new perspective but also poses new challenges.

5.1 Behavior Type Definition

The first issue of multi-type UBM is the definition of behavior types, which is still an open question as the recommendation scenarios can vary drastically from each other. Researchers usually define different behavior types through expert analysis, which can be roughly divided into three categories.

- **Macro behaviors** refer to interactive behaviors with explicit motivations, which are defined by the logical design of recommendation scenarios. For example, the *click*, *add-to-cart*, *add-to-favorite*, and *buy* behaviors together implement the basic functions of e-commerce services [Xia *et*

al., 2021a; Xia *et al.*, 2020]. Users conduct a macro behavior to fulfill a distinctive purpose, which is well recorded (including the behavior type) and can be directly used for subsequent multi-type UBM methods.

- **Micro behaviors** are extracted from macro behaviors by experts [Zhou *et al.*, 2018b; Meng *et al.*, 2020], which provides a fine-grained understanding of user behaviors. For example, the *click* behavior can be divided into different micro behaviors according to the click source, which helps to explain users’ interests in specific items [Zhou *et al.*, 2018b]. However, definitions of micro behaviors are non-trivial and require sophisticated knowledge about both the particular scenarios and the targeted applications.
- **Behaviors from different domains or scenarios** can also be incorporated to provide complementary information for user interest mining [Gu *et al.*, 2021; Ma *et al.*, 2019]. For example, the ZEUS model combines the *query* behavior in a search scenario with the *click* behavior in a recommendation scenario to mine users’ implicit feedback [Gu *et al.*, 2021]. Besides, the same *watch* behavior in two domains is regarded as different behaviors to highlight the domain distinctions [Ma *et al.*, 2019]. Depending on the applications, expert guidance is needed to select complementary behaviors from related domains or scenarios.

5.2 Multi-Behavior Fusion

After the behavior types are determined, the next challenge of multi-type UBM is how to further model the complicated cross-type behavior dependencies on top of conventional UBM models. Existing researches fall into two groups according to how and when the cross-type behavior relations are fused with the intra-type behavior relations, i.e., the early-fusion models and the late-fusion models.

- **Late-fusion multi-type UBM** explores the intra-type and cross-type behavior relations under a two-step model architecture [Gao *et al.*, 2019; Gu *et al.*, 2020; Gu *et al.*, 2021; Chen *et al.*, 2021a]. For example, the NMTR model first separately predicts the user-item interactions concerning different behavior types, then assembles them in a cascading manner to account for the cross-type behavior relations [Gao *et al.*, 2019]. DMT and ZEUS propose to learn cross-type user representations based on the concatenation of intra-type sequential representations [Gu *et al.*, 2020; Gu *et al.*, 2021]. GHCF conducts multi-behavior predictions based on the behavior-wise fusion of intra-type collaborative representations [Chen *et al.*, 2021a]. Conventional UBM techniques (e.g., RNN, Transformer, and GCN) are readily usable for the intra-type learning process due to the homogeneity of behavior records. Besides, the independent intra-type learning processes are parallelizable, which brings fast computational efficiency. However, the common limitation of late-fusion models is that they ignore the item-level cross-type behavior modeling, which may negatively impact the performance.
- **Early-fusion multi-type UBM** learns both intra-type and cross-type behavior relations jointly in a hybrid manner [Zhou *et al.*, 2018b; Yuan *et al.*, 2022; Guo *et al.*, 2021;

Wei *et al.*, 2022]. For example, in RIB [Zhou *et al.*, 2018b], MB-STR [Yuan *et al.*, 2022], and π -Net [Ma *et al.*, 2019] models, the multi-behavior sequential patterns are learned from the multi-type hybrid behavior sequences. While in MBGCN [Jin *et al.*, 2020], MB-GMN [Xia *et al.*, 2021b], and MATN [Xia *et al.*, 2020], researchers attempt to explore the multiplex user-item interactive semantics with (heterogeneous) graph learning techniques. Building on top of conventional UBM, various modifications are proposed to incorporate behavior-aware information to deal with the item-level behavior differences. On one hand, the performances are improved by the enumerated exploration of intra-type and cross-type behavior dependencies at the item level. On the other hand, the computation complexity also increases with the sophisticated model architecture.

5.3 Multi-Behavior Prediction

Another unique problem of multi-type UBM is that sometimes multiple types of behaviors are required to predict in the same model [Gao *et al.*, 2019; Gu *et al.*, 2020; Gu *et al.*, 2021; Guo *et al.*, 2021]. Joint prediction of different types of behaviors is challenging because the label distributions of different behaviors are not aligned in the same space or even mutually exclusive. A common practice is to construct separated prediction modules for different behavior types, which is widely applicable regardless of the model paradigms [Guo *et al.*, 2019; Guo *et al.*, 2021; Chen *et al.*, 2021a]. However, the separated prediction head neglects the potential task relations, which leads to sub-optimal performances. To avoid the negative transfer across different behavior prediction tasks, the MMoE [Ma *et al.*, 2018] and PLE [Tang *et al.*, 2020] methods are employed to promote task relevance and suppress task conflicts [Gu *et al.*, 2020; Gu *et al.*, 2021; Yuan *et al.*, 2022]. When rigorous behavior dependencies are given by domain experts, the cascaded prediction structure can be applied [Gao *et al.*, 2019].

5.4 Discussion

The above three issues are closely related in multi-type UBM. For example, the ZEUS model utilizes behavior types from different scenarios, which causes different label distributions and demands for sophisticated multi-behavior prediction modules [Gu *et al.*, 2021]. DIPN simultaneously models macro and micro behaviors that express user interests at different levels, causing the early-fusion design incompetent to fuse them well [Guo *et al.*, 2019]. Therefore, multi-type UBM solutions should make a trade-off between the above three key designs, tailoring for the specific applications.

6 UBM with Side Information

Despite the great successes, most of the above-mentioned UBM methods overlook or underuse the rich side information associated with each behavior record. To fill the gap, UBM with side information aims to design specialized components to exploit the rich knowledge contained in f_k ,

$$P(u, i) = F_{\Theta}^{UBM}(u, i, M(\mathcal{H}_u)), \forall u \in \mathcal{U}, i \in \mathcal{I}, \quad (6)$$

where $M(\cdot)$ stands for the new side information fusing component. When dealing with a behavior record $h_{u,k} =$

$\{v_k, t_k, f_k\}$, $M(\cdot)$ fuses the side information (t_k and f_k) with item ID (v_k) to obtain a fine-grained representation. Apparently, the design of $M(\cdot)$ relies on the side information.

6.1 Side Information Sources

Various sources of side information can be utilized in UBM, which can be mainly classified into three categories, i.e., time information, item attributes, and multi-modal information. In conventional UBM methods, such as SASRec [Kang and McAuley, 2018] and BERT4REC [Sun *et al.*, 2019], the time information is used to sort the behavior records, which only influences the position encoding before sequence modeling. However, TiSASRec [Li *et al.*, 2020] finds that the time intervals between different item pairs convey crucial knowledge, thus proposing a novel time interval aware mechanism for the attentive weight calculation. With a similar motive, TISSA [Lei *et al.*, 2019] firstly proposes a time interval-based GRU to obtain session-level behavior representations, then splits them into slices with multi-scale time windows to better capture temporal dependencies. The item attributes also influence user behaviors significantly, providing necessary context information. FDSA [Zhang *et al.*, 2019] proposes to combine the item ID with attributes like category, brand, and description text for sequential recommendations. To track the changes of attributes over time, Trans2D [Singer *et al.*, 2022] performs feature transformation over item ID and attributes to learn complex item-attribute patterns. Different from the above methods, the use of multi-modal side information is more complicated. p-RNN [Hidasi *et al.*, 2016b] first extracts image features from video thumbnails and text features from product descriptions. Then it adopts existing CV and NLP approaches (GoogLeNet and bag-of-words) to learn multi-modal feature representations separately. Inspired by the achievements of multi-modal pre-training, SEMI [Lei *et al.*, 2021] obtains video and text representations through the direct use of pre-trained CV and NLP SOTAs.

6.2 Side Information Utilization

Once obtaining the side information representations, the effective use of them is another key issue. Early works usually merge different representation vectors with the simple addition or concatenation operations, then feed the mixed item vector to subsequent learning modules. For example, p-RNN [Hidasi *et al.*, 2016b] proposes to fuse the item ID information and side information at the input level with a concatenated feature or at the output level with a weighted summation output. RNN units are applied for sequence processing. SC-CNN [Zhang *et al.*, 2022] regards side information as different views, concatenates them as a 3D cube, then uses a semi-causal CNN to simultaneously capture the relations of different views. Trans2D [Singer *et al.*, 2022] also transforms items with attributes as a 3D cube, but applies a modified Transformer with 2D self-attention to handle the 3D data. CARCA [Rashed *et al.*, 2022] further extends the self-attention network with a two-branch multi-head self-attention-based framework to capture the dynamic user preferences hidden in users' context and attribute-aware profiles. The left branch is used to extract the dependencies among the rich historical behaviors and profiles, and the right branch is

used to capture the influence of the behavior sequence on the target item. However, fusing side information with simple fusion operators may impact the original item ID representations. As a result, NOVA-BERT [Liu *et al.*, 2021] leverages side information as an auxiliary for the self-attention module to learn better attention distribution, instead of being fused into item representations. DIF-SR [Xie *et al.*, 2022] argues that integrating side information before attention calculation will limit the learning of attention matrices thus decouples various side information with separate attention calculations to further improve NOVA-BERT. Inspired by the success of Self-Supervised Learning (SSL) in CV and NLP fields, researchers also adopt SSL for side information modeling. To mine the item-attribute relations, S3Rec [Zhou *et al.*, 2020] proposes two SSL objectives in pre-training, i.e., associated attribute prediction and masked attribute prediction. To further consider item attributes at the interest level, MISS [Guo *et al.*, 2022] proposes a CNN-based extractor to capture interest representations, dependencies, and correlations.

6.3 Discussion

It can be observed that the heterogeneous side information play a decisive role in the subsequent extraction and fusion processes. The simple time and attribute information can be seamlessly integrated into behavior modeling, while the complex image and text information needs to be processed with existing techniques in an ad-hoc manner. Thus, there is a lot of room to explore how to efficiently and effectively integrate the side information into the interest representation space.

7 Industrial Practices

The research of recommender systems is highly application-oriented, which is challenged by various practical problems. Therefore, pure research-oriented design is not sufficient enough to reflect the potential values of UBM research. In this section, we discuss the industrial practices of representative UBM methods. To ensure data authority and accuracy, online performance statistics are collected from published papers only. The results are presented in Table 2, from which we have the following observations.

- Besides the frequently mentioned online advertising and e-commerce, UBM methods are also applicable in other scenarios, such as APP store and coupon allocation, which reflects the potential research and application value of UBM.
- The consideration of long-range behaviors, multi-type behaviors, and side information all achieve improvements over conventional UBM baselines (e.g., UIC v.s. DIEN, DMT v.s. DIEN, SEMI v.s. BST), which demonstrates the value of them. However, if we consider the successfully deployed models, long-sequence UBM takes the lead.
- Computational efficiency is another key concern in UBM industrial practices. Despite that BST outperforms WDL and DIN in terms of CTR, the Transformer unit also causes a higher response time (RT), which limits its application to high throughput scenarios. When it comes to long behavior sequences, hashing (ETA), sampling (SDIM), and denoising (ADFM) methods demonstrate their superiority in balancing between performances and costs.

Model	Application Scenario	Baseline	Gains & Costs
DIN	Online advertising	Embedding&MLP	+10.0% CTR, +3.8% RPM
DIEN	Online advertising	Embedding&MLP DIN	+20.7% CTR, +17.1% eCPM, -3.0% PPC +11.8% CTR, +10.4% eCPM, -1.0% PPC
GRU4Rec+	Online video	Strategy	+5% Watch time, +5% Video play, +4% Click
BST	E-commerce	DIN	+3.02% CTR, +4ms RT
(a) Conventional UBM			
UIC	Online advertising	DIEN	+7.5% CTR, +6% RPM
SIM	Online advertising	UIC DIEN	+7.1% CTR, +4.4% RPM, +53 times MSL +2.1 $d_{category}$
UBR	App store	w/o UBR	+6.6% eCPM, +11.1% CTR
ETA	E-commerce	SIM	+1.8% CTR, +3.1% GMV, -2ms IT
SDIM	Online search	w/o Long sequence	+2.98% CTR, +2.69% VBR, +1ms IT
ADFM	Online advertising	SIM	+4.7% CTR, +3.1% RPM, -70.8% Storage
(b) Long-sequence UBM			
DMT	E-commerce	DIEN	+4.5% CTR, +4.6% CVR, +6.0% GMV
ZEUS	E-commerce	DMT	+6.0% CTR, +9.7% CVR, 11.7% GMV
DIPN	Coupon allocation	Strategy	+41.1% Usage Rate, +39.8% GMV
(c) Multi-type UBM			
NOVA-BERT	App store	BERT	+0.192 $\times 10^9$ FLOPs, +7.1 Mb Model Size
SEMI	E-commerce	BST	+9.32% NBV, +10.45% DT, +12.10% CWR
TiSSA	E-commerce	w/o TiSSA	+1.56% CTR, +2.09% CVR, +3.66% GMV
(d) UBM with side information			
CTR: Click-Through Rate; RPM: Revenue Per Mille; eCPM: effective Cost Per Mille; PPC: Pay Per Click; RT: Response Time; MSL: Max Sequence Length; $d_{category}$: Days till Last Same Category Behavior; IT: Inference Time; VBR: Visited Buy Rate; CVR: ConVersion Rate; GMV: Gross Merchandise Value; NBV: Number of Browsing Videos; DT: Dwell Time; CWR: Complete-Watch Ratio.			
(e) Explanation of evaluation metrics			

Table 2: Industrial online deployments of UBM solutions claimed in published papers.

- Despite the popularity of GNN techniques in academic research (e.g., SURGE, MBGCN, MB-GMN, etc.), none of them have been deployed online. The process of graph data requires a large amount of computational resources, which causes a heavy burden for the online environment.

8 Summary and Future Prospects

This survey summarizes the recent advances in user behavior modeling. Rich user preferences can be discovered from the behavior logs. With the help of RNN, CNN, and Attention-based techniques, conventional UBM methods are capable of finding the implicit feedback from the behavior sequences, thus largely improving the performances. Further improvements are obtained by taking into account the long-term behavior histories, the multi-type behaviors, and the side information accompanied with user behaviors. Despite the appreciable achievements, however, UBM research still faces challenges from both academic and industrial perspectives.

- **Deeper information fusion.** Despite the existing explorations on long behavior sequences, multi-type behaviors, and side information, the combination of them has seldom been considered. For example, modeling long multi-type behavior sequences might lead to new improvements,

but also demands novel retrieval or sampling techniques. Therefore, there still exists much exploration space for the wider and deeper fusion of various useful information.

- **More efficient learning method.** As application-oriented research, there exists a trade-off between effectiveness and efficiency for online serving. Along with the growing behavior length, types, and kinds of side information, the computation and storage burdens of UBM solutions also increase significantly. With the purpose of maintaining compelling performances, it is always important to seek UBM solutions with a light computation burden.
- **More interpretable user representations.** Despite the achieved improvements, the learned user interest representations are not well interpretable, which restricts the use case of UBM to other applications, such as user profiling and causal analysis. Thus interpretable UBM is another promising future research direction.
- **More advanced techniques.** Advanced deep learning techniques like pre-training and big models have been academically explored for side information learning. However, there is still a big gap before their industrial deployments.

Contribution Statement

Zhicheng He and Weiwen Liu contribute equally to this work. Ruiming Tang is the corresponding author.

References

- [Cao *et al.*, 2022] Yue Cao, Xiaojiang Zhou, and Jiaqi Feng et al. Sampling is all you need on modeling long-term user behaviors for CTR prediction. In *CIKM*, 2022.
- [Chang *et al.*, 2021] Jianxin Chang, Chen Gao, and Yu Zheng et al. Sequential recommendation with graph neural networks. In *SIGIR*, 2021.
- [Chen *et al.*, 2018] Xu Chen, Hongteng Xu, and Yongfeng Zhang et al. Sequential recommendation with user memory networks. In *WSDM*, 2018.
- [Chen *et al.*, 2019] Qiwei Chen, Huan Zhao, and Wei Li et al. Behavior sequence transformer for e-commerce recommendation in alibaba. In *DLP-KDD*, 2019.
- [Chen *et al.*, 2021a] Chong Chen, Weizhi Ma, and Min Zhang et al. Graph heterogeneous multi-relational recommendation. In *AAAI*, 2021.
- [Chen *et al.*, 2021b] Qiwei Chen, Changhua Pei, and Chao Li et al. End-to-end user behavior retrieval in click-through rate prediction model. *CoRR*, 2021.
- [Fang *et al.*, 2020] Hui Fang, Danning Zhang, and Yiheng Shu et al. Deep learning for sequential recommendation: Algorithms, influential factors, and evaluations. *ACM Trans. Inf. Syst.*, 2020.
- [Feng *et al.*, 2019] Yufei Feng, Fuyu Lv, and Weichen Shen et al. Deep session interest network for click-through rate prediction. In *IJCAI*, 2019.
- [Gao *et al.*, 2019] Chen Gao, Xiangnan He, and Dahua Gan et al. Neural multi-task recommendation from multi-behavior data. In *ICDE*, 2019.
- [Graves *et al.*, 2014] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural Turing machines. *CoRR*, 2014.
- [Gu *et al.*, 2020] Yulong Gu, Zhuoye Ding, and Shuaiqiang Wang et al. Deep multifaceted transformers for multi-objective ranking in large-scale e-commerce recommender systems. In *CIKM*, 2020.
- [Gu *et al.*, 2021] Yulong Gu, Wentian Bao, and Dan Ou et al. Self-supervised learning on users' spontaneous behaviors for multi-scenario ranking in e-commerce. In *CIKM*, 2021.
- [Guo *et al.*, 2019] Long Guo, Lifeng Hua, and Rongfei Jia et al. Buying or browsing?: Predicting real-time purchasing intent using attention-based deep network with multiple behavior. In *KDD*, 2019.
- [Guo *et al.*, 2021] Lei Guo, Li Tang, and Tong Chen et al. DA-GCN: A domain-aware attentive graph convolution network for shared-account cross-domain sequential recommendation. In *IJCAI*, 2021.
- [Guo *et al.*, 2022] Wei Guo, Can Zhang, and Zhicheng He et al. MISS: multi-interest self-supervised learning framework for click-through rate prediction. In *ICDE*, 2022.
- [Hidasi and Karatzoglou, 2018] Balázs Hidasi and Alexandros Karatzoglou. Recurrent neural networks with top-k gains for session-based recommendations. In *CIKM*, 2018.
- [Hidasi *et al.*, 2016a] Balázs Hidasi, Alexandros Karatzoglou, and Linas Baltrunas et al. Session-based recommendations with recurrent neural networks. In *ICLR*, 2016.
- [Hidasi *et al.*, 2016b] Balázs Hidasi, Massimo Quadrana, and Alexandros Karatzoglou et al. Parallel recurrent neural network architectures for feature-rich session-based recommendations. In *RecSys*, 2016.
- [Hu *et al.*, 2008] Yifan Hu, Yehuda Koren, and Chris Volinsky. Collaborative filtering for implicit feedback datasets. In *ICDM*, 2008.
- [Huang *et al.*, 2018] Jin Huang, Wayne Xin Zhao, and Hongjian Dou et al. Improving sequential recommendation with knowledge-enhanced memory networks. In *SIGIR*, 2018.
- [Jin *et al.*, 2020] Bowen Jin, Chen Gao, and Xiangnan He et al. Multi-behavior recommendation with graph convolutional networks. In *SIGIR*, 2020.
- [Kang and McAuley, 2018] Wang-Cheng Kang and Julian J. McAuley. Self-attentive sequential recommendation. In *ICDM*, 2018.
- [Kumar *et al.*, 2016] Ankit Kumar, Ozan Irsoy, and Peter Ondruska et al. Ask me anything: Dynamic memory networks for natural language processing. In *ICML*, 2016.
- [Lei *et al.*, 2019] Chenyi Lei, Shouling Ji, and Zhao Li. Tissa: A time slice self-attention approach for modeling sequential user behaviors. In *WWW*, 2019.
- [Lei *et al.*, 2021] Chenyi Lei, Yong Liu, and Lingzi Zhang et al. Semi: a sequential multi-modal information transfer network for e-commerce micro-video recommendations. In *KDD*, 2021.
- [Li and et al, 2017] Jing Li and Pengjie Ren, et al. Neural attentive session-based recommendation. In *CIKM*, 2017.
- [Li *et al.*, 2020] Jiacheng Li, Yujie Wang, and Julian J. McAuley. Time interval aware self-attention for sequential recommendation. In *WSDM*, 2020.
- [Li *et al.*, 2022] Xiaochen Li, Jian Liang, and Xialong Liu et al. Adversarial filtering modeling on long-term user behavior sequences for click-through rate prediction. In *SIGIR*, 2022.
- [Liu *et al.*, 2021] Chang Liu, Xiaoguang Li, and Guohao Cai et al. Noninvasive self-attention for side information fusion in sequential recommendation. In *AAAI*, 2021.
- [Ma *et al.*, 2018] Jiaqi Ma, Zhe Zhao, and Xinyang Yi et al. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *KDD*, 2018.
- [Ma *et al.*, 2019] MUYANG Ma, Pengjie Ren, and Yujie Lin et al. π -net: A parallel information-sharing network for shared-account cross-domain sequential recommendations. In *SIGIR*, 2019.

- [Meng *et al.*, 2020] Wenjing Meng, Deqing Yang, and Yanghua Xiao. Incorporating user micro-behaviors and item knowledge into multi-task learning for session-based recommendation. In *SIGIR*, 2020.
- [Pi *et al.*, 2019] Qi Pi, Weijie Bian, and Guorui Zhou et al. Practice on long sequential user behavior modeling for click-through rate prediction. In *KDD*, 2019.
- [Pi *et al.*, 2020] Qi Pi, Guorui Zhou, and Yujing Zhang et al. Search-based user interest modeling with lifelong sequential behavior data for click-through rate prediction. In *CIKM*, 2020.
- [Qin *et al.*, 2020] Jiarui Qin, Weinan Zhang, and Xin Wu et al. User behavior retrieval for click-through rate prediction. In *SIGIR*, 2020.
- [Qin *et al.*, 2023] Jiarui Qin, Weinan Zhang, and Rong Su et al. Learning to retrieve user behaviors for click-through rate estimation. In *TOIS*, 2023.
- [Quadrana *et al.*, 2017] Massimo Quadrana, Alexandros Karatzoglou, and Balázs Hidasi et al. Personalizing session-based recommendations with hierarchical recurrent neural networks. In *RecSys*, 2017.
- [Rashed *et al.*, 2022] Ahmed Rashed, Shereen Elsayed, and Lars Schmidt-Thieme. Context and attribute-aware sequential recommendation via cross-attention. In *RecSys*, 2022.
- [Ren *et al.*, 2019] Kan Ren, Jiarui Qin, and Yuchen Fang et al. Lifelong sequential modeling with personalized memorization for user response prediction. In *SIGIR*, 2019.
- [Singer *et al.*, 2022] Uriel Singer, Haggai Roitman, and Yotam Eshel et al. Sequential modeling with multiple attributes for watchlist recommendation in e-commerce. In *WSDM*, 2022.
- [Sun *et al.*, 2019] Fei Sun, Jun Liu, and Jian Wu et al. Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer. In *CIKM*, 2019.
- [Tan *et al.*, 2016] Yong Kiam Tan, Xinxing Xu, and Yong Liu. Improved recurrent neural networks for session-based recommendations. In *DLRS@RecSys*, 2016.
- [Tang and Wang, 2018] Jiayi Tang and Ke Wang. Personalized top-n sequential recommendation via convolutional sequence embedding. In *WSDM*, 2018.
- [Tang *et al.*, 2020] Hongyan Tang, Junning Liu, and Ming Zhao et al. Progressive layered extraction (PLE): A novel multi-task learning (MTL) model for personalized recommendations. In *RecSys*, 2020.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, and Niki Parmar et al. Attention is all you need. In *NIPS*, 2017.
- [Wang *et al.*, 2014] Pengfei Wang, Jiafeng Guo, and Yanyan Lan et al. Learning hierarchical representation model for next basket recommendation. In *SIGIR*, 2014.
- [Wang *et al.*, 2018] Qinyong Wang, Hongzhi Yin, and Zhit-ing Hu et al. Neural memory streaming recommender networks with adversarial training. In *KDD*, 2018.
- [Wang *et al.*, 2019] Shoujin Wang, Liang Hu, and Yan Wang et al. Sequential recommender systems: Challenges, progress and prospects. In *IJCAI*, 2019.
- [Wei *et al.*, 2022] Wei Wei, Chao Huang, and Lianghao Xia et al. Contrastive meta learning with behavior multiplicity for recommendation. In *WSDM*, 2022.
- [Xia *et al.*, 2020] Lianghao Xia, Chao Huang, and Yong Xu et al. Multiplex behavioral relation learning for recommendation via memory augmented transformer network. In *SIGIR*, 2020.
- [Xia *et al.*, 2021a] Lianghao Xia, Chao Huang, and Yong Xu et al. Knowledge-enhanced hierarchical graph transformer network for multi-behavior recommendation. In *AAAI*, 2021.
- [Xia *et al.*, 2021b] Lianghao Xia, Yong Xu, and Chao Huang et al. Graph meta network for multi-behavior recommendation. In *SIGIR*, 2021.
- [Xie *et al.*, 2022] Yueqi Xie, Peilin Zhou, and Sunghun Kim. Decoupled side information fusion for sequential recommendation. In *SIGIR*, 2022.
- [Yuan *et al.*, 2019] Fajie Yuan, Alexandros Karatzoglou, and Ioannis Arapakis et al. A simple convolutional generative network for next item recommendation. In *WSDM*, 2019.
- [Yuan *et al.*, 2022] Enming Yuan, Wei Guo, and Zhicheng He et al. Multi-behavior sequential transformer recommender. In *SIGIR*, 2022.
- [Zhang *et al.*, 2019] Tingting Zhang, Pengpeng Zhao, and Yanchi Liu et al. Feature-level deeper self-attention network for sequential recommendation. In *IJCAI*, 2019.
- [Zhang *et al.*, 2021] Weinan Zhang, Jiarui Qin, and Wei Guo et al. Deep learning for click-through rate estimation. In *IJCAI*, 2021.
- [Zhang *et al.*, 2022] Xuanyu Zhang, Qing Yang, and Dongliang Xu et al. Deepvt: Deep view-temporal interaction network for news recommendation. In *CIKM*, 2022.
- [Zhou *et al.*, 2018a] Guorui Zhou, Xiaoqiang Zhu, and Chengru Song et al. Deep interest network for click-through rate prediction. In *KDD*, 2018.
- [Zhou *et al.*, 2018b] Meizi Zhou, Zhuoye Ding, and Jiliang Tang et al. Micro behaviors: A new perspective in e-commerce recommender systems. In *WSDM*, 2018.
- [Zhou *et al.*, 2019] Guorui Zhou, Na Mou, and Ying Fan et al. Deep interest evolution network for click-through rate prediction. In *AAAI*, 2019.
- [Zhou *et al.*, 2020] Kun Zhou, Hui Wang, and Wayne Xin Zhao et al. S3-rec: Self-supervised learning for sequential recommendation with mutual information maximization. In *CIKM*, 2020.